

Fisher's Exact Test in Excel: A Practical Guide for Data Analysis

Authored by
Mohammed loot

November 8, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Fisher's Exact Test in Excel: A Practical Guide for Data Analysis*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=13475>

The [Fisher's Exact Test](#) stands as a cornerstone in statistical analysis, designed specifically to evaluate the existence of a non-random association between two [categorical variables](#). This powerful methodology becomes indispensable when researchers encounter data derived from [small sample sizes](#) or when the expected frequencies within a 2x2 [contingency table](#) fall below standard thresholds (typically five). In these specific instances, the traditional and often-used [Chi-Square Test of Independence](#) may yield statistically invalid or unreliable results because its fundamental assumptions regarding large sample size and asymptotic distribution are violated. Consequently, mastering the correct implementation of the Fisher's Exact Test, even outside of specialized statistical software, is a critical skill for conducting rigorous research and drawing accurate conclusions.

While many advanced statistical packages offer direct, single-function solutions for executing this test, countless analysts and researchers still rely heavily on Microsoft Excel for initial data management, manipulation, and analysis. This comprehensive guide provides a precise, step-by-step methodology for executing the [Fisher's Exact Test](#) entirely within the Excel environment. We achieve this by leveraging the platform's built-in function dedicated to the [Hypergeometric Distribution](#), which serves as the foundational mathematical model underpinning the entire test.

The Statistical Rationale for Fisher's Exact Test

The primary objective of the [Fisher's Exact Test](#) is to rigorously evaluate the [null hypothesis](#) (H_0), which posits that the observed frequencies within the 2x2 table are strictly proportional to their respective marginal totals, implying zero association or independence between the two variables under investigation. Unlike tests that rely on approximations, Fisher's method calculates the precise probability of observing the specific data set, or data that are considered more extreme, based on the crucial assumption that the marginal totals (the fixed row and column sums) remain constant.

This exact probability calculation is rooted deeply in the principles of combinatorial mathematics, specifically drawing upon the [Hypergeometric Distribution](#). This distribution mathematically models the likelihood of obtaining a specific number of "successes" (represented by the frequency in a cell) within a defined sample size (the row or column total). Critically, this sample is drawn without replacement from a finite population (the grand total), based on the known total number of successes within that population (the other marginal total). This framework ensures that the probability is calculated exactly, circumventing the limitations associated with asymptotic approximations.

The application of this test is not merely optional but mandatory across disciplines such as epidemiology, behavioral sciences, and biological research, particularly whenever outcome data is binary (e.g., pass/fail, disease present/absent, treatment received/not received) and the overall

sample size is low. Ignoring the constraints imposed by [small sample sizes](#) and relying incorrectly on the [Chi-Square Test of Independence](#) in these situations can lead directly to the miscalculation of the [P-value](#) and, consequently, result in scientifically flawed conclusions or erroneous policy decisions. Thus, selecting the appropriate statistical tool based on sample characteristics is paramount.

Setting Up the Analysis: A Practical Case Study

To clearly demonstrate the required computational steps within Excel, we will utilize a practical, illustrative example. Consider a scenario where an academic researcher seeks to determine if there is a statistically significant relationship between a student's gender and their declared political party preference at a specific university institution. To collect the necessary data, a randomized sample of 25 students is surveyed, and their responses are systematically organized into a standard 2x2 contingency table. The two variables being analyzed are Gender (categorized as Male/Female) and Political Preference (categorized as Democrat/Republican).

The resulting observed frequencies derived directly from this poll are structured precisely as shown in the table below:

	A	B	C	D	E	F
1		Democrat	Republican	Total		
2	Male	4	9	13		
3	Female	8	4	12		
4	Total	12	13	25		
5						
6						
7						
8						
9						
10						
11						
12						
13						
14						
15						
16						
17						

The core analytical objective is to determine whether this observed distribution of counts provides sufficient statistical evidence to confidently reject the [Null Hypothesis](#) (H_0). The null hypothesis, in this context, asserts that gender and political party preference are entirely independent

variables. Given that the total sample size is small ($N=25$) and several individual cell counts fall below the standard threshold of five, the [Fisher's Exact Test](#) is unequivocally the statistically correct and most appropriate method for conducting this specific analysis.

