A Step-by-Step Guide to Paired Samples T-Tests in Stata

Authored by

Mohammed looti

November 8, 2025

A paired samples t-test is a fundamental statistical procedure utilized when researchers aim to compare the means of two related samples. This test is crucial when each observation in the first sample is directly and logically paired with a corresponding observation in the second sample, often occurring in before-and-after studies or repeated measures designs.

This comprehensive tutorial provides an expert explanation of how to efficiently conduct and interpret a paired samples t-test using the powerful statistical software, Stata. We will walk through a practical example, ensuring clarity in every step of the process.

Understanding the Paired Samples t-Test

The core distinction of the paired samples t-test, compared to the independent samples t-test, is the inherent dependency between the two groups being analyzed. When the same subjects or items are measured under two different conditions--such as measuring performance before an intervention and then again after the intervention--the observations are intrinsically linked. This linkage means that the variance between pairs must be accounted for, establishing the paired t-test as the appropriate analytical tool.

The primary purpose of this test is to assess whether the mean difference between the paired observations is **statistically significant**. The underlying statistical assumption, or null hypothesis (H?), is that the true population mean difference between the paired measurements is zero. The alternative hypothesis (H?) posits that a meaningful difference exists, either directional or non-directional, depending on the specific research question being addressed by the investigation.

Case Study: Evaluating Fuel Treatment Efficiency

To demonstrate the utility of the paired t-test, consider a scenario where automotive researchers are investigating the efficacy of a new specialized fuel treatment. Their objective is to determine if this treatment causes a statistically significant change in the average miles per gallon (mpg) achieved by a specific model of car. To execute this test rigorously, they recruit twelve vehicles and measure the mpg for each car first without the treatment, and subsequently with the treatment applied.

Because the measurement is taken twice on the exact same set of twelve cars--once in the control condition (without treatment) and once in the experimental condition (with treatment)--we must treat these observations as **paired data**. Using a paired t-test allows us to accurately isolate the effect of the fuel treatment while controlling for the inherent variability across different individual vehicles. The following steps will guide us through executing this analysis using Stata.

Step-by-Step Implementation in Stata

Performing a paired t-test in Stata is straightforward, whether you prefer using the command line or the graphical user interface (GUI). We begin by ensuring the relevant dataset is properly loaded and inspected.

Step 1: Load the Data

To begin the analysis, the necessary dataset must be loaded into the Stata environment. This can be accomplished by typing the following command directly into the command box and pressing Enter:

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

This action retrieves the dataset needed for our fuel efficiency example, preparing the environment for statistical testing.

