

How to Compare Two Groups Using a Two Sample t-test in Stata

Authored by
stats writer

December 28, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Compare Two Groups Using a Two Sample t-test in Stata*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=109616>

A Two-Sample t-test in Stata is a fundamental statistical tool used to determine if the means of two distinct, independent populations are statistically equivalent. This procedure is essential when researchers need to assess whether a treatment, intervention, or inherent group difference results in a significant shift in the dependent variable's average value. By comparing the means of two groups, this test allows us to ascertain if the observed difference is likely due to genuine variation or merely random sampling error.

The test is typically executed using the dedicated Stata command, `ttest`. This command requires specifying the variable of interest, the grouping variable that defines the two samples, and any specific options relating to assumptions (like unequal variances). The resulting output provides crucial metrics--the t-statistic, the associated degrees of freedom, and the p-value--which are interpreted together to draw conclusions about the null hypothesis regarding mean equality.

The Two-Sample t-test is specifically designed to test whether the population means of two groups are equal, based on samples drawn from each population. The subsequent sections will provide a step-by-step guide on how to successfully execute and interpret this test using the Stata statistical software package.

Prerequisites and Assumptions of the Two-Sample t-test

Before proceeding with the analysis in Stata, it is critical to ensure that the data meets the core assumptions of the independent samples t-test. Failure to meet these assumptions may necessitate the use of a non-parametric alternative, such as the Mann-Whitney U test.

Independence: Observations within each sample, and between the two samples, must be independent.

Normality: The dependent variable should be approximately normally distributed within each of the two groups. While the t-test is robust to minor violations, severe skewness may require data transformation or alternative methods.

Homogeneity of Variance: The population variances of the two groups should be equal. Stata provides an option (the Welch's t-test, denoted by the `unequal` option) if this assumption is violated, which adjusts the degrees of freedom calculation.

Case Study: Fuel Treatment Effectiveness

To illustrate the application of the Two-Sample t-test, consider a practical research scenario. A team of automotive engineers is investigating whether a newly developed fuel additive significantly alters the average miles per gallon (mpg) performance of a specific type of vehicle. To test this, they design a controlled experiment:

A total of 24 identical cars are randomly selected. Twelve cars are assigned to the treatment group (receiving the new fuel additive), and the remaining twelve cars serve as the control group (using standard fuel). The dependent variable is the average mpg recorded for each vehicle.

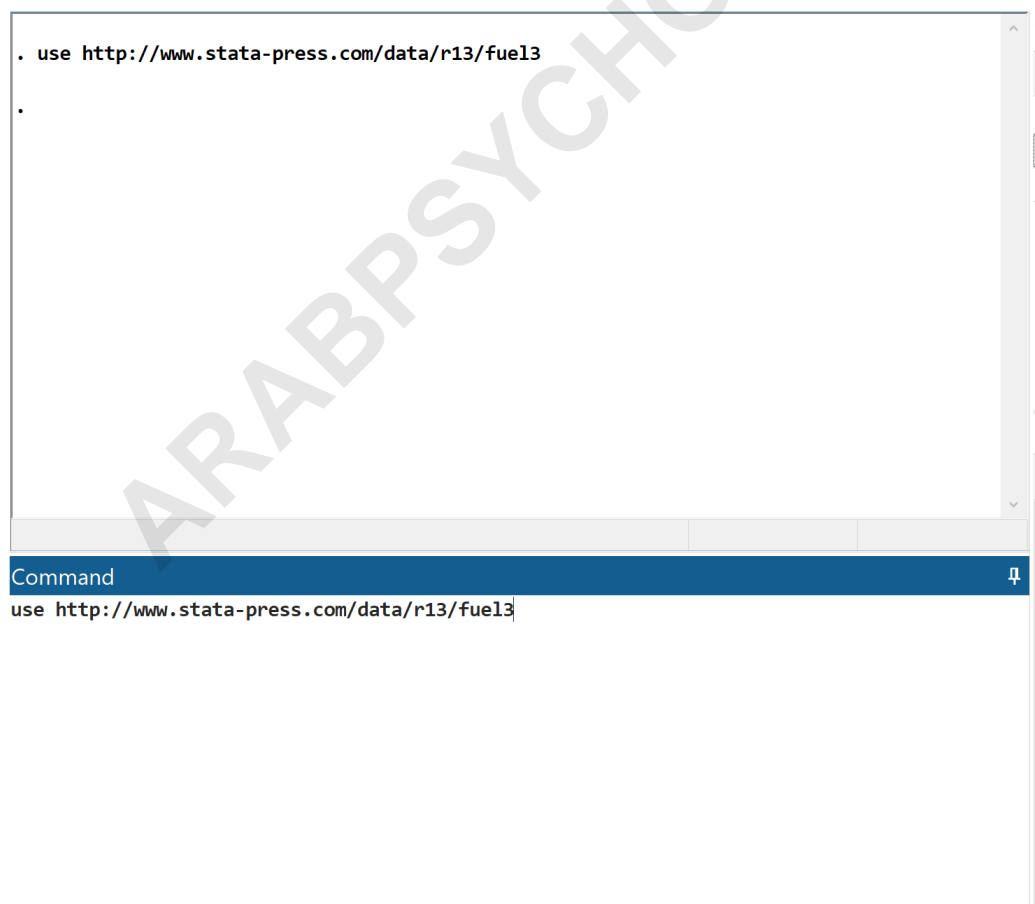
The goal is to conduct a Two-Sample t-test to formally determine if the difference in average mpg between the treated and non-treated groups is statistically significant, thus supporting the efficacy of the new additive.

