

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

Authored by
stats writer

December 28, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Run Fisher's Exact Test in Stata: A Step-by-Step Guide*.
PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=109525>

The Fisher's Exact Test is a powerful statistical method used within Stata to analyze the relationship between two categorical variables. This test is essential for researchers, particularly when dealing with small sample sizes where traditional methods fail. It operates by testing the null hypothesis that the two variables under examination are statistically independent. The core outcome of this procedure is a p-value, which provides the evidence needed to determine the significance of the association between the variables. Importantly, Fisher's Exact Test serves as the preferred alternative to the commonly used chi-squared test, especially when the assumption regarding sufficient expected cell frequencies is violated.

Understanding the Necessity of Fisher's Exact Test

Statistical inference often requires us to determine if an association exists between classification factors. When analyzing associations between two nominal variables summarized in a contingency table, the standard procedure is typically the Chi-Squared Test of Independence. However, this test relies on asymptotic theory and assumes that the sample size is large enough such that the majority of the expected cell counts are greater than five. When this critical assumption is breached--specifically, when one or more cells in a 2x2 table have an expected frequency less than five--the calculated Chi-Squared test statistic becomes unreliable, leading to inaccurate p-values.

The Fisher's Exact Test resolves this issue by calculating the exact probability of observing the data, or data more extreme, given the marginal totals are fixed. This calculation is derived from the hypergeometric distribution, making it precise even for very small samples. This makes it the authoritative choice for analyzing small, non-parametric datasets summarized in 2x2 contingency tables, ensuring the validity of the statistical conclusion regarding the independence of the two factors.

Prerequisites and Data Structure in Stata

To successfully implement Fisher's Exact Test in Stata, you must first ensure your data is prepared correctly. While Stata usually works with raw datasets, the command utilized for the exact test often requires inputting the data directly as a summarized table, especially when using the tabi command variant. This approach is highly efficient when the counts for the 2x2 matrix are already known, bypassing the need to define individual variables within the dataset.

The structure of the data must conform to a 2x2 contingency table, meaning we are comparing two binary categorical variables. For instance, if examining the relationship between Gender (Male/Female) and Outcome (Success/Failure), the counts representing the intersection of these categories must be clearly identified and ordered sequentially for input into the command line. Failure to maintain the correct row and column order during input will result in a misinterpretation of

the final statistical output. This reliance on fixed marginal totals is central to the exact calculation methodology, differentiating it significantly from approximation tests like the chi-squared test.

Illustrative Example: Political Preference and Gender

Consider a typical research scenario where we aim to investigate whether there is a statistically significant association between a student's gender and their political party preference (Democrat or Republican) at a specific university. To gather preliminary evidence, a random sample of 25 students was surveyed. The resulting counts, categorized by gender and preference, are meticulously summarized in the 2x2 contingency table presented below.

The primary goal is to formally test the independence of these two variables using robust methodology. If the resulting p-value from the Exact Test is sufficiently small (typically less than the chosen alpha level of 0.05), we can confidently reject the notion of independence and conclude that gender and political preference are significantly associated within this sampled population.

	Democrat	Republican
Male	4	9
Female	8	4

This specific data configuration yields a total count of 25 observations (N=25). Given the modest sample size and the presence of several small cell counts (4 in two cells), the assumption required for the expected cell frequencies necessary for the standard chi-squared test is likely violated. Therefore, applying the rigorous methods of Fisher's Exact Test is mandatory to ensure the reliability and validity of our statistical conclusion regarding the association between gender and political party preference.

Executing Fisher's Exact Test using Stata's `tabi` Command

In Stata, the most straightforward method for performing Fisher's Exact Test on summarized, immediate data is by utilizing the built-in `tabi` command. The `tabi` (tabulate immediate) command is designed to perform cross-tabulation when you enter the cell frequencies directly into the command line, eliminating the intermediate step of loading or generating a structured dataset. This is highly efficient when only the summary counts are readily available from external sources or preliminary analysis.

