

Learning Z-Tests in R: A Tutorial for One and Two Sample Tests

Authored by
Mohammed loot

November 2, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning Z-Tests in R: A Tutorial for One and Two Sample Tests*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=8543>

Introduction to Z-Tests in the R Environment

The [Z-test](#) represents a foundational procedure in inferential statistics, serving the essential purpose of determining whether the means of two populations are statistically dissimilar, given that the population variance (or standard deviation) is known. This powerful statistical tool is indispensable across numerous scientific and professional disciplines, including quality control, financial modeling, and academic research, providing a reliable method for making data-driven decisions regarding population parameters. When transitioning this methodology into practical application, especially within a high-powered environment, leveraging the capabilities of the [R programming environment](#) becomes crucial for efficient and robust computation.

While R provides numerous base functions for statistical procedures, performing Z-tests requires accessing specialized external libraries. This tutorial focuses specifically on utilizing the robust **z.test()** function, which is conveniently packaged within the widely used [BSDA](#) package (Basic Statistics and Data Analysis). The primary advantage of employing this function is its ability to streamline the entire process of [hypothesis testing](#), simplifying the steps required when researchers are working with scenarios where the true [population standard deviation](#) is an established value. This structure allows practitioners to focus more intently on the interpretation of results rather than complex manual calculations.

Mastering the accurate implementation of the **z.test()** function is paramount for anyone involved in rigorous [statistical analysis](#). We will provide a thorough demonstration of its application in two distinct contexts: the comparison of a single sample mean against a hypothesized population mean (one-sample Z-test), and the comparison of the means derived from two independent samples (two-sample Z-test). Understanding the subtle differences in parameter specification between these two cases is the key to executing valid statistical comparisons and generating trustworthy analytical insights.

Essential Parameters of the z.test() Function

Before proceeding with practical examples, it is critical to gain a comprehensive understanding of the syntax and the role of each argument within the **z.test()** function provided by the BSDA package. This function's adaptability allows it to seamlessly handle both one-sample and two-sample test scenarios, adjusting its computational logic based on the specific inputs supplied by the user. A clear grasp of these parameters ensures the test is correctly aligned with the underlying research question and assumptions regarding population variances.

The core structure of the command reveals the necessary components for defining the test conditions, including the data inputs, the type of hypothesis being tested, and known population characteristics. It is essential to note that unlike the t-test, the Z-test fundamentally requires the user to input the known population standard deviations, which distinguishes it as a procedure

suitable for large samples or situations where historical population data is readily available and reliable.

The following template illustrates the full range of parameters available for defining the Z-test criteria, showcasing how the function is structured to maximize versatility while maintaining statistical rigor:

```
z.test(x, y, alternative='two.sided', mu=0, sigma.x=NULL, sigma.y=NULL, conf.level=.95)
```

A detailed explanation of each parameter is necessary for proper implementation, particularly when transitioning between the one-sample and two-sample frameworks:

x: This required argument represents the vector containing the observed data values for the primary sample being analyzed. In the case of a one-sample Z-test, this is the only sample data input provided.

y: This argument, which is optional, is exclusively reserved for two-sample Z-tests. It must contain the observed data values for the second independent sample, which is being compared against the sample represented by x .

alternative: This parameter specifies the directionality of the alternative hypothesis ($H?$). Users must select one of the three valid strings: `'greater'` (testing if the mean is greater than the hypothesized value), `'less'` (testing if the mean is less than the hypothesized value), or `'two.sided'`, which is the default setting and tests for any difference (not equal to) between the means.

mu: This critical value defines the hypothesized population mean under the [null hypothesis](#) ($H?$). For a one-sample test, it is the population mean the sample is compared against. For a two-sample test, `mu` represents the hypothesized difference between the two population means, which is typically set to zero (0).

sigma.x: This parameter requires the input of the known [population standard deviation](#) associated with the sample data provided in x . This value cannot be estimated from the sample data itself for a true Z-test application.

sigma.y: Similarly, this parameter requires the known [population standard deviation](#) for the second sample (y). This must be specified only when conducting a two-sample comparison.

conf.level: This optional parameter dictates the desired significance level for the resulting [confidence interval](#). The standard industry practice dictates a value of `.95`, corresponding to a 95% confidence level, which is the function's default.

