

Perform a Two Sample t-Test in SAS

Authored by
Mohammed loot

November 1, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Perform a Two Sample t-Test in SAS*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=7631>

The Foundation of Comparison: The Two-Sample t-Test

The [two-sample t-test](#) serves as a cornerstone in inferential statistics, providing a robust method to determine whether the average values (means) of two separate and independent [populations](#) exhibit a statistically significant difference. This analytical tool is indispensable across diverse fields, including medical research, engineering quality control, and social sciences, whenever a direct comparison between two distinct groups is necessary to validate experimental findings.

To employ this test accurately, certain assumptions must hold true. Primarily, the underlying data from both groups should ideally follow a [normal distribution](#), and the observations collected must be independent of one another. The procedure relies on formal statistical [hypotheses](#)--the null hypothesis (H_0) and the alternative hypothesis (H_A)--to guide the decision process regarding the equality of the population means.

This comprehensive guide is dedicated to outlining the precise, step-by-step methodology required to successfully execute and interpret this fundamental test utilizing the powerful statistical analysis system, [SAS](#) (Statistical Analysis System).

Practical Application: Defining the Botanical Experiment

To provide a clear context for the application of the two-sample t-test, we will model a real-world scenario involving a biological investigation. Imagine a botanist conducting an experiment aimed at assessing whether two distinct plant species (Species 1 and Species 2) possess significantly different mean heights. This requires gathering empirical data from representative samples of both species under identical growth conditions.

The botanist meticulously collects a sample of 12 plants from each species. The height of each sampled plant is recorded in inches, establishing a clean, structured dataset ready for importation and analysis within the [SAS](#) environment. It is crucial that these two samples are treated as independent groups, as the growth of one species does not influence the growth of the other.

The collected height measurements for the two independent samples are provided below, representing the raw data foundation for our statistical inquiry:

Sample 1 (Species 1) Heights: 13, 15, 15, 16, 16, 16, 17, 18, 18, 19, 20, 21

Sample 2 (Species 2) Heights: 15, 15, 16, 18, 19, 19, 19, 20, 21, 23, 23, 24

The primary objective of the subsequent steps is to leverage SAS programming to formally conduct the two-sample t-test and definitively determine whether the observed difference in mean height between Species 1 and Species 2 is statistically meaningful.

Step 1: Structuring and Loading Data into SAS

Before any statistical procedure can commence, the raw measurements must be organized into a format that SAS can process efficiently. For a two-sample comparison, we need a dataset structured with a minimum of two variables: a categorical variable to identify the group (Species) and a continuous variable holding the measurement (Height).

We initiate the data preparation using the [DATA step](#) in conjunction with the `DATALINES` statement. This allows us to input the botanical height measurements directly into a new SAS dataset, which we name `my_data`, effectively preparing the data for subsequent analytical routines.

/*create dataset: Defining the variables Species and Height*/

```
data my_data;  
input Species $ Height;  
datalines;  
1 13  
1 15  
1 15  
1 16  
1 16  
1 16  
1 17  
1 18  
1 18  
1 19  
1 20  
1 21  
2 15  
2 15  
2 16  
2 18  
2 19  
2 19  
2 20  
2 21  
2 23  
2 23  
2 24  
;
```

run;

A crucial technical detail is the inclusion of the `$` sign immediately following the variable `Species` in the `input` statement. This designation signals to SAS that `Species` is a character variable (or categorical identifier), even though we use the numerical codes 1 and 2. This correct declaration ensures that the variable is treated as a grouping factor rather than a continuous measurement during the comparison procedure.

Step 2: Executing the Analysis with PROC TTEST

With the data successfully loaded and structured, the next stage involves invoking the specific SAS procedure designed for mean comparison: [PROC TTEST](#). This procedure automates the complex calculations required for the t-test and produces all necessary output tables.

The command line must be carefully configured using specific options that define the scope and parameters of the test. We use `data=my_data` to specify the input source; `sides=2` to enforce a two-sided test (checking for inequality, $\mu_1 \neq \mu_2$); `alpha=0.05` to set the standard [significance level](#); and `h0=0` to stipulate the [null hypothesis](#) that the true difference between the population means is zero.

