Learning Independent Samples t-Tests in Stata: A Step-by-Step Guide

Authored by

Mohammed looti

November 8, 2025

RECOMMENDED CITATION

Mohammed looti (2025). *Learning Independent Samples t-Tests in Stata: A Step-by-Step Guide*. PSYCHOLOGICAL STATISTICS. Retrieved from https://statistics.arabpsychology.com/?p=13741

The **Independent Samples t-test**, commonly referred to as the two-sample t-test, is a fundamental statistical procedure used widely in quantitative research. Its primary function is to determine whether the population means of two distinct, independent groups are statistically different from one another. This test is crucial for drawing robust conclusions when comparing average outcomes--for instance, between a controlled experimental group and a baseline control group, or between two naturally occurring populations under study.

Mastering this analytical technique is essential for any data analyst or quantitative researcher. This comprehensive guide provides a detailed, step-by-step methodology for executing and correctly interpreting the independent samples t-test using the powerful statistical software package, Stata. We will follow a rigorous five-step process, ensuring complete clarity from initial data preparation and visualization through to the final statistical conclusion and professional reporting.

Theoretical Foundations of the Independent Samples T-Test

The core objective of the two-sample t-test is to rigorously evaluate the strength of evidence against the null hypothesis, symbolized as H_0. The null hypothesis asserts that there is no true difference between the population means of the two groups (i.e., $mu_1 = mu_2$). Conversely, the alternative hypothesis (H_a) posits that a genuine difference does exist ($mu_1 neq mu_2$). A successful analysis hinges on calculating a measure that quantifies the observed difference relative to the expected random variation.

Like all parametric tests, the t-test relies on several key **assumptions** that must be reasonably met for the results to be valid. These include the independence of observations, the measurement of the dependent variable on a continuous scale, and the requirement that the populations from which the samples are drawn are approximately **normally distributed**. While the t-test is generally robust to minor violations of the normality assumption, significant deviations can undermine the reliability of the calculated probabilities and necessitate non-parametric alternatives.

The calculation process yields a **t-value**, which is the test statistic. This value represents the standardized difference between the two sample means. This t-value, combined with the calculation of the degrees of freedom, is then used to determine the critical element of the analysis: the p-value. The p-value provides the essential probability of observing a difference as large as, or larger than, the one calculated, assuming the null hypothesis is actually true.

Practical Application: Fuel Treatment Effectiveness Case Study

To illustrate the two-sample t-test procedure effectively within the Stata environment, we will utilize a practical research scenario focused on automotive performance. This case study investigates whether a newly developed fuel treatment product significantly impacts a vehicle's average **miles per gallon** (mpg). The experiment was conducted under controlled conditions, involving 24 vehicles of the identical make and model to ensure initial homogeneity.

The total sample was systematically divided into two equally sized groups, with $n=12$ vehicles assigned to each. The **experimental group** received the new fuel treatment, while the **control group** received no treatment. Our primary analytical goal is to use Stata to execute the required statistical test and determine if the observed difference in mean mpg between the treated and untreated vehicles is statistically significant when judged against the conventional $alpha=0.05$ significance level.

Step 1: Data Preparation and Importation in Stata

The critical first step in conducting statistical analysis is the successful loading of the dataset into the statistical environment. For maximum reproducibility and ease of access, the data required for this specific case study are hosted remotely by Stata Press. This method allows analysts to pull the required variables directly into their current session without manual file downloads or complex path management.

To load the required dataset, navigate directly to the Stata Command window and execute the following explicit command. Executing this command retrieves the data efficiently and prepares the variables for initial inspection and subsequent hypothesis testing:

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

Once the command is successfully processed, Stata will confirm the data load status in the command output. The variables are then accessible in the memory, ready for the next phase of exploratory data analysis (EDA).

