

McNemar's Test in R: A Step-by-Step Guide for Paired Data Analysis

Authored by
Mohammed looti

November 8, 2025

RECOMMENDED CITATION

Mohammed looti (2025). *McNemar's Test in R: A Step-by-Step Guide for Paired Data Analysis*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=13381>

The [McNemar's Test](#) stands as a cornerstone in non-parametric statistics, expertly utilized to determine whether a [statistically significant difference](#) exists between proportions derived from [paired data](#). This test is indispensable in fields ranging from medicine to market research, particularly when analyzing designs such as 'before-and-after' interventions, crossover trials, or matched-pair case-control studies where subjects effectively serve as their own controls.

Mastering this technique requires a solid understanding of its theoretical basis and practical implementation. This comprehensive guide details the necessary requirements, the underlying assumptions, and provides a precise, step-by-step tutorial for executing the [McNemar's Test](#) using the powerful statistical programming environment, [R](#). We will transition from conceptual understanding to generating and interpreting actionable results.

Distinguishing McNemar's Test for Paired Nominal Data

The fundamental distinction between McNemar's Test and standard tests for categorical variables, like the ubiquitous Chi-Squared test, lies in the nature of the data collection. While the standard Chi-Squared test assumes sample independence, [McNemar's Test](#) is specifically formulated for dependent samples or matched pairs. It requires [nominal data](#)--data that can be classified into categories--collected from the same individuals at two different points in time or under two different conditions.

The test's statistical power is derived from its unique focus on the [contingency table](#), specifically examining the marginal homogeneity of the classification variables in a 2x2 square matrix. Crucially, the test disregards the concordant pairs (those who maintained the same classification) and centers its calculation entirely on the discordant pairs--subjects whose classification changed between the initial and final measurement. This methodology effectively controls for inherent variability between subjects, making it highly sensitive to true shifts in proportion.

Proper experimental design is paramount. If researchers mistakenly apply McNemar's Test to independent data, the results will be invalid. The strength of this approach is its ability to isolate the effect of an intervention by analyzing only the transitions within the same observational unit. Common applications include evaluating changes in diagnostic status, voting intention, or consumer preference following a specific event or treatment.

Setting Up the Experimental Scenario: Marketing Intervention

To fully grasp the practical application of the test, we will establish a concrete scenario based on market research principles. Imagine a team aiming to measure the direct persuasive impact of a new marketing video promoting a specific piece of legislation. A cohort of 100 participants is selected. Initially, their stance (Support or Do Not Support) on the law is recorded. This measurement establishes the baseline proportion of supporters.

Immediately following the baseline survey, all 100 participants are exposed to the marketing video, which is specifically designed to influence their perspective. A second survey is then administered to record their opinion again. Because the two measurements (before and after the intervention) are taken from the exact same individuals, the data generated is intrinsically [paired data](#).

The analytical objective is to determine if the intervention--the marketing video--caused a [statistically significant difference](#) in the overall proportion of individuals supporting the law. If the change observed is merely due to chance, the test should yield a non-significant result. Conversely, if the video was effective, we expect a statistically significant shift in opinions, captured by the discordance in the resulting 2x2 table.

Data Preparation and Constructing the R Matrix

Before any statistical testing can commence, the raw frequency counts must be structured into a 2x2 matrix format, which serves as the input [contingency table](#) for the [R](#) function. Each cell in this matrix represents the count of individuals characterized by their joint status (Before vs. After the video). The structure must accurately reflect the transitional states of the participants.

Let's review the hypothetical results summarized below. The four cells detail how opinions transitioned. The cells of interest for McNemar's calculation are the discordant cells: 40 individuals switched from "Do Not Support" (Before) to "Support" (After), while 12 individuals switched from "Support" (Before) to "Do Not Support" (After). McNemar's Test fundamentally compares the frequencies of these two discordant counts (40 vs 12) to assess marginal homogeneity.

