

# Learn How to Perform a Chi-Square Goodness of Fit Test in R

Authored by  
**Mohammed loot**

November 7, 2025

## RECOMMENDED CITATION

Mohammed loot (2025). *Learn How to Perform a Chi-Square Goodness of Fit Test in R*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=11961>

The [Chi-Square Goodness of Fit Test](#) is one of the most fundamental and widely utilized non-parametric statistical procedures. Its primary purpose is to determine if the observed frequency distribution of a single [categorical variable](#) deviates significantly from a specified theoretical or hypothesized distribution. This powerful test is essential for researchers and analysts who need to validate whether empirical data aligns with predefined expectations, models, or null hypotheses.

This comprehensive guide is designed to transform complex statistical theory into actionable steps, demonstrating precisely how to execute and interpret the Chi-Square Goodness of Fit Test using the robust capabilities of the [R statistical environment](#). We will walk through data preparation, function execution, and the critical steps of drawing statistically sound conclusions from R's output.

## Practical Example: Customer Traffic Analysis

To illustrate the application of this test, let us consider a common business scenario. A shop owner hypothesizes that customer visits are evenly distributed across the five weekdays (Monday through Friday). Statistically, this represents a uniform distribution, meaning that the proportion of customers expected on any given day is  $1/5$ , or 20%.

The shop owner's assertion serves as our [null hypothesis](#) ( $H_0$ ). To rigorously test this claim against reality, a market researcher meticulously collected data on customer arrivals over a full week, resulting in the following observed frequencies:

**Monday:** 50 customers

**Tuesday:** 60 customers

**Wednesday:** 40 customers

**Thursday:** 47 customers

**Friday:** 53 customers

The core challenge is determining whether the differences between these observed counts (e.g., 60 customers on Tuesday versus 40 on Wednesday) are substantial enough to reject the idea of an equal distribution, or if these variations are merely the result of expected random sampling fluctuation. We will employ the Chi-Square Goodness of Fit Test in R to provide a definitive, statistical answer to this question.

## Statistical Foundation: Understanding the Test's Core Mechanism

The Chi-Square Goodness of Fit Test operates by calculating a single statistic, denoted as  $\chi^2$  (Chi-squared), which quantifies the discrepancy between the observed frequencies ( $O_i$ ) and the expected frequencies ( $E_i$ ). The formula essentially squares these differences and weights them by the expected frequency to ensure that larger expected counts do not disproportionately influence the result.

A small  $\chi^2$  value suggests that the observed data aligns closely with the expected distribution specified by the null hypothesis. Conversely, a large  $\chi^2$  value indicates a significant divergence, suggesting that the hypothesized distribution is likely incorrect. The distribution of this test statistic is governed by a parameter known as the [degrees of freedom](#) (df), which is calculated as the number of categories minus one ( $k - 1$ ). In our customer traffic example, with five weekdays, the degrees of freedom will be  $5 - 1 = 4$ .

It is vital to understand that the test requires sufficient data; specifically, the expected frequency in each category should ideally be 5 or greater to ensure that the Chi-Square approximation is valid. In our scenario, the total observed count is  $50 + 60 + 40 + 47 + 53 = 250$ . Since the expected proportion for each day is 0.2, the expected count for each day is  $250 \times 0.2 = 50$ , which easily satisfies this critical assumption.

## Step 1: Defining Observed Frequencies and Expected Proportions in R

The initial and perhaps most critical phase in conducting any statistical test in R is the proper structuring of the input data. For the Chi-Square Goodness of Fit Test, we require two distinct numerical vectors. The first vector must contain the actual observed counts from the experiment or survey, and the second must explicitly define the expected proportions under the assumption of the [null hypothesis](#) ( $H_0$ ).

Since the shop owner claims a perfectly equal distribution across five weekdays, the expected proportion for each day is calculated as  $1/5 = 0.2$ . It is an absolute requirement that the vector defining the expected proportions sums precisely to 1.0, representing 100% of the total probability distribution. This meticulous setup ensures the subsequent statistical calculation is accurate and correctly models the hypothesized population.

We use the R `c()` function (combine values into a vector) to store both sets of data:

```
observed <- c(50, 60, 40, 47, 53)  
expected <- c(.2, .2, .2, .2, .2) #must add up to 1
```

Ensuring the correct order of elements in both vectors is crucial; the first element of the observed vector (50, Monday) must correspond directly to the first element of the expected proportion vector (0.2, Monday), and so forth. This alignment guarantees that R correctly pairs the observed counts with their theoretical probabilities during the calculation.

## Step 2: Executing the Chi-Square Test using `chisq.test()`

The [R statistical environment](#) streamlines complex statistical calculations through its extensive library of built-in functions. The core function for executing the Chi-Square Goodness of Fit Test is

**chisq.test()**. This function automatically handles the calculation of the Chi-Square statistic ( $\chi^2$ ) and determines the corresponding [P-value](#) based on the appropriate degrees of freedom.

The function is highly intuitive but requires two specific arguments to operate correctly when performing a Goodness of Fit test:

The standard syntax utilized for this specific test is:

**chisq.test(x, p)**

Where:

**x**: This parameter accepts the numerical vector representing the **observed frequencies** (our `observed` vector).

**p**: This parameter accepts the numerical vector representing the **expected proportions** (our `expected` vector). Note that R handles the conversion of these proportions into expected counts internally, based on the total sum of the observed frequencies.

Applying this function to our customer traffic data, we execute the test and receive the following structured output from R:

**#perform Chi-Square Goodness of Fit Test**

**chisq.test(x=observed, p=expected)**

Chi-squared test for given probabilities

data: observed

X-squared = 4.36, df = 4, p-value = 0.3595

This output provides the three critical components necessary for statistical inference: the calculated test statistic ( $\chi^2$ ), the [degrees of freedom](#) (df), and the resultant [p-value](#).

## Interpreting the Results and Drawing Conclusions

The R output reveals that our calculated Chi-Square test statistic (labeled X-squared) is **4.36**, associated with 4 degrees of freedom (df), and yields a corresponding p-value of **0.3595**. Interpreting these numerical results requires careful consideration of the test's underlying hypotheses:

**H0 (Null Hypothesis):** The population distribution of customer traffic is consistent with the hypothesized uniform distribution (i.e., customer traffic is equal across all weekdays).

**H1 (Alternative Hypothesis):** The population distribution of customer traffic is significantly different from the hypothesized distribution (i.e., customer traffic is not equally distributed across all weekdays).

The standard procedure in inferential statistics requires us to compare the calculated p-value against a predefined [significance level](#) ( $\alpha$ ), which is typically set at 0.05. The decision rule is straightforward: if the p-value is less than  $\alpha$ , we have sufficient evidence to reject the [null hypothesis](#) ( $H_0$ ).

In our analysis, the calculated p-value (0.3595) is substantially greater than the conventional significance level ( $\alpha = 0.05$ ). Because  $0.3595 > 0.05$ , we must **fail to reject the null hypothesis**. This statistical decision implies that the observed variations in customer numbers throughout the week (e.g., the difference between 60 and 40) are not statistically significant enough to conclude that the true underlying distribution of customer traffic is unequal.

Therefore, we conclude that based on the collected data, there is insufficient evidence to dispute the shop owner's claim of equal daily traffic. The observed fluctuations are attributed to random sampling variability rather than a genuine difference in the population's daily traffic patterns. For verification purposes, you can confirm using statistical tables or a specialized calculator that the p-value corresponding to the Chi-Square statistic ( $\chi^2 = 4.36$  with  $df = 4$ ) is indeed approximately **0.3595**.

## Additional Statistical Resources

To further enhance your mastery of Chi-Square tests and related statistical functions within [R statistical environment](#), we recommend exploring these valuable resources. These tutorials delve into related applications and specific calculations essential for comprehensive data analysis:

For verification, you can confirm that the p-value corresponding to the Chi-Square statistic ( $\chi^2 = 4.36$  with  $df = 4$ ) is indeed **0.35947** using a standard Chi-Square calculator tool.

You can use the [Chi-Square to P Value Calculator](#) to confirm that the p-value that corresponds to  $\chi^2 = 4.36$  with  $dof = 4$  is **0.35947**.

[How to Perform a Chi-Square Test of Independence in R](#)

[How to Calculate the P-Value of a Chi-Square Statistic in R](#)

[How to Find the Chi-Square Critical Value in R](#)