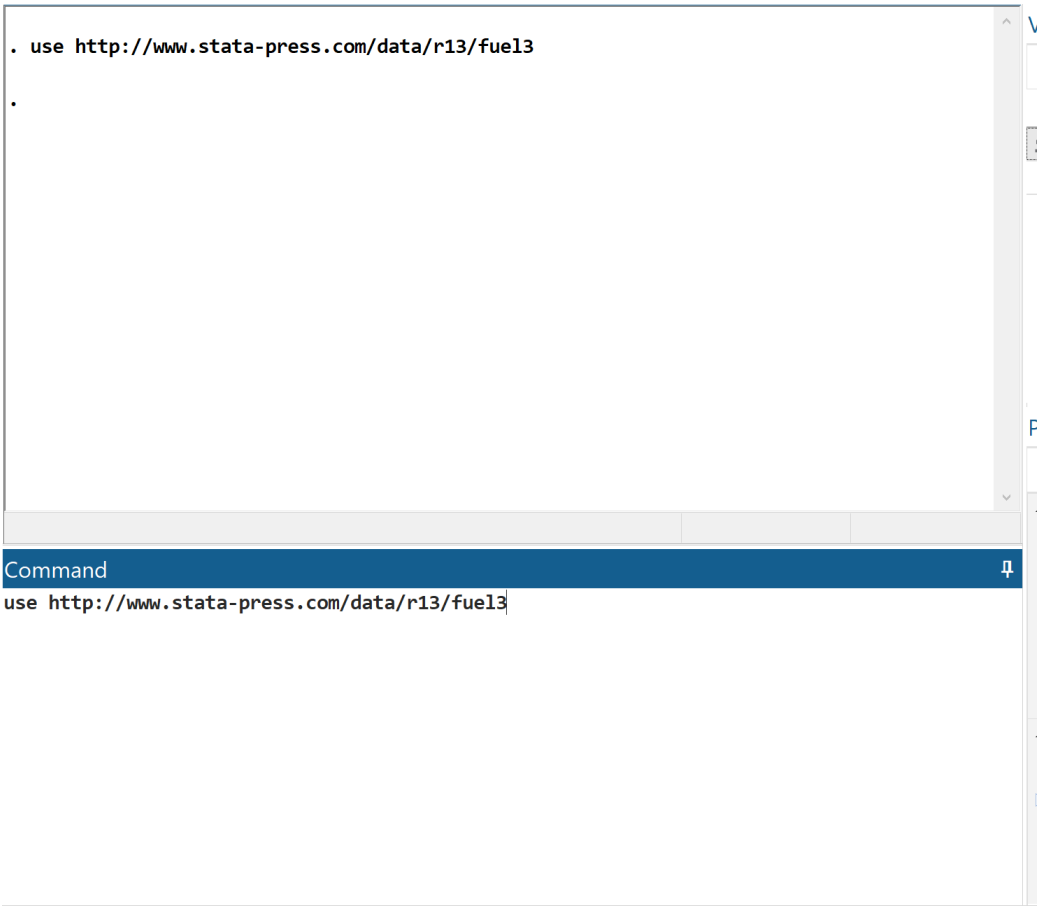

Step 2: Reviewing Dataset Structure and Variables

Before proceeding to any inferential statistics, it is mandatory to review the structure, coding, and content of the imported dataset. This essential process ensures that the variables are correctly identified (dependent vs. independent) and properly coded (continuous vs. categorical). We achieve this most easily by using the Data Editor (Browse) feature in Stata.

Access the data viewer through the main navigation menu by selecting **Data > Data Editor > Data Editor (Browse)**. A quick visual inspection of the variables confirms the presence of the two crucial columns required for our t-test analysis:

mpg: This is the **dependent variable** (outcome). It is a continuous variable recording the measured miles per gallon achieved by each individual vehicle in the study.

treated: This is the **independent grouping variable**. It is a categorical variable coded as a binary indicator, where a value of 0 signifies the control group (untreated), and a value of 1 signifies the experimental group (treated).

	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 Distributions using Box Plots

Exploratory data visualization is a powerful precursor to formal statistical testing. Creating side-by-side box plots provides immediate preliminary insights into the central tendency, dispersion, and potential symmetry of the distribution of mpg for both the control and treated groups. This visual check helps confirm whether the data structure aligns with the assumptions of the t-test.

To generate this visualization using Stata's graphical user interface (GUI), follow the specified path and settings carefully:

Start by navigating to **Graphics > Box plot** from the main menu.

Within the primary dialogue box, ensure that the outcome variable, *mpg*, is correctly specified in

the Variables section.

