Learn How to Conduct a One Sample t-test in Stata

Authored by

Mohammed looti

November 8, 2025

The One Sample t-test represents a foundational statistical procedure. Researchers utilize this powerful tool to rigorously determine if an unknown population mean significantly deviates from a specific, hypothesized value. This test is indispensable across numerous quantitative disciplines, providing an objective and reliable method for evaluating hypotheses about central tendencies, particularly when the population's standard deviation remains unknown.

This comprehensive article provides an expert, step-by-step guide on how to accurately conduct, analyze, and formally report a One Sample t-test using the sophisticated statistical software package, Stata. We will navigate the entire analytical workflow, beginning with data preparation and loading, proceeding through the command execution, and concluding with a professional interpretation of the generated statistical output.

Theoretical Foundation of the One Sample t-Test

Statistical hypothesis testing is the core mechanism used to draw robust inferences about a larger population based solely on data collected from a smaller sample. The one-sample t-test is specifically engineered for scenarios involving a single random sample, where the objective is to compare the calculated sample mean against a known or assumed value, often denoted as mu_0. The validity of the test rests on certain underlying assumptions, primarily that the data are drawn from an approximately normally distributed population, or alternatively, that the sample size is sufficiently large to invoke the tenets of the Central Limit Theorem.

The primary function of this test is the evaluation of two mutually exclusive statements: the null hypothesis ($text{H}_0$) and the alternative hypothesis ($text{H}_text{a}$). The $text{H}_0$ always posits that no difference exists; formally, it states that the true population mean (mu) is precisely equal to the hypothesized value ($mu = mu_0$). Conversely, the alternative hypothesis asserts that a genuine, statistically significant difference exists ($mu neq mu_0$). For directional tests, the alternative hypothesis may specify that the mean is either greater than ($mu > mu_0$) or less than ($mu < mu_0$) the hypothesized value.

The decision to accept or reject the null hypothesis hinges on the computed test statistic, or the t-value. This value quantifies the magnitude of the difference between the observed sample mean and the hypothesized population mean, relative to the standard error of the mean. Once calculated, this t-value is compared against a critical value derived from the t-distribution. This comparison is mediated by the specified degrees of freedom (df), allowing the researcher to determine whether the observed deviation is likely attributable to random chance or if it represents a systematic and genuine effect within the population being studied.

Case Study Implementation: Automotive Fuel Efficiency

To demonstrate the practical utility of the one-sample t-test, we utilize a realistic scenario involving quality control in automotive manufacturing. A team of engineers seeks rigorous confirmation of a long-standing industry assumption: that a specific model of automobile achieves an average fuel efficiency of precisely 20 miles per gallon (mpg). While acknowledging inherent variation in individual vehicle performance, the engineers must establish whether the true underlying average for the **entire population** of these vehicles statistically differs from 20 mpg. To gather evidence, they conduct controlled testing on a randomly selected sample of 74 automobiles.

The specific objective is to execute a One Sample t-test to formally ascertain if the population's true average mpg is statistically equivalent to 20 mpg. A statistically significant finding would necessitate the immediate revision of the prevailing industry standard assumption. Because the researchers are interested in whether the mean is simply **different** from 20 (it could be higher or lower), this mandates the use of a two-tailed test.

The statistical hypotheses guiding this investigation are formally structured as follows:

H_0 (Null Hypothesis): The true population mean miles per gallon is exactly 20 ($\mu = 20$).
H_a (Alternative Hypothesis): The true population mean miles per gallon is not equal to 20 ($\mu \neq 20$).

For the decision rule, the researchers have adopted a standard significance level (α) of 0.05. This acceptance threshold means they are prepared to accept a 5% probability of incorrectly rejecting the null hypothesis when it is, in fact, true--a statistical error known as a Type I error.

Setting Up the Analysis and Data Preparation in Stata

Before any statistical computations can commence, the data must be correctly loaded and verified within the software environment. Stata provides a robust platform for managing and analyzing data, and we will begin by accessing the relevant built-in automotive dataset, which contains the necessary fuel efficiency measurements for our case study.