Step 1: Loading the Dataset into Stata

The initial requirement for any analysis in Stata is loading the appropriate dataset. For this example, we will use a sample dataset available directly through Stata's repository. In the Stata command window, type the following command precisely and execute it by pressing Enter:

```
use http://www.stata-press.com/data/r13/fuel3
```

This command retrieves and loads the dataset named `fuel3`, making the variables accessible for immediate analysis.



```
. use http://www.stata-press.com/data/r13/fuel3
.
```

Command

```
use http://www.stata-press.com/data/r13/fuel3
```

Step 2: Reviewing the Raw Data Structure

Before running the statistical test, it is good practice to visually inspect the raw data to understand its organization and coding scheme. Navigate to the top menu bar in Stata and select **Data > Data Editor > Data Editor (Browse)**. This action opens a read-only view of the dataset.

The dataset contains two critical columns for our analysis. The first column, **mpg**, represents the measured miles per gallon--this is our continuous dependent variable. The second column, **treated**, is the binary grouping variable, indicating sample membership (where **0 = control/no treatment** and **1 = treatment received**). Understanding this coding is crucial for correct analysis execution.

	mpg	treated		
1	20	0		
2	23	0		
3	21	0		
4	25	0		
5	18	0		
6	17	0		
7	18	0		
8	24	0		
9	20	0		
10	24	0		
11	23	0		
12	19	0		
13	24	1		
14	25	1		
15	21	1		
16	22	1		
17	23	1		
18	18	1		
19	17	1		
20	28	1		
21	24	1		
22	27	1		
23	21	1		
24	23	1		

Step 3: Visualizing Group Differences with Box Plots

A necessary preliminary step in comparative statistics is data visualization, which offers an initial, intuitive assessment of group differences and data spread. We will construct side-by-side box plots to visualize the distribution of mpg values for both the treated and control groups.