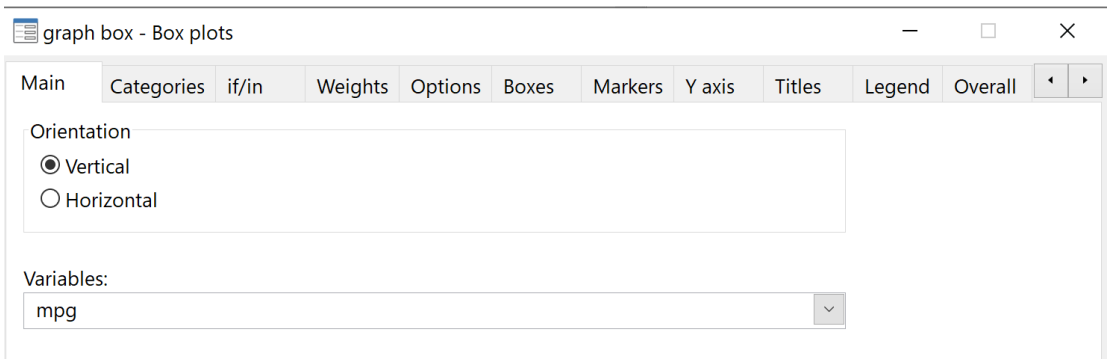

Next, move to the Categories subheading. Here, the critical step is to specify the Grouping variable as *treated*. This instruction compels Stata to produce two distinct box plots, one for each level (0 and 1) of the categorical treatment variable, facilitating direct visual comparison of the distributions.

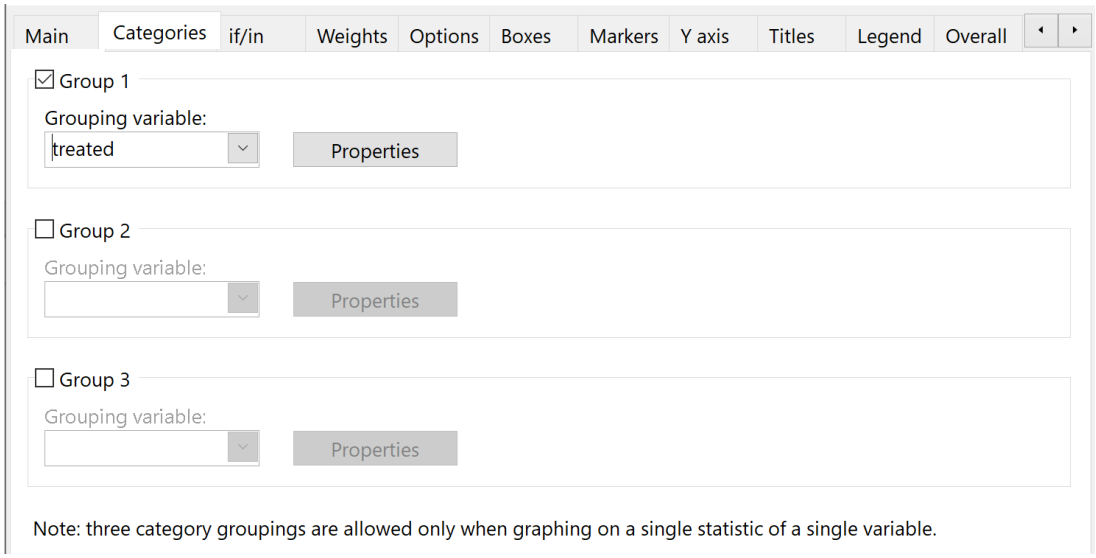

The resulting graph visually confirms that the average mpg is numerically higher for the treated group (1) compared to the control group (0). While this observation provides directional insight, this visual evidence is not sufficient proof of a true effect. We must proceed to the formal independent samples t-test to rigorously test the hypothesis.