Case Study 1: Conducting a One Sample Z-Test

The one-sample Z-test is specifically designed for scenarios where a researcher aims to assess whether the mean of a single collected sample differs significantly from a well-established, known

population mean. To illustrate this, consider a classic scenario involving IQ scores. It is widely accepted that IQ scores in the general population are [normally distributed](#), possessing a known population mean (μ) of 100 and a known population standard deviation (σ) of 15. This established distribution forms the basis of our [null hypothesis](#).

Imagine a pharmaceutical researcher developing a new cognitive-enhancing medication. The researcher hypothesizes that this medication might, in fact, alter the average IQ level of those taking it. To test this claim, an independent sample of 20 patients is recruited, who then use the medication for a specified duration, after which their final IQ scores are meticulously recorded. The core objective is to determine if the resulting sample mean obtained from these 20 patients exhibits a statistically significant departure from the benchmark population mean of 100. This requires setting up the [null hypothesis](#) ($H_0: \mu = 100$) against the two-sided alternative hypothesis ($H_1: \mu \neq 100$).

The following R code snippet executes the one-sample Z-test. It first loads the essential BSDA library, defines the sample data collected from the 20 patients, and then invokes the `z.test()` function. Crucially, we must specify the hypothesized population mean using the `mu` argument and the known population standard deviation using `sigma.x`. Since this is a one-sample test, the `y` and `sigma.y` arguments are omitted, prompting the function to automatically perform the one-sample calculation:

library(BSDA)

```
#enter IQ levels for 20 patients
data = c(88, 92, 94, 94, 96, 97, 97, 97, 99, 99,
105, 109, 109, 109, 110, 112, 112, 113, 114, 115)

#perform one sample z-test
z.test(data, mu=100, sigma.x=15)
```

One-sample z-Test

```
data: data
z = 0.90933, p-value = 0.3632
alternative hypothesis: true mean is not equal to 100
95 percent confidence interval:
96.47608 109.62392
sample estimates:
mean of x
103.05
```

The resulting output clearly presents the computed [test statistic](#) (Z), the corresponding [p-value](#), the precise 95% confidence interval for the population mean, and the calculated sample mean. These elements are essential for rigorously evaluating the initial hypothesis.

Interpreting Results: One-Sample Z-Test Analysis

The interpretation of the output from the `z.test()` function requires careful consideration of the calculated Z statistic and, most importantly, the associated [p-value](#). In the preceding example, the calculated Z statistic, which quantifies the difference between the sample mean (103.05) and the hypothesized population mean (100) in terms of standard errors, is determined to be **0.90933**. The critical metric for decision-making, the [p-value](#), is reported as **0.3632**.

The standard procedure for hypothesis testing involves comparing the calculated [p-value](#) against a predetermined significance level, often denoted as alpha (α). Conventionally, this significance level is set at 0.05, representing a 5% risk of incorrectly rejecting a true [null hypothesis](#) (Type I error). The fundamental decision rule in this context is straightforward: if the p-value is less than or equal to α ($P \leq 0.05$), we reject the null hypothesis; otherwise, if the p-value is greater than α ($P > 0.05$), we fail to reject the null hypothesis.

In our cognitive medication example, the calculated p-value of 0.3632 is substantially greater than the established significance threshold of 0.05. This statistical outcome mandates that we fail to reject the [null hypothesis](#). The conclusion drawn from this analysis is that the observed mean IQ score of 103.05 in the sample of 20 patients is not sufficiently different from the population mean of 100 to be considered statistically significant. Therefore, based on the available evidence, the researcher must conclude that the new medication did not cause a statistically significant alteration in the average IQ level of the patient group.

Case Study 2: Executing the Two-Sample Z-Test

While the one-sample test compares a sample to a known population, the two-sample Z-test is utilized when the objective is to assess whether the means of two distinct, independent samples are statistically different from one another. This test is appropriate when the population variances for both groups are known. As a practical scenario, let us consider a comparison of the average IQ levels between individuals residing in two different geographical regions, designated as City A and City B.

