

Fisher's Exact Test in Stata: A Comprehensive Tutorial

Authored by
Mohammed looti

November 8, 2025

RECOMMENDED CITATION

Mohammed looti (2025). *Fisher's Exact Test in Stata: A Comprehensive Tutorial*.
PSYCHOLOGICAL STATISTICS. Retrieved from
<https://statistics.arabpsychology.com/?p=13630>

The Statistical Imperative: Why Choose Fisher's Exact Test?

The analysis of association between two nominal or **categorical variables** is a foundational exercise in statistics across diverse disciplines, including medical research, sociology, and marketing. When researchers seek to determine whether a statistically significant relationship exists between two such variables, the [Fisher's Exact Test](#) (FET) stands out as an exceptionally reliable tool. Unlike methods that rely on large-sample approximations, FET is predicated on calculating the precise probability of observing the data, or data more extreme, given fixed marginal totals. This exact approach grants it superior accuracy, particularly in challenging data environments.

This powerful test is specifically designed for analyzing data arranged in a 2x2 contingency table. Such tables organize observations based on two dichotomous characteristics, allowing for a structured investigation into their interdependence. The central utility of Fisher's test emerges when sample sizes are inherently small, or when the data distribution is sparse. In these scenarios, the test provides a safeguard against the inflated risk of Type I errors that plague approximate methods. By meticulously calculating the probability based on the known distribution of the data, the FET ensures that any statistical conclusion drawn is robust and mathematically valid, irrespective of the sample size constraints.

For statisticians and researchers utilizing powerful software like **Stata**, mastering the execution of the Fisher's Exact Test is essential for responsible data analysis. This comprehensive guide will navigate the theoretical rationale for choosing the exact test over approximations and provide a rigorous, step-by-step methodology for executing and interpreting the results within the [Stata](#) computing environment. By the conclusion of this tutorial, users will be equipped to confidently analyze small-sample contingency data, ensuring precise and defensible statistical inference.

The Critical Distinction: Fisher's Test vs. the Chi-Square Approximation

To truly appreciate the necessity of the exact test, one must first understand the limitations inherent in the ubiquitous [Chi-Square Test of Independence](#). The Chi-Square test operates under the crucial assumption that the sampling distribution of its test statistic closely approximates the theoretical chi-square distribution. This approximation holds true only when the sample size is sufficiently large, ensuring that the expected frequency within every cell of the contingency table meets a specific threshold. The conventional statistical rule dictates that if one or more expected cell counts fall below five, the Chi-Square approximation becomes unreliable.

When this critical assumption of minimum expected counts is violated, relying on the standard Chi-Square test introduces significant risk. The resulting p-value may be inaccurate, often leading to an inflated rate of false positives--a situation where the researcher incorrectly rejects the [null hypothesis](#) of independence. This failure to maintain statistical validity can undermine research

findings and lead to flawed scientific conclusions. Consequently, the presence of sparse data or small cohorts mandates a pivot to a more stringent analytical approach.

The [Fisher's Exact Test](#) resolves this issue by entirely bypassing the need for asymptotic approximations. Its theoretical foundation is rooted in the [Hypergeometric Distribution](#), which allows it to calculate the precise probability of observing the specific data configuration (and all configurations more extreme), assuming that both the row and column totals are fixed. Because it deals with discrete probabilities rather than continuous approximations, the resulting [p-value](#) is mathematically exact. This level of precision is invaluable in fields where gathering large sample sizes is impractical, expensive, or ethically challenging, such as clinical trials involving rare diseases or specialized social science surveys. The choice between the two tests is therefore not a matter of preference, but an absolute requirement dictated by the empirical characteristics of the collected data.

Setting Up the Analysis: A Practical Research Example

To demonstrate the practical application of Fisher's Exact Test, we will utilize a compelling example drawn from a small-scale social science study. Imagine a university researcher who wants to explore whether a student's gender is associated with their declared [political party preference](#). This investigation focuses on two key variables: **Gender** (a binary categorical variable: Male or Female) and **Party Preference** (another binary categorical variable: Democrat or Republican).

Given the constraints of time and resources, the researcher conducts a small, targeted survey, gathering responses from only 25 randomly selected students. This deliberately small sample size immediately raises a red flag regarding the use of the Chi-Square test, confirming the need for the exact methodology. The small N increases the likelihood of low cell counts in the resulting contingency table, which is the foundational structure for the analysis.

The data collected is aggregated into the standard 2x2 table format shown below. The objective of our analysis in [Stata](#) will be to determine if the pattern of counts observed in this table suggests a genuine statistical dependency between gender and political preference, or if the observed disparity is merely a product of chance variation under the assumption of independence. The small count of 4 Republican females, for instance, confirms that we must use the exact method to produce a reliable statistical measure of association.

	Democrat	Republican
Male	4	9
Female	8	4

This contingency table is the direct input required by **Stata**. The structure clearly defines the row variables (Male/Female) and the column variables (Democrat/Republican), with the cell entries representing the frequency of students falling into each combination. The next step involves translating this summarized data into a precise command for execution within the statistical software.