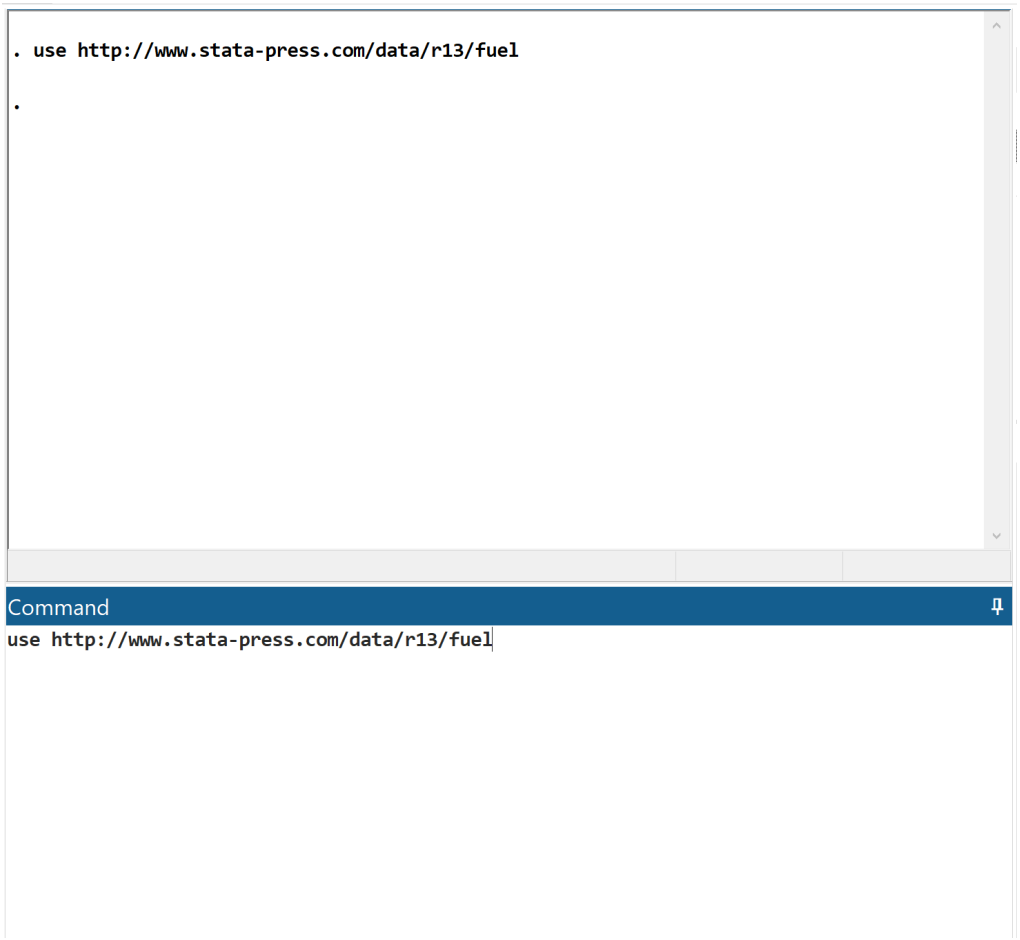

Step 2: Reviewing the Raw Data Structure

Before proceeding with any test, it is best practice to examine the structure and values of the raw data. To view the data, navigate through the top menu bar by selecting **Data > Data Editor > Data Editor (Browse)**. This window will display the two critical variables:

The first column, *mpg1*, represents the miles per gallon achieved by each car **without** the fuel treatment.
The second column, *mpg2*, represents the miles per gallon achieved by the same car **with** the fuel treatment applied.

Observing the data confirms the paired structure, where each row represents a single vehicle measured under two distinct conditions.

	mpg1	mpg2
1	20	24
2	23	25
3	21	21
4	25	22
5	18	23
6	17	18
7	18	17
8	24	28
9	20	24
10	24	27
11	23	21
12	19	23

Step 3: Executing the Paired t-Test

The paired t-test is initiated through the Stata GUI using the following menu path: **Statistics > Summaries, tables, and tests > Classical tests of hypotheses > t test (mean-comparison test)**. Once the t-test dialogue box appears, ensure the following selections are made:

Choose the test type: Select *Paired*.
Define the variables: Input *mpg1* for the First variable and *mpg2* for the Second variable.

Set the Confidence level: The default 95% is standard, corresponding to a significance level (α) of 0.05. We will retain this standard value.

After confirming these parameters, click *OK* to execute the analysis and generate the test results.