For this comparison, we continue to assume that the IQ levels in both populations are [normally distributed](#), and, crucially for the Z-test, we assume that the known [population standard deviation](#) for both City A's population and City B's population is 15. A scientist collects two independent random samples, each consisting of 20 individuals--one sample from City A and one from City B--and records their respective IQ scores. The central research question is whether the mean IQ

levels between the two cities differ significantly.

In the context of the two-sample test, the [null hypothesis](#) (H_0) posits that the true difference between the population means ($\mu_x - \mu_y$) is zero. To execute this test using the `z.test()` function, we must now input data for both \bar{x} (City A) and \bar{y} (City B) and explicitly specify the known population standard deviations for both groups using both the `sigma.x` and `sigma.y` parameters. The code below demonstrates the setup and execution, maintaining the default two-sided alternative hypothesis:

library(BSDA)

```
#enter IQ levels for 20 individuals from each city
cityA = c(82, 84, 85, 89, 91, 91, 92, 94, 99, 99,
105, 109, 109, 109, 110, 112, 112, 113, 114, 114)

cityB = c(90, 91, 91, 91, 95, 95, 99, 99, 108, 109,
109, 114, 115, 116, 117, 117, 128, 129, 130, 133)

#perform two sample z-test
z.test(x=cityA, y=cityB, mu=0, sigma.x=15, sigma.y=15)
```

Two-sample z-Test

```
data: cityA and cityB
z = -1.7182, p-value = 0.08577
alternative hypothesis: true difference in means is not equal to 0
95 percent confidence interval:
-17.446925 1.146925
sample estimates:
mean of x mean of y
100.65 108.80
```

The resulting output clearly displays the crucial elements for comparison: the calculated Z statistic, the [p-value](#) specific to the test of difference, the 95% confidence interval for the difference between population means, and the individual sample means derived from City A (100.65) and City B (108.80).

Interpreting Results: Two-Sample Z-Test Analysis

Analyzing the results of the two-sample Z-test requires focusing on whether the observed difference between the sample means ($108.80 - 100.65 = 8.15$) is large enough, given the known

population variability, to reject the assertion that the true population means are equal. The output provides a calculated Z statistic of **-1.7182**, which indicates that the difference between the means is slightly less than two standard errors away from zero. More importantly, the corresponding [p-value](#) is calculated as **0.08577**.

As in the one-sample case, we compare this p-value against the conventional significance level of $\alpha = 0.05$. Since 0.08577 is clearly larger than 0.05, the data does not provide sufficient statistical evidence to meet the threshold required to reject the [null hypothesis](#). Had the p-value been less than 0.05, we would have concluded that the difference in average IQs was statistically significant.

Based on this comprehensive analysis, the appropriate statistical conclusion is that we fail to reject the null hypothesis. This means there is insufficient statistical evidence from the collected samples to confidently suggest that the true mean IQ levels between the populations of City A and City B are significantly different at the 95% confidence level. While a difference exists in the sample means, it can reasonably be attributed to random sampling variation rather than a true population effect.

Conclusion and Advanced Statistical Testing in R

The Z-test, whether employed in its one-sample or two-sample format, serves as a powerful foundational technique for hypothesis testing when the population variance is known. The **z.test()** function within the **BSDA** package significantly simplifies the implementation of this procedure in R, allowing researchers to quickly and accurately determine the significance of differences between means. Successful execution relies heavily on the correct specification of parameters, particularly the known population standard deviation(s) and the appropriate definition of the null hypothesis.

While mastering the Z-test is a crucial step, R's extensive ecosystem supports a vast array of other statistical analyses tailored for different scenarios, such as the t-test (when population variance is unknown) and ANOVA (for comparing three or more means). Analysts are encouraged to build upon this foundation to tackle more complex data challenges, exploring statistical models that do not rely on the stringent assumption of known population variance.

To further enhance data manipulation and analytical capabilities within R, the following resources provide pathways for expanding knowledge beyond the Z-test. These resources often delve into methodologies for advanced statistical inference, model building, and robust data visualization techniques essential for professional practice:

Tutorials detailing the practical application of other common statistical tests in R, such as the Student's t-test and Chi-Squared tests.

Official documentation and guides covering advanced hypothesis testing techniques, including non-parametric methods for data that violate assumptions of normality.

Resources focused on data preprocessing and visualization, which are necessary prerequisites for any robust statistical modeling effort.