The following table visualizes the collected data, showing how participants transitioned across the intervention period:

	Opinion Before Marketing Video	
Opinion After Marketing Video	Support	Do Not Support
Support	30	40
Do Not Support	12	18

The initial step in [R](#) is to construct this matrix accurately using the appropriate commands, ensuring the counts are placed in the correct cell position (following R's column-major ordering) and that the row and column names are descriptive of the experimental variables:

Create the 2x2 contingency table matrix

```
data <- matrix(c(30, 12, 40, 18), nrow = 2,  
dimnames = list("After Video" = c("Support", "Do Not Support"),  
"Before Video" = c("Support", "Do Not Support")))
```

```
# View the resulting data matrix  
data
```

Before Video

After Video Support Do Not Support

Support 30 40

Do Not Support 12 18

Implementing McNemar's Test in R

With the data correctly organized in the matrix variable `data`, we proceed to execute the core statistical analysis using R's built-in function, [mcnemar.test\(\)](#). This function automatically calculates the appropriate [Chi-squared](#) statistic based exclusively on the frequencies of the discordant pairs. The test is highly efficient and requires minimal input once the matrix is prepared.

The standard syntax for invoking the function is defined as follows:

```
mcnemar.test(x, y = NULL, correct = TRUE)
```

A clear understanding of the parameters is essential for precise execution. The arguments control how the function processes the data and whether specific statistical adjustments are applied:

x: This required argument represents the primary input data. It should be the two-dimensional [contingency table](#) provided in matrix format (as we created) or a factor object.

y: This is an optional factor object, which is entirely disregarded if the input `x` is provided as a matrix.

correct: This logical argument (defaulting to **TRUE**) dictates whether the [continuity correction](#) should be applied to the test statistic calculation.

The application of the [continuity correction](#) is a statistical safeguard used to improve the approximation when mapping the discrete counts to the continuous [Chi-squared](#) distribution. While best practice often suggests applying the correction when discordant cell expected counts are low (e.g., less than 5), given our observed counts (12 and 40), the difference will be marginal. For thoroughness, we will demonstrate the execution both with and without the correction to observe its effect on the final results.

Analyzing the Results and P-value Interpretation

We now execute the [mcnemar.test\(\)](#) function twice: first using the default correction (`correct = TRUE`), and second by explicitly disabling it (`correct = FALSE`). Comparing the outputs allows us to observe the minute impact of the [continuity correction](#) on both the calculated test statistic and

the resulting [p-value](#).

```
# Perform McNemar's Test with default continuity correction (correct = TRUE)
mcnemar.test(data)
```

McNemar's Chi-squared test with continuity correction

data: data

McNemar's chi-squared = 14.019, df = 1, p-value = 0.000181

```
# Perform McNemar's Test without continuity correction (correct = FALSE)
mcnemar.test(data, correct=FALSE)
```

McNemar's Chi-squared test

data: data

McNemar's chi-squared = 15.077, df = 1, p-value = 0.0001032

In both computations, R returns the calculated McNemar's [Chi-squared](#) statistic, the degrees of freedom (always 1 for a 2x2 table), and the corresponding [p-value](#). The central purpose of this analysis is to test the [null hypothesis](#), which posits that there is marginal homogeneity--meaning the probability of an individual switching from category A to B is equal to the probability of switching from B to A. In context, the null hypothesis suggests the marketing video had no measurable effect on the population proportion.

Drawing Substantive Conclusions

The interpretation of the results hinges entirely on comparing the calculated [p-value](#) against a pre-established significance level (α), conventionally set at 0.05. If the calculated [p-value](#) is less than α , the evidence is strong enough to justify rejecting the [null hypothesis](#), thereby asserting that the observed marginal differences are [statistically significant](#).

Examining our R output, both computations yielded extremely small [p-values](#) (0.000181 and 0.0001032). Since both values are substantially smaller than the 0.05 threshold, we confidently reject the [null hypothesis](#). This rejection leads to the firm conclusion that the intervention had a non-random, measurable effect.

Specifically, we conclude that there is a [statistically significant difference](#) in the proportion of people supporting the law before and after exposure to the marketing video. Given that the count of individuals who switched to Support (40) vastly outweighs those who switched away from Support (12), we can infer that the marketing video was successful in influencing public opinion

toward supporting the legislation. This outcome powerfully demonstrates the utility of [McNemar's Test](#) in evaluating the efficacy of interventions on paired nominal outcomes.