Step 1: Loading the Official Stata Data

To access the required data, specifically the official automobile dataset, users should type the following command directly into the Stata command window and press Enter. This procedure loads the dataset, which is hosted on the Stata Press website, ensuring access to a reliable, standardized set of observations:

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

Upon successful execution of this command, Stata will confirm that 74 observations (representing

74 cars) and several variables, including our key variable of interest, $text{mpg}$ (miles per gallon), have been successfully loaded into the working memory.

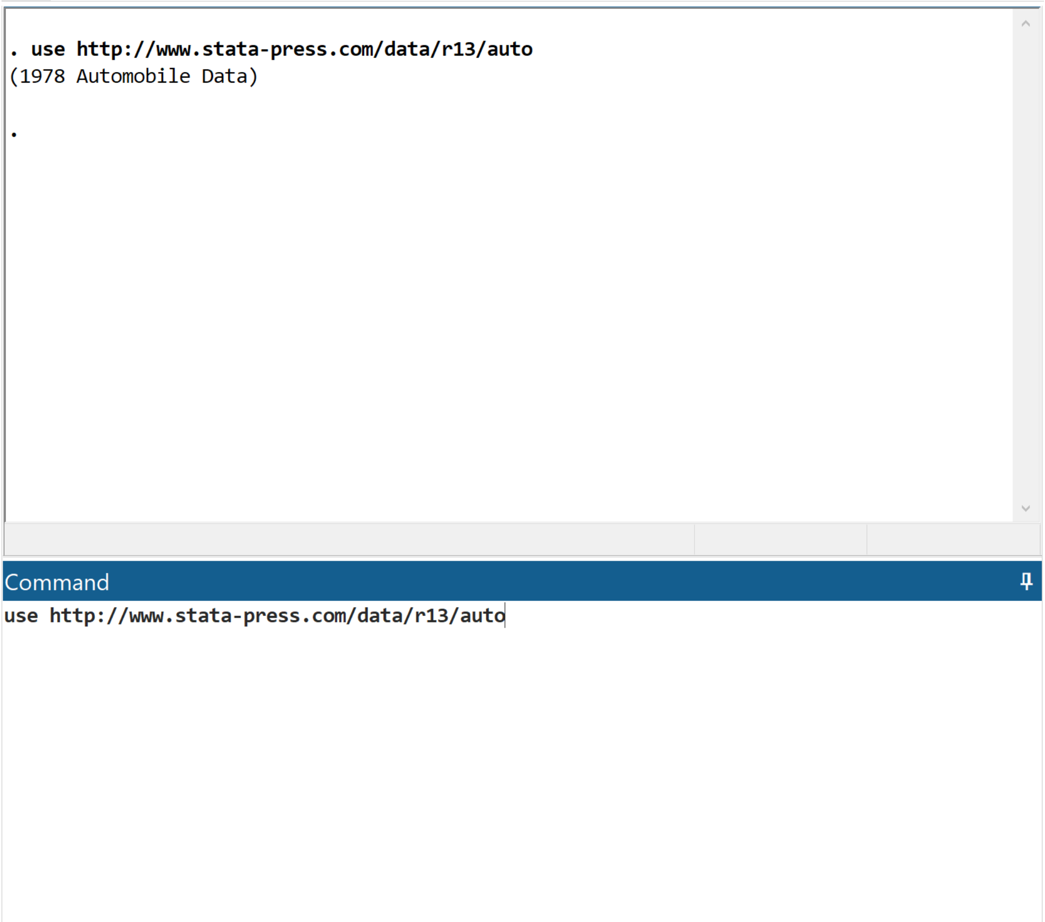

Step 2: Inspecting the Raw Data for Integrity

A crucial preliminary step in any analysis is the visual inspection of the raw data. This step confirms the data's structural integrity, verifies that variables are correctly formatted, and familiarizes the analyst with the sample values. In the Stata interface, navigate to the top menu bar: **Data > Data Editor > Data Editor (Browse)**. This action opens the dataset in a spreadsheet format. While the dataset contains various metrics (e.g., price, weight), our focus is exclusively on the $text{mpg}$ column, which contains the 74 numerical values that constitute our sample data for the t-test:

	make	price	mpg	rep78	headroom	trunk	weight	length	turn	displacement	gear_ratio	foreign
1	AMC Concord	4,099	22	3	2.5	11	2,930	186	40	121	3.58	Domestic
2	AMC Pacer	4,749	17	3	3.0	11	3,350	173	40	258	2.53	Domestic
3	AMC Spirit	3,799	22	.	3.0	12	2,640	168	35	121	3.08	Domestic
4	Buick Century	4,816	20	3	4.5	16	3,250	196	40	196	2.93	Domestic
5	Buick Electra	7,827	15	4	4.0	20	4,080	222	43	350	2.41	Domestic
6	Buick LeSabre	5,788	18	3	4.0	21	3,670	218	43	231	2.73	Domestic
7	Buick Opel	4,453	26	.	3.0	10	2,230	170	34	304	2.87	Domestic
8	Buick Regal	5,189	20	3	2.0	16	3,280	200	42	196	2.93	Domestic
9	Buick Riviera	10,372	16	3	3.5	17	3,880	207	43	231	2.93	Domestic
10	Buick Skylark	4,082	19	3	3.5	13	3,400	200	42	231	3.08	Domestic
11	Cad. Deville	11,385	14	3	4.0	20	4,330	221	44	425	2.28	Domestic
12	Cad. Eldorado	14,500	14	2	3.5	16	3,900	204	43	350	2.19	Domestic
13	Cad. Seville	15,906	21	3	3.0	13	4,290	204	45	350	2.24	Domestic
14	Chev. Chevette	3,299	29	3	2.5	9	2,110	163	34	231	2.93	Domestic
15	Chev. Impala	5,705	16	4	4.0	20	3,690	212	43	250	2.56	Domestic
16	Chev. Malibu	4,504	22	3	3.5	17	3,180	193	31	200	2.73	Domestic
17	Chev. Monte Carlo	5,104	22	2	2.0	16	3,220	200	41	200	2.73	Domestic
18	Chev. Monza	3,667	24	2	2.0	7	2,750	179	40	151	2.73	Domestic
19	Chev. Nova	3,955	19	3	3.5	13	3,430	197	43	250	2.56	Domestic
20	Dodge Colt	3,984	30	5	2.0	8	2,120	163	35	98	3.54	Domestic
21	Dodge Diplomat	4,010	18	2	4.0	17	3,600	206	46	318	2.47	Domestic
22	Dodge Magnum	5,886	16	2	4.0	17	3,600	206	46	318	2.47	Domestic
23	Dodge St. Regis	6,342	17	2	4.5	21	3,740	220	46	225	2.94	Domestic
24	Ford Fiesta	4,389	28	4	1.5	9	1,800	147	33	98	3.15	Domestic
25	Ford Mustang	4,187	21	3	2.0	10	2,650	179	43	140	3.08	Domestic
26	Linc. Continental	11,497	12	3	3.5	22	4,840	233	51	400	2.47	Domestic
27	Linc. Mark V	13,594	12	3	2.5	18	4,720	230	48	400	2.47	Domestic
28	Linc. Versailles	13,466	14	3	3.5	15	3,830	201	41	302	2.47	Domestic
29	Merc. Bobcat	3,829	22	4	3.0	9	2,580	169	39	140	2.73	Domestic
30	Merc. Cougar	5,379	14	4	3.5	16	4,060	221	48	302	2.75	Domestic
31	Merc. Marquis	6,165	15	3	3.5	23	3,720	212	44	302	2.26	Domestic
32	Merc. Monarch	4,516	18	3	3.0	15	3,370	198	41	250	2.43	Domestic
33	Merc. XR-7	6,303	14	4	3.0	16	4,130	217	45	302	2.75	Domestic

Visual confirmation ensures that the $text{mpg}$ variable is correctly recognized as a continuous, numerical input, making it appropriate for the subsequent t-test procedure.

Executing the t-Test and Interpreting Stata Output

Once the data is securely loaded and its structure verified, the next phase involves running the statistical test itself. We will utilize Stata's intuitive graphical user interface (GUI) to specify the test parameters: the variable to be tested, the type of comparison, and the hypothesized value against which the sample mean is being measured.

Step 3: Performing the One Sample t-test via GUI

To initiate the test dialogue box, follow this menu sequence: **Statistics > Summaries, tables, and tests > Classical tests of hypotheses > t test (mean-comparison test)**. Within the resulting dialogue box, precise settings must be applied to ensure the correct execution of our two-tailed test:

The primary test type must be confirmed as **One-sample** (typically the default setting).
For the **Variable name** field, select the variable of interest, $text{mpg}$.
In the **Hypothesized mean** (mu_0) input box, enter the value $mathbf{20}$, corresponding to the industry standard assumption.
The **Confidence level** should be maintained at the standard 95%, which aligns directly with our chosen significance level ($alpha$) of 0.05.