Furthermore, the `CLASS` statement is mandatory for identifying the grouping variable (`Species`), and the `VAR` statement designates the dependent, continuous variable being analyzed (`Height`). These statements ensure the procedure correctly partitions the data for comparison.

```
/*perform two sample t-test*/  
proc ttest data=my_data sides=2 alpha=0.05 h0=0;  
class Species;  
var Height;  
run;
```

Upon execution, this code generates a detailed output that includes descriptive statistics for each species, a formal test for the assumption of variance equality, and the core t-test results. The resulting statistical summary from SAS is visually represented below:

The TTEST Procedure

Variable: Height

Species	Method	N	Mean	Std Dev	Std Err	Minimum	Maximum
1		12	17.0000	2.2962	0.6629	13.0000	21.0000
2		12	19.3333	3.0551	0.8819	15.0000	24.0000
Diff (1-2)	Pooled		-2.3333	2.7024	1.1033		
Diff (1-2)	Satterthwaite		-2.3333		1.1033		

Species	Method	Mean	95% CL Mean		Std Dev	95% CL Std Dev	
1		17.0000	15.5410	18.4590	2.2962	1.6266	3.8987
2		19.3333	17.3922	21.2744	3.0551	2.1642	5.1871
Diff (1-2)	Pooled	-2.3333	-4.6213	-0.0453	2.7024	2.0900	3.8249
Diff (1-2)	Satterthwaite	-2.3333	-4.6316	-0.0350			

Method	Variances	DF	t Value	Pr > t
Pooled	Equal	22	-2.11	0.0460
Satterthwaite	Unequal	20.422	-2.11	0.0469

Equality of Variances				
Method	Num DF	Den DF	F Value	Pr > F
Folded F	11	11	1.77	0.3577

Interpreting Results: Variance Homogeneity and P-Values

The interpretation phase begins by rigorously testing one of the critical underlying assumptions of the t-test: the homogeneity of variances (or equality of variances). SAS provides a dedicated section for this, typically utilizing the F-test to compare the variability (variance) between Species 1 and Species 2.

This variance test is pivotal because the subsequent t-test calculation depends on its outcome. If the variances are significantly unequal, the Satterthwaite method must be used; if they are equal, the more powerful pooled method is preferred. For our botanical data, the F-test yields a [p-value](#) of **.3577**. Given that this p-value (.3577) is considerably larger than our pre-set [significance level](#) (**alpha = 0.05**), we retain the null hypothesis for the variance test. This confirms that we can safely assume the population variances are equal, instructing us to proceed with the results derived from the pooled variance estimate.

We now focus on the pooled t-test results, which address the primary research question regarding

the mean heights. The statistical output provides two key metrics:

t Value: -2.11 (This is the standardized test statistic, measuring how many standard errors the observed difference between the sample means is away from the hypothesized zero difference.)

p-value: .0460 (This probability indicates the likelihood of observing a mean difference as extreme as ours, assuming the [null hypothesis](#) is actually true.)

Recall the formal statistical [hypotheses](#):

H0 (Null Hypothesis): $\mu_1 = \mu_2$ (The mean height of Species 1 is equal to the mean height of Species 2.)

HA (Alternative Hypothesis): $\mu_1 \neq \mu_2$ (The mean height of Species 1 is not equal to the mean height of Species 2.)

Final Conclusion and Reporting of Statistical Findings

The decisive step in hypothesis testing involves comparing the calculated t-test [p-value](#) to the critical [significance level](#) (alpha = 0.05). Our analysis, utilizing the pooled variance method, yielded a p-value of **.0460**.

Because the calculated p-value of .0460 is less than the predetermined alpha level of .05, we possess sufficient statistical evidence to confidently reject the [null hypothesis](#) (H0). Rejecting H0 means that we accept the alternative hypothesis (HA): the difference in mean heights is statistically significant.

For the botanist, this result translates into a definitive conclusion: there is a statistically significant difference in the mean height between the two plant species. The observed disparity is highly unlikely to be attributed solely to random sampling variation. When reporting these statistical findings, it is professional practice to include the complete statistical citation, such as the t-value, the associated degrees of freedom (df=22, in this case), and the p-value (t = -2.11, p = .0460).

Expanding Your Statistical Toolkit

Mastering the two-sample [t-test](#) provides a strong foundation for tackling more intricate statistical problems. [SAS](#) is equipped with numerous other procedures designed to handle more than two groups (ANOVA), non-parametric data, or scenarios where data distribution assumptions are severely violated.

For those interested in expanding their proficiency, the following tutorials detail the execution of other essential statistical tests available within the SAS platform: