

How to Easily Perform a Two-Sample t-Test in SAS

Authored by
stats writer

December 1, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Easily Perform a Two-Sample t-Test in SAS*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=103401>

The two sample t-test in SAS is a fundamental procedure in statistical analysis used to conduct a form of hypothesis testing. Its core purpose is to rigorously compare the means of two distinct and independent samples to ascertain whether a statistically significant difference exists between them. This test is indispensable when data is collected from two separate groups that operate independently, and the objective is to measure the extent of the difference between their underlying population averages.

SAS provides robust statistical tools necessary for executing this complex analysis efficiently. The primary tool employed is the PROC TTEST procedure, which automates the calculation of essential statistics, including the t-statistic, degrees of freedom, and the crucial p-value required for decision-making. Furthermore, supplementary procedures, such as PROC SGPLOT, can be utilized to generate high-quality visual representations, enhancing the interpretation and communication of the results derived from the test.

Theoretical Foundation of the Two-Sample t-Test

A two-sample t-test determines whether the mean value of a variable measured across two independent samples is statistically equivalent or different. This determination relies on comparing the calculated t-statistic against a critical value from the t-distribution, or more commonly, by evaluating the probability associated with the observed difference (the p-value).

For this test to be valid, several key statistical assumptions must be met. First, the data in both samples should ideally follow a normal distribution, although the t-test is relatively robust against minor deviations from normality, especially with larger sample sizes. Second, the observations within each group must be independent, meaning the measurement of one observation does not influence the measurement of any other. Third, and most importantly for the interpretation of the results table in SAS, the assumption of homogeneity of variances must be considered. This requires determining if the population variances for both groups are equal, which dictates which specific form of the t-test (pooled or unpooled variance) should be interpreted.

Understanding these assumptions is critical to correctly selecting and interpreting the SAS output, ensuring that the conclusions drawn about the comparison of population means are statistically sound and reliable for decision-making.

Defining the Null and Alternative Hypotheses

Every statistical test begins with the formulation of the null and alternative hypotheses. These formal statements guide the analysis and define the objective of the study. For a standard two-sided two-sample t-test, the hypotheses are structured as follows:

Null Hypothesis (H_0): This hypothesis posits that there is no true difference between the

mean of the first population (μ_1) and the mean of the second population (μ_2). Mathematically, this is expressed as $H_0: \mu_1 = \mu_2$ (or $\mu_1 - \mu_2 = 0$). We assume this hypothesis is true unless the evidence strongly suggests otherwise.

Alternative Hypothesis (H_A): This hypothesis is the complement of the null hypothesis. For a two-sided test, it suggests that the mean of the first population is not equal to the mean of the second population. Mathematically, this is expressed as $H_A: \mu_1 \neq \mu_2$. Rejecting the null hypothesis means accepting this alternative, concluding that a genuine difference exists.

The goal of using the SAS `PROC TTEST` procedure is to gather sufficient statistical evidence to either reject the null hypothesis in favor of the alternative, or fail to reject the null hypothesis, thereby concluding that any observed difference is likely due to random chance.

Case Study: Comparing Plant Species Heights

To demonstrate the practical application of the two-sample t-test, consider a scenario in botany. Suppose a researcher is investigating two genetically distinct species of plants (Species 1 and Species 2) and wishes to determine if the average mature height differs significantly between them. The botanist assumes that the environmental factors are controlled and that any variation in height is primarily attributable to species differences.

The botanist collects an independent and random sample of 12 plants from each species. The recorded heights, measured in inches, form the dataset for our analysis. The use of an independent random sample is crucial to maintain the statistical validity required by the t-test framework, ensuring that the results can be generalized back to the entire population of each species.

The raw data collected for each sample is presented below, categorized by species:

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

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

We will now use a sequence of steps in SAS, starting with data creation and culminating in the execution and interpretation of the two-sample t-test, to formally determine if the mean height is statistically equivalent between Species 1 and Species 2.

Step 1: Creating and Preparing the SAS Dataset

The initial and necessary step in any SAS analysis is defining the dataset structure and populating it with the raw observational data. We must organize the height measurements into a structure that SAS's procedural steps can readily recognize and process. This typically involves defining two

variables: one categorical variable to identify the group (Species) and one continuous variable containing the measurement (Height).

The following SAS code utilizes the `DATA` step and the `DATALINES` statement to create a temporary dataset named `my_data`. The species variable is defined as a character variable using the dollar sign (`$`) to accommodate non-numeric identifiers, although in this case, we use '1' and '2' as indicators.

The input format specifies that SAS should read the Species identifier first, followed immediately by the corresponding Height measurement. This structure is essential for running the subsequent comparison procedures correctly.

```
/*create dataset for plant heights comparison*/
```

```
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 19
```

```
2 20
```

```
2 21
```

```
2 23
```

```
2 23
```

```
2 24
```

```
;  
run;
```

Execution of this code loads the 24 observations (12 for Species 1 and 12 for Species 2) into the SAS environment, making them ready for statistical processing. The data is now properly formatted for the `PROC TTEST` procedure, which requires a classification variable (`Species`) and a measurement variable (`Height`).