Implementing the Hypergeometric Distribution Function in Excel

A crucial point for Excel users is that Microsoft Excel does not natively include a dedicated function explicitly named "FISHER.TEST." Therefore, analysts must instead utilize the mathematically equivalent function for calculating probabilities based on the [Hypergeometric Distribution](#). This function is designated as **HYPGEOM.DIST**. This function calculates the probability of obtaining a specific number of successes (or a cumulative number, if specified) given the fixed parameters of sample size, population size, and total population successes. Understanding the exact syntax and parameter mapping is vital when applying this function to our 2x2 table data:

=HYPGEOM.DIST(sample_s, number_sample, population_s, number_pop, cumulative)

The parameters must be assigned with extreme care, corresponding directly to the values within the contingency table:

sample_s: This denotes the number of "successes" observed within the sample. Within the 2x2 table structure, this corresponds to the specific frequency value of the cell chosen as the starting point for the calculation (conventionally, the top-left cell, often labeled 'A').

number_sample: This represents the total size of the sample being drawn, which must correspond precisely to the column total containing the chosen cell.

population_s: This parameter signifies the total number of "successes" present in the overall population, corresponding to the row total containing the chosen cell.

number_pop: This is the total population size, which is equivalent to the grand total of all observations across the entire 2x2 table.

cumulative: This parameter requires a logical value. Setting it to **TRUE** instructs the function to return the cumulative distribution function (the probability of observing **sample_s** or fewer successes). Conversely, setting it to **FALSE** returns the probability mass function (the probability of observing exactly **sample_s** successes). For calculating the one-tailed [P-value](#) in Fisher's test, we must use **TRUE** initially to find the cumulative probability of the observed outcome and all more extreme outcomes in the direction being tested.

Statistically, the selection of which cell to use for **sample_s** is arbitrary, as the final calculated [P-value](#) will remain the same irrespective of the cell chosen. For instructional consistency throughout this guide, we will utilize the top-left cell, which corresponds to Male Democrats and holds the frequency value of 4.

Calculating the Initial One-Tailed P-Value

The critical first step in the analysis is determining the one-tailed [P-value](#) that corresponds to our observed data. This value is fundamentally calculated by summing the probabilities of the specific observed distribution and all other possible distributions that are considered more extreme in the same direction (i.e., those tables having a lower frequency in the chosen cell, given that the marginal totals are fixed). By setting the **cumulative** parameter within the function to **TRUE**, the **HYPGEOM.DIST** function efficiently performs this summation automatically.

Using the observed cell value of 4 (Male Democrats) as our reference point, the parameters map precisely to the following values:

sample_s (Value in the target cell): 4

number_sample (Total column count for Democrats): 12

population_s (Total row count for Males): 10

number_pop (Total sample size): 25

cumulative: TRUE

Therefore, the exact Excel formula applied to our specific data set is:

```
=HYPGEOM.DIST(4, 12, 10, 25, TRUE)
```

When this calculation is executed in Excel, the result directly provides the one-tailed probability. The visualization of the calculation setup and the resulting value are shown below:

	A	B	C	D	E	F	G
1		Democrat	Republican	Total			
2	Male	4	9	13			
3	Female	8	4	12			
4	Total	12	13	25			
5							
6							
7	One-tailed p-value:	0.0812	=HYPGEOM.DIST(B2, \$B\$4, \$D\$2, \$D\$4, TRUE)				
8							
9							
10							
11							
12							
13							
14							
15							
16							
17							
18							

Executing this formula yields a one-tailed P-value of approximately **0.0812**. This value represents the exact probability of observing 4 or fewer Male Democrats, assuming the fixed marginal totals. If the research had been based on an a priori directional hypothesis (a one-tailed test), this value would be used immediately for the hypothesis testing conclusion.

Determining the Comprehensive Two-Tailed P-Value

In the vast majority of research applications, a two-tailed test is the required standard, as the analyst is typically interested in detecting an association regardless of its specific direction (i.e., whether males are significantly more or significantly less likely to be Democrats compared to females). Calculating the two-tailed [P-value](#) for the [Fisher's Exact Test](#) is notably more intricate than simply doubling the one-tailed P-value. It requires the identification and summation of the probabilities of all possible tables that are equally or more extreme than the observed table in *both* tails of the probability distribution.