The required syntax demands that the four cell counts be entered sequentially: the first row's entries followed by the second row's entries. A critical separator, the backward slash (`()`), must be used to clearly delineate the end of the top row and the start of the bottom row data entry. This structured input allows Stata to accurately reconstruct the 2x2 contingency table necessary for the

statistical calculation.

Based on our example data (Row 1: 4, 9; Row 2: 8, 4), the specific command input into the Stata Command Window should appear precisely as follows, including the optional but recommended `exact` modifier to ensure the correct test is performed:

tabi 4 9 8 4, exact

Upon executing this command, Stata will generate output that includes the reconstructed contingency table, along with the results of the statistical tests. This output provides the raw counts, marginal totals, and, most importantly, the exact probabilities (p-values) required for hypothesis testing, as shown in the visualization below.

```
. 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**

Detailed Interpretation of Stata Output

The generated output provides essential statistical information necessary for drawing a valid conclusion regarding the association between the two variables. Understanding each component is critical for accurate reporting and decision-making in quantitative research.

Output Table (Cross-Tabulation): This introductory section serves as a verification of the input data. Stata displays the 2x2 table showing the counts for each cell (4, 9, 8, 4), along with the computed row totals, column totals, and the grand total (N=25). This confirms that the numerical input provided via the `tabi` command was correctly interpreted by the software.

Fisher's Exact (Two-Sided): This is the key measure for non-directional hypothesis testing. The two-sided p-value tests the general null hypothesis of independence against the alternative hypothesis that the variables are related, without specifying the direction of the relationship. In our example output, the two-sided p-value is reported as **0.115**. This value represents the exact probability of observing the current distribution of counts, or a more extreme distribution, assuming the two variables are truly independent.

One-Sided Fisher's Exact: This p-value, reported here as **0.081**, is only relevant when the research question is strictly directional. For instance, if the hypothesis was specifically that "Males are more likely to be Republicans than Females," we might consider the one-sided test. Since the standard test of independence uses a non-directional alternative hypothesis, the two-sided p-value is the default and appropriate measure for assessing the overall association.

We must emphasize the importance of selecting the correct p-value based on the research design. Since the objective was to determine **whether or not** an association exists, the two-sided test is the statistically correct choice. Reliance on the one-sided test without prior justification introduces potential bias into the analysis.

Formal Hypothesis Testing and Decision Criteria

The formal process of hypothesis testing hinges on defining the null and alternative hypotheses clearly. For the analysis performed using Fisher's Exact Test on our sample data:

Null Hypothesis (\$H_0\$): Gender and political party preference are statistically independent.

Alternative Hypothesis (\$H_a\$): Gender and political party preference are statistically dependent (associated).

We use the calculated two-sided p-value ($P = 0.115$) and compare it against the conventional significance level, alpha (α), set at 0.05. The decision rule dictates that we reject H_0 only if the P-value is less than or equal to 0.05. If $P > 0.05$, we fail to reject H_0 .

Conclusion: Interpreting the Non-Significant Result

Comparing the observed probability (P-value = 0.115) with the standard threshold ($\alpha = 0.05$), we find that 0.115 is greater than 0.05. Therefore, we must fail to reject the null hypothesis of independence.

This statistical outcome leads to the formal conclusion that, based on the sampled data of 25 students, there is insufficient statistical evidence to assert that a significant association exists between a student's gender and their political party preference. The observed difference in party affiliation proportions between males and females in the sample could reasonably have occurred simply due to random chance, even if the variables were truly independent in the larger population.

It is crucial to understand that failing to reject the null hypothesis does not equate to proving independence. Rather, it means the collected data does not provide the strong evidence (a P-value below 0.05) required to overturn the assumption of independence. This example perfectly illustrates why Stata users should employ Fisher's Exact Test rigorously in situations involving small expected cell frequencies, ensuring all statistical inferences are sound and justifiable, unlike

the unreliable results that might be generated by the standard chi-squared test in such constrained sample conditions.

ARABPSYCHOLOGY.COM