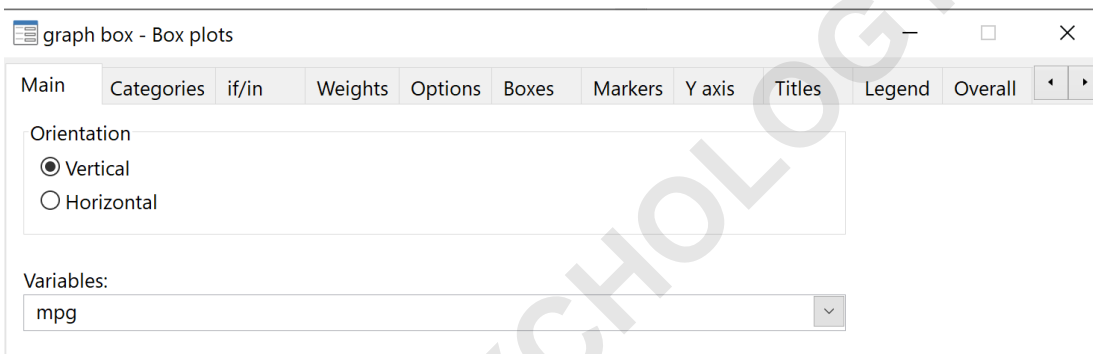
To generate the visualization, proceed through the following menu path: **Graphics > Box plot**.

In the resulting dialogue box:

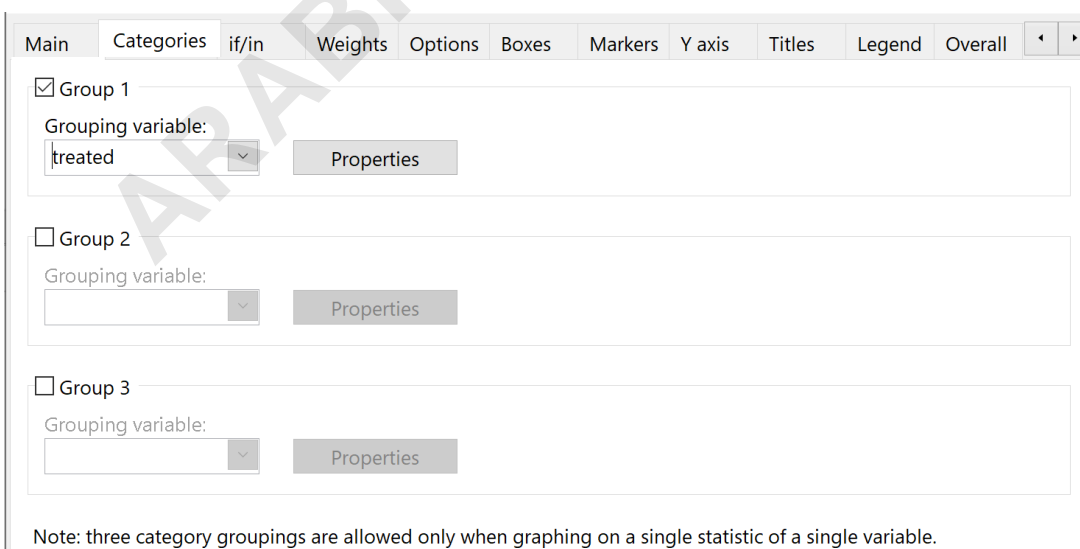
Under the main variables section, select **mpg** as the variable to be plotted.

In the Categories subheading, specify **treated** as the Grouping variable.

This setup instructs Stata to create separate box plots for the mpg variable, categorized by the treatment status.




Next, confirm the Grouping variable selection:



The resulting graphical output immediately suggests a potential difference: the box plot for the

treated group (coded 1) appears shifted upward, indicating a higher average mpg compared to the non-treated group (coded 0). However, this visual evidence must be confirmed through formal hypothesis testing to ensure the difference is statistically robust rather than random fluctuation.



Step 4: Executing the Two-Sample t-test

To formally test the null hypothesis ($H_0: \mu_0 = \mu_1$) that the population means are equal, we proceed with the Two-Sample t-test using Stata's dialogue box interface. Follow this menu path:

Statistics > Summaries, tables, and tests > Classical tests of hypotheses > t test (mean-comparison test).

In the dialogue window that opens, configure the test parameters as follows:

Under the type of t-test, select **Two-sample using groups**.

For the Variable name (the dependent variable), choose **mpg**.

For the Group variable name (the independent variable), choose **treated**.

The Confidence level is typically set to 95%, corresponding to an alpha (significance) level of 0.05. We will retain the default 95% level.