Specifically, the two-tailed P-value is constructed by summing two distinct components:

The one-tailed probability that has already been calculated (the probability of the observed table and all tables more extreme in that initial direction). In our specific case, this is the probability of observing 4 or fewer Male Democrats (**0.0812**).

The cumulative probability of observing tables in the opposite tail that possess a probability equal to or less than the probability of the original observed table. In this opposite tail, the number of Male Democrats would necessarily be greater than 4.

To accurately calculate the probability of the opposite extreme tail, we must first calculate the cumulative probability of the tables that are *less extreme* than the observed table (in the opposite direction) and then subtract this intermediate value from 1. Since our observed cell count is 4, the opposite extreme begins at a cell count of 8 (as the maximum possible count is 10, given the fixed margins, and 8 is the point where the probability of the table is less than or equal to the probability of the observed table). We are therefore interested in the probability of getting 8 or more Male Democrats.

The necessary calculation to find the cumulative probability of 8 or more involves finding the probability of all tables from 0 up to 7 successes in the Male Democrat cell, and then subtracting that cumulative probability from 1. This second component can be expressed in Excel using the formula:

1 - HYPGEOM.DIST(7, 12, 10, 25, TRUE)

The final two-tailed [P-value](#) is the sum of the probabilities derived from the two tails:

=HYPGEOM.DIST(4, 12, 10, 25, TRUE) + (1 - HYPGEOM.DIST(7, 12, 10, 25, TRUE))

The comprehensive setup for this complex two-tailed calculation within an Excel spreadsheet environment is precisely illustrated below:

	A	B	C	D	E	F	G	H	I	J
1		Democrat	Republican	Total						
2	Male	4	9	13						
3	Female	8	4	12						
4	Total	12	13	25						
5										
6										
7	One-tailed p-value:	0.0812								
8	Two-tailed p-value:	0.1152	=HYPGEOM.DIST(B2, \$B\$4, \$D\$2, \$D\$4, TRUE)+1-HYPGEOM.DIST(B3, \$B\$4, \$D\$2, \$D\$4, TRUE)							
9										
10										
11										
12										
13										
14										
15										

Performing this critical summation produces the final two-tailed P-value of approximately **0.1152**.

Interpreting the Results and Drawing Statistical Conclusions

The conclusive step of any statistical hypothesis test involves interpreting the calculated [P-value](#) in relation to the predefined significance level, which is conventionally set at $\alpha = 0.05$. This

significance level establishes the maximum acceptable risk of making a Type I error--the error of incorrectly rejecting a true [Null Hypothesis](#).

In the context of our case study, when comparing the resulting P-values to the critical threshold of 0.05, we observe that both the one-tailed P-value (**0.0812**) and the generally preferred two-tailed P-value (**0.1152**) are greater than the threshold. The two-tailed P-value of 0.1152 specifically indicates that if there were genuinely no underlying association between gender and political preference (i.e., if the [Null Hypothesis](#) were perfectly true), there would still be an 11.52% chance of observing data as extreme or even more extreme than the specific data collected in our sample.

Since the obtained P-value (0.1152) is not less than the established significance level of 0.05, the appropriate statistical decision is to fail to reject the [Null Hypothesis](#). Consequently, based on this sample of 25 students, the researcher does not possess sufficient statistically significant evidence to confidently conclude that an association exists between a student's gender and their political party preference at this university. The minor differences observed within the 2x2 contingency table are highly likely attributable simply to random sampling variability rather than an actual, underlying relationship between the two [categorical variables](#).

Summary and Recommendations for Further Analysis

While Excel offers a technically feasible and accessible method for calculating the [Fisher's Exact Test](#) by utilizing the [Hypergeometric Distribution](#), analysts who routinely conduct frequent or large-scale exact tests are strongly advised to transition to specialized statistical software packages. Dedicated tools such as R, SPSS, or automated online calculators provide superior efficiency, reduce the risk of manual calculation errors, and are essential when dealing with tables larger than 2x2 or when automating the derivation of the complex two-tailed P-value calculation.