Executing the Test in Stata: The Power of the `tabi` Command

The statistical software [Stata](#) offers a remarkably efficient way to conduct the [Fisher's Exact Test](#) when the data is already aggregated into a summary table format. The primary command utilized for this purpose is **tabi**, which is an abbreviation for "table immediate." This command allows the researcher to input the cell frequencies directly into the command line, bypassing the need to create and load a full raw dataset file. This immediacy is extremely convenient for quick calculations or when replicating results from published tables.

The syntax for the [tabi command](#) requires careful sequencing of the cell counts. The cell frequencies must be entered row by row, moving from left to right across the table. Crucially, a backslash (`\`) is used as a delimiter to signify the end of one row and the beginning of the next. For our example, the first row contains the counts 4 and 9, and the second row contains 8 and 4.

To request the Fisher's Exact Test specifically, the **exact** option is appended to the command. Although Stata often defaults to the exact test when marginal totals are small, explicitly stating the option ensures the correct test is performed and documented in the command log. Following these guidelines, the complete command input into the Stata console for our data example is structured as follows:

tabi 4 9 8 4, exact

It is important to note the inclusion of the backslash separating the first row (4 9) from the second row (8 4). This structure clearly defines the 2x2 matrix for Stata. Upon execution, the software immediately processes the data and generates a detailed output that includes the reconstructed contingency table and the essential statistical metrics needed for inference. This ease of execution highlights why [Stata](#) remains a preferred tool for rigorous statistical analysis.

Deconstructing the Stata Output: Understanding the P-Values

The output generated by **Stata** following the execution of the **tabi** command is rich with information, but requires careful interpretation to extract the definitive conclusion regarding the association between the two [categorical variables](#). The output begins by displaying the contingency table exactly as entered, allowing for immediate verification, followed by the calculated measures, including the p-values from Fisher's Exact Test.

```
. 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**

The output specifically highlights three critical statistics necessary for drawing inferences about the [null hypothesis](#). These elements are presented clearly and must be understood in the context of the initial research question:

Verification Table: The top section of the output is a restatement of the 2x2 table (4, 9, 8, 4), along with the computed row and column totals and the overall sample size (\$N=25\$). This step is crucial for confirming that the correct data matrix was successfully processed by the **tabi** command.

Fisher's exact (Two-Sided): This is arguably the most frequently used statistic. The two-sided [p-value](#) evaluates the possibility of observing the current data distribution, or any distribution showing an equally or more extreme difference in proportions in **either direction**, assuming independence. For our example, the two-sided p-value is reported as **0.115**. This figure is used when the researcher does not have a specific, pre-determined direction for the hypothesized association.

1-sided Fisher's exact: This value represents the p-value for a one-tailed test. A one-sided test is only appropriate if the researcher had established a strong, directional hypothesis before data collection (e.g., hypothesizing specifically that males are statistically more likely to favor the Republican party than females). In the provided output, the one-sided p-value is **0.081**. Unless a directional hypothesis was explicitly stated and justified, the researcher must adhere to the two-sided result for conservative and unbiased inference.

Since our initial research question was broad--"Is there an association?"--we must rely exclusively on the two-sided p-value of 0.115. This figure quantifies the probability that the observed pattern of political preferences across genders occurred purely by random chance if, in the underlying population, gender and preference were entirely independent.

Statistical Inference: Drawing Robust and Valid Conclusions

The process of statistical inference begins with clearly articulating the hypotheses. For the [Fisher's Exact Test](#), the structure is as follows: The **Null Hypothesis** (\$H_0\$) posits that the two variables

are statistically independent. In our case, H_0 states: "Gender and [political party preference](#) are independent (i.e., there is no association) among students." The **Alternative Hypothesis** (H_a) asserts that a statistically significant association exists between the two variables.

To reach a decision, the calculated two-sided [p-value](#) (0.115) is compared against a predetermined significance level, denoted as alpha (α). The standard convention in social science and medical research sets α at 0.05. The decision rule is straightforward: if the p-value is less than α (0.05), we reject H_0 and conclude that an association exists. If the p-value is greater than or equal to α , we fail to reject H_0 .

In our specific scenario, the p-value of 0.115 is notably greater than the significance level of 0.05. Therefore, we **fail to reject the null hypothesis**. This crucial statistical outcome means that, despite the observed differences in the sample (where males appeared slightly more Republican and females slightly more Democrat), the small size of the sample dictates that this pattern is not strong enough to be deemed statistically significant. The probability of observing this distribution by random chance, even if the variables were truly independent, is 11.5%, which exceeds the acceptable threshold of 5%.

In conclusion, the analysis performed in **Stata** using Fisher's Exact Test indicates that there is insufficient statistical evidence, at the 5% significance level, to establish a genuine association between gender and [political party preference](#) among the students sampled. Researchers must report this result accurately, emphasizing that the conclusion is based on the exact probability calculation, which was necessary due to the low expected cell counts inherent in the small sample size.