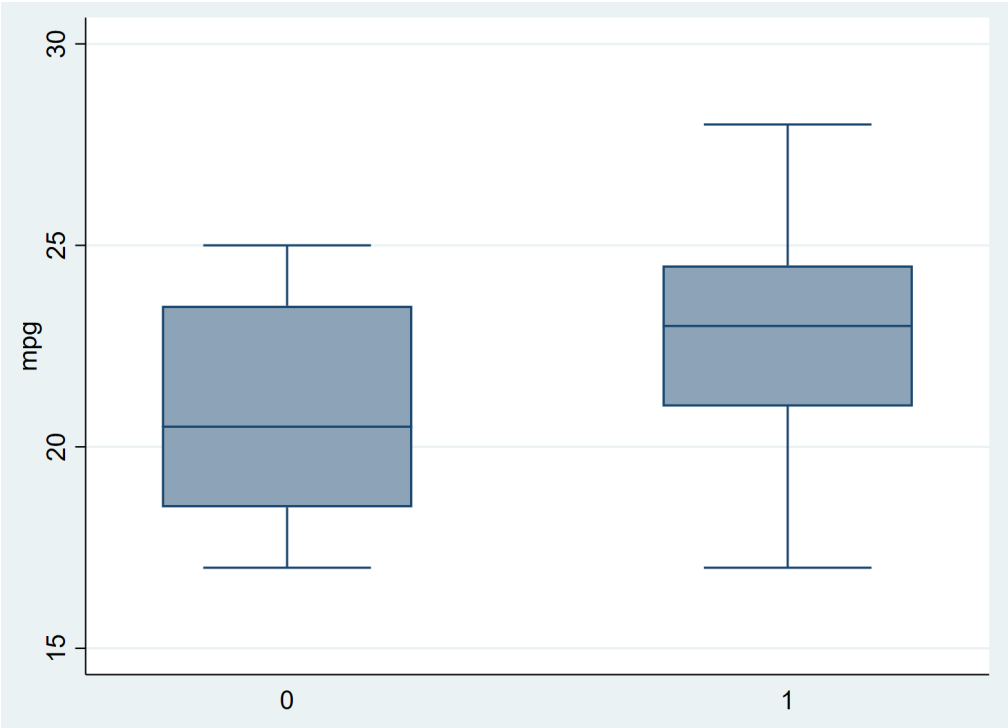

Step 4: Executing and Interpreting the T-Test Output

The formal execution of the t-test is where we quantitatively assess the probability that the difference between the two sample means is merely due to random sampling variation. This procedure is initiated via the following menu sequence: **Statistics > Summaries, tables, and tests > Classical tests of hypotheses > t test (mean-comparison test)**.

Within the resulting dialogue box, precise specification of the test parameters is necessary:

Test Type: Select *Two-sample using groups*.

Outcome Variable: Specify the continuous outcome, *mpg*.

Grouping Variable: Specify the independent variable, *treated*.

Significance Level: Retain the default 95% Confidence level, which aligns with the standard $alpha=0.05$ threshold.