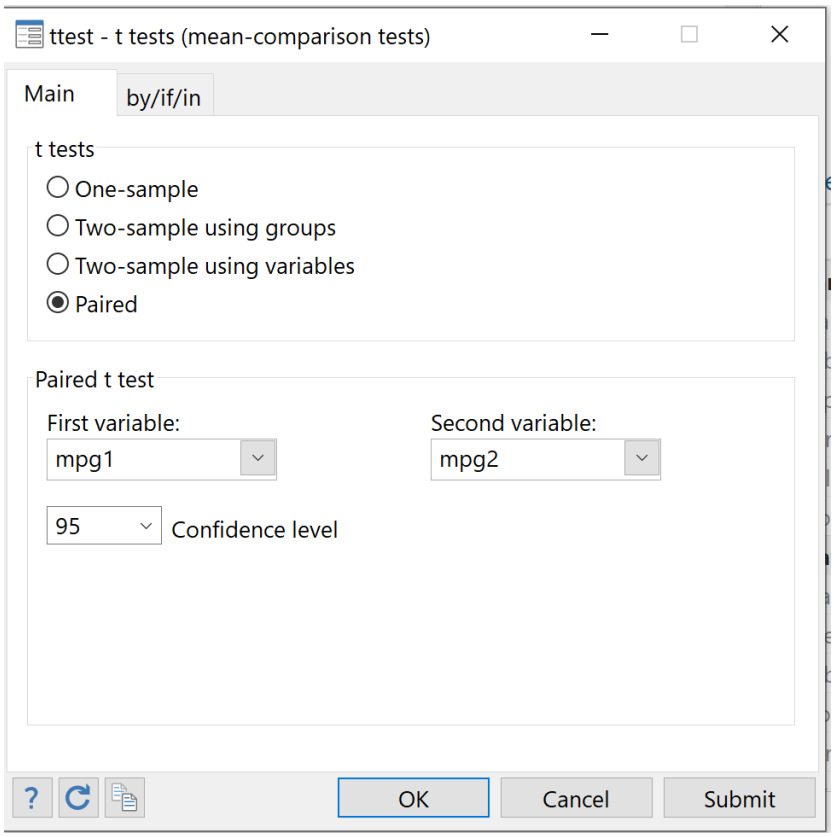

Interpreting the Stata Output

Upon execution, Stata provides a detailed output summarizing both the descriptive statistics for the groups and the inferential results from the paired t-test. Understanding each component is essential for drawing accurate conclusions about the fuel treatment's effect.

The output presents key descriptive statistics for the untreated (mpg1) and treated (mpg2) groups, as well as the difference between the pairs:

Obs: The number of paired observations (n=12).
Mean: The arithmetic average mpg. The untreated mean is 21.00, and the treated mean is 22.75.
Std. Err: The standard error of the mean, which estimates the standard deviation of the sampling distribution.
Std. Dev: The **standard deviation** of mpg, quantifying the spread of data points.
95% Conf. Interval: The 95% confidence interval for the true population mean of mpg for that

specific group.

The core inferential results focusing on the paired difference are:

t: The calculated test statistic of the paired t-test (t = -2.2444).
degrees of freedom (df): Calculated as the number of pairs minus one (n-1 = 12-1 = 11).

```
. ttest mpg1 == mpg2

Paired t test
```

Variable	Obs	Mean	Std. Err.	Std. Dev.	[95% Conf.	Interval]
mpg1	12	21	.7881701	2.730301	19.26525	22.73475
mpg2	12	22.75	.9384465	3.250874	20.68449	24.81551
diff	12	-1.75	.7797144	2.70101	-3.46614	-.0338602

```
    mean(diff) = mean(mpg1 - mpg2)                                t =  -2.2444
Ho: mean(diff) = 0                             degrees of freedom =       11

Ha: mean(diff) < 0           Ha: mean(diff) != 0           Ha: mean(diff) > 0
Pr(T < t) = 0.0232         Pr(|T| > |t|) = 0.0463          Pr(T > t) = 0.9768
```

Step 4: Making the Decision (P-Value Analysis)

The lower section of the output displays the p-values for three different alternative hypotheses. Since our research goal is simply to see if the average mpg is **different** between the two groups, we examine the two-tailed test, identified by the alternative hypothesis H?: diff ? 0. The corresponding p-value is **0.0463**.

We compare this p-value against our predetermined significance level (α) of 0.05. Because 0.0463 is less than 0.05, we must reject the null hypothesis. This result provides sufficient statistical evidence to conclude that the true mean mpg is significantly different between the two conditions, confirming an effect of the fuel treatment.

Drawing Conclusions and Reporting Results

The final and most crucial step in any statistical analysis is reporting the findings accurately and professionally. The conclusion must clearly state the test performed, the relevant statistics, and the interpretation in the context of the original research question. Our analysis indicates that the fuel treatment had a statistically significant positive effect on car efficiency.

Step 5: Report the Results

When reporting a paired t-test, standard academic practice requires including the means, the test statistic (t), the degrees of freedom (df), the p-value, and the confidence interval for the mean difference. Below is an exemplary template for reporting these findings:

A paired t-test was conducted on a sample of 12 automobiles to empirically evaluate whether a new fuel treatment introduced a difference in mean miles per gallon (mpg).

Results showed that the mean mpg was **statistically significantly** different between the two groups. Specifically, the treatment group (M = 22.75) showed a higher average mpg compared to the untreated group (M = 21.00). The inferential statistics supported this finding (t = -2.2444 with df = 11, p = .0463) when assessed against a significance level of 0.05.

Furthermore, a 95% confidence interval for the true population difference in means resulted in the interval of (-3.466, -0.034). Since this interval does not contain zero, it reinforces the rejection of the null hypothesis.

Based on this comprehensive analysis, it is concluded that the new fuel treatment leads to a statistically significantly higher mpg for cars, demonstrating its effectiveness.