After configuring these options, click **OK** to run the analysis.

ttest - t tests (mean-comparison tests)

Main by/if/in

t tests

One-sample

Two-sample using groups

Two-sample using variables

Paired

Two-sample t test using groups

Variable name: mpg

Group variable name: treated

Reverse group order

Unequal variances

Welch's approximation

95 Confidence level

? ↻ 📄 OK Cancel Submit

Interpreting the Stata Output

The output generated by *Stata* is comprehensive and provides summary statistics alongside the formal test results. Understanding each component is essential for accurate interpretation.

```
. ttest mpg, by(treated)
```

Two-sample t test with equal variances

Group	Obs	Mean	Std. Err.	Std. Dev.	[95% Conf. Interval]	
0	12	21	.7881701	2.730301	19.26525	22.73475
1	12	22.75	.9384465	3.250874	20.68449	24.81551
combined	24	21.875	.6264476	3.068954	20.57909	23.17091
diff		-1.75	1.225518		-4.291568	.7915684

diff = mean(0) - mean(1)

t = -1.4280

Ho: diff = 0

degrees of freedom = 22

Ha: diff < 0

Ha: diff != 0

Ha: diff > 0

Pr(T < t) = 0.0837

Pr(|T| > |t|) = 0.1673

Pr(T > t) = 0.9163

The upper section summarizes descriptive statistics for the two groups:

Obs: The number of observations (sample size) in each group. Here, both the control (0) and treated (1) groups have **12 observations**.

Mean: The arithmetic average of the dependent variable. The mean mpg for the control group is 21.00, while the mean mpg for the treated group is 22.75.

Std. Err: The standard error of the mean, calculated as the standard deviation divided by the square root of the sample size (σ / \sqrt{n}).

Std. Dev: The standard deviation, measuring the variability or spread of the mpg values within each group.

95% Conf. Interval: This represents the 95% confidence interval for the true population mean of mpg for that specific group.

The lower section displays the core inferential statistics:

t: The calculated t-statistic. In this case, \$t = -1.4284\$. This value quantifies the difference between the sample means relative to the variability within the samples.

degrees of freedom: The degrees of freedom (df) used in the test, calculated as \$n_1 + n_2 - 2\$. With 12 cars in each group, \$df = 12 + 12 - 2 = 22\$.

At the bottom, three p-values are provided, corresponding to different alternative hypotheses:

P-value for $H_1: \text{diff} < 0$ (One-tailed test: Is the treated mean significantly lower?).

P-value for $H_1: \text{diff} \neq 0$ (Two-tailed test: Is there any difference, positive or negative?).

P-value for $H_1: \text{diff} > 0$ (One-tailed test: Is the treated mean significantly higher?).

Since the initial research question was whether the new fuel treatment leads to a *change* (i.e., a difference in either direction) in average mpg, we focus on the two-tailed test result ($H_1: \text{diff} \neq 0$). This test yields a p-value of **0.1673**.

To make a decision, we compare the p-value to the chosen significance level ($\alpha = 0.05$). Since $0.1673 > 0.05$, we fail to reject the null hypothesis. Therefore, based on this sample, there is insufficient statistical evidence to conclude that the true mean mpg differs between the treated and non-treated groups.

Reporting the Findings

The final step in any statistical analysis is clearly and concisely reporting the results, ensuring all necessary statistical metrics are included. The report should summarize the findings, the test used, the decision regarding the null hypothesis, and the confidence interval for the difference between the means.

A Two-Sample t-test was conducted on 24 cars ($n=12$ per group) to determine if the application of a new fuel treatment led to a significant difference in mean miles per gallon (mpg).

The descriptive statistics showed that the mean mpg for the treated group ($M = 22.75$, $SD = 2.45$) was numerically higher than the control group ($M = 21.00$, $SD = 2.14$). However, the results of the t-test indicated that this difference was **not** statistically significant ($t_{22} = -1.43$, $p = 0.1673$) at the standard $\alpha = 0.05$ level.

The 95% confidence interval for the true difference in population means (Treated - Control) was calculated as $(-4.29, 0.79)$. Since this interval includes zero, it further supports the finding that the difference between the two population means is not statistically significant.