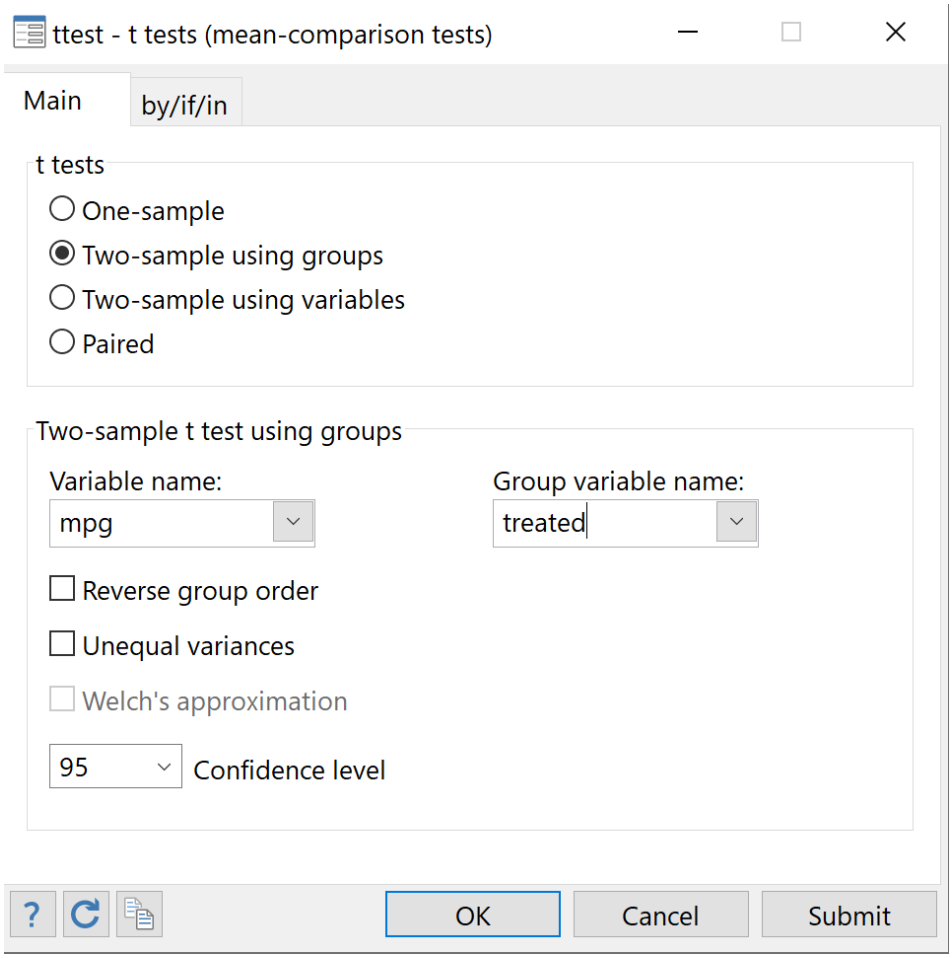

Clicking *OK* generates the comprehensive output window, which is logically divided into descriptive statistics and the core inferential test results. The first section provides key summary measures for each group:

```
. 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
```

From the descriptive statistics, we confirm that the treated group (Group 1) achieved a higher sample mean mpg of 22.75, compared to the control group (Group 0) mean of 21.00. Additional useful measures presented include the **standard deviation** (Std. Dev), **standard error** (Std. Err), and the 95% confidence interval for the true population mean of each group.

Drawing Statistical Conclusions

The second, and most crucial, part of the Stata output is dedicated to the hypothesis test results. Here, we locate the calculated test statistic, which is $t = -1.428$, and the corresponding degrees of freedom, $df = 22$. Since our primary research question involves assessing whether the means are simply unequal (a two-sided test), we focus specifically on the central row labeled H_a: diff neq 0.

The associated p-value for this two-sided test is reported as **0.1673**. To make a final decision, we compare this value to our predetermined significance threshold ($alpha=0.05$). Because 0.1673 is substantially greater than 0.05, the result is deemed **not statistically significant**. Therefore, we must **fail to reject the null hypothesis**.

This finding indicates that the observed numerical difference of 1.75 mpg between the two groups is highly likely attributable to random sampling error rather than being a true, measurable effect of the fuel treatment. In practical terms, the data does not provide sufficient statistical evidence to

conclude that the new fuel treatment effectively changes a car's mileage at the standard level of significance.

Step 5: Professional Reporting of Findings

The final step in the analytical process is to communicate the findings clearly and consistently using standardized statistical notation. This practice is essential for ensuring that the conclusions are accessible, reproducible, and verifiable by the scientific community. A complete statistical report must integrate the key metrics derived from the t-test output, including means, t-statistic, degrees of freedom, and the p-value.

The following block demonstrates the appropriate professional format for summarizing the results of this independent samples t-test, adhering to common reporting standards:

An independent samples t-test was performed on a sample of 24 vehicles ($n=12$ per group) to evaluate the effect of a new fuel treatment on mean miles per gallon (mpg). Descriptive analysis showed that the treated group had a numerically higher mean (M = 22.75, SD = 2.45) compared to the control group (M = 21.00, SD = 2.16).

However, this difference was not statistically significant, $t(22) = -1.428$, $p = 0.1673$. Based on this result, we **fail to reject the null hypothesis**. There is insufficient statistical evidence, at the $alpha = 0.05$ level, to conclude that the fuel treatment has a true impact on vehicle mileage.

The 95% confidence interval for the true difference between the population means was calculated as . Crucially, because this interval spans zero (meaning the true difference could plausibly be zero), it provides strong corroboration for the non-significant outcome obtained from the p-value test.