Step 2: Executing the PROC TTEST Procedure

Once the data is prepared, the next crucial stage is invoking the statistical procedure itself. We use `PROC TTEST` specifically designed for mean comparisons. This procedure automatically handles the calculations for both the assumption check (equality of variances) and the final mean comparison test.

The code below includes several key options within the `PROC TTEST` statement to ensure the test is conducted according to standard scientific practice:

DATA=my_data: Specifies the dataset we created in Step 1.

SIDES=2: Indicates a two-sided test, meaning we are testing if μ_1 is simply not equal to μ_2 (as opposed to being strictly greater than or strictly less than).

ALPHA=0.05: Sets the significance level (α) for the test. This is the maximum acceptable probability of rejecting a true null hypothesis (Type I error). A conventional value of 0.05 is used here.

H0=0: Specifies the null hypothesis difference. Since we are testing if the means are equal ($\mu_1 - \mu_2 = 0$), the hypothesized difference is zero.

The subsequent statements, `CLASS Species` and `VAR Height`, tell SAS which variable defines the two groups being compared and which variable holds the measurements to be tested, respectively. This configuration directs SAS to compare the mean Height between the classes defined by the Species variable.

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

Upon execution, SAS generates extensive output, including descriptive statistics, an assumption test, and the actual t-test results. The following image represents a typical output snippet generated by SAS:

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 Output: Assessing Equality of Variances

Before interpreting the t-test results themselves, it is statistically mandatory to examine the section of the output titled **Equality of Variances**. This section uses the F-test (specifically the Folded F statistic in many SAS versions) to test the critical assumption that the variance of the Height measurements is equal across both Species populations.

The null hypothesis for the F-test is that the population variances are equal ($\sigma_1^2 = \sigma_2^2$). The alternative hypothesis is that they are unequal ($\sigma_1^2 \neq \sigma_2^2$). We assess this using the F-test's corresponding p-value.

In the provided output, we look for the p-value associated with the F-test. If this p-value is less than the chosen significance level ($\alpha = 0.05$), we reject the null hypothesis of equal variances,

concluding that the variances are heterogeneous (unequal). Conversely, if the p-value is greater than or equal to 0.05, we fail to reject the null hypothesis, assuming that the variances are homogeneous (equal).

In our example output, the p-value for the Equality of Variances F-test is **0.3577**. Since $0.3577 > 0.05$, we do not have sufficient evidence to conclude that the variances are unequal. Therefore, we proceed by assuming that the two population variances are equal. This assumption directs us to interpret the row in the t-test results labeled "Pooled" rather than the row labeled "Satterthwaite" (which is used when variances are unequal).

Interpreting Output: Drawing the Final Conclusion

The final step involves analyzing the core t-test results table, specifically focusing on the row corresponding to the assumption we just confirmed (Equal variances, or "Pooled"). This row provides the calculated t Value and the final p-value for the comparison of the means of the two species.

Key Results from the Pooled (Equal Variances) row:

t Value: -2.11

p-value: 0.0460

We recall the central hypothesis testing framework:

H_0 : $\mu_1 = \mu_2$ (Mean heights are equal)

H_A : $\mu_1 \neq \mu_2$ (Mean heights are not equal)

To make a decision, we compare the calculated p-value (0.0460) against the predetermined significance level ($\alpha = 0.05$). The rule is simple: if p-value $\leq \alpha$, we reject H_0 . If p-value $> \alpha$, we fail to reject H_0 .

In this case, since the p-value of **0.0460** is less than 0.05, we must reject the Null Hypothesis. This statistical decision indicates that the observed difference between the mean heights of Species 1 and Species 2 is too large to be attributed merely to random sampling variability.

Conclusion and Implications of the Test

The rejection of the null hypothesis leads to the statistically sound conclusion that there is a significant difference in the mean height between the two species of plants under investigation. Specifically, based on the sample data and the rigorous two-sample t-test performed in SAS, we

have sufficient evidence at the 5% significance level to assert that the average height of Species 1 is not equal to the average height of Species 2.

The p-value of 0.0460 quantifies the probability of observing the current data (or more extreme data) if the null hypothesis of equal means were actually true. Since this probability is low (less than 4.6%), the null hypothesis is highly improbable given the collected data.

For the botanist, this conclusion validates the need for further research into the underlying genetic or environmental factors driving this height disparity. The utilization of SAS procedures provides a reproducible and mathematically precise method for reaching this important scientific conclusion.

The following resources offer detailed steps on performing other common statistical tests using SAS:

Tutorial on One-Way ANOVA in SAS

Implementing Linear Regression using PROC REG

Performing Chi-Square Tests in SAS