After setting these required parameters, clicking the **OK** button executes the command and immediately generates the statistical results in the Stata output window.

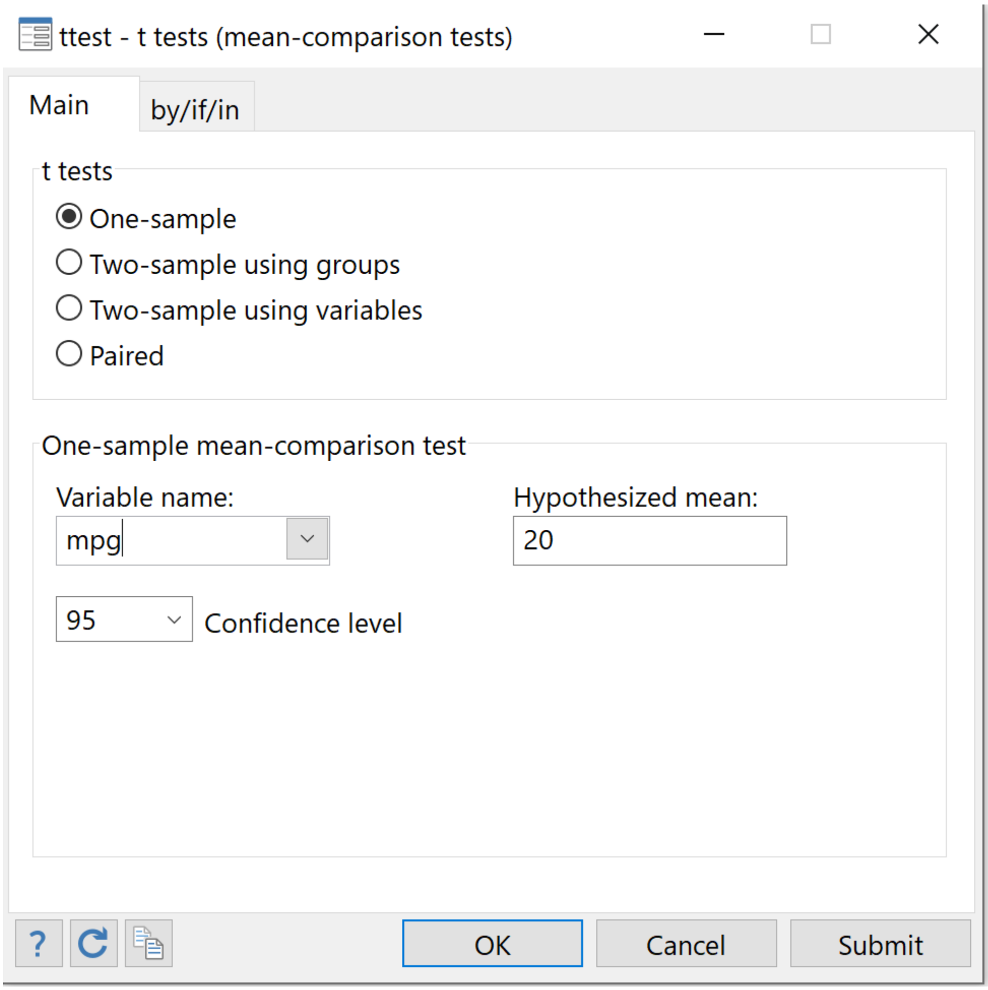

The resulting output window contains all the necessary quantitative metrics required for making a formal decision regarding the null hypothesis. The structured output should resemble the image below:

```
One-sample t test

Variable |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
     mpg |      74     21.2973     .6725511    5.785503     19.9569    22.63769
------------------------------------------------------------------------------
    mean = mean(mpg)                                              t =   1.9289
Ho: mean = 20                                    degrees of freedom =       73

    Ha: mean < 20               Ha: mean != 20                Ha: mean > 20
 Pr(T < t) = 0.9712         Pr(|T| > |t|) = 0.0576         Pr(T > t) = 0.0288
```

Detailed Interpretation of Key Output Metrics:

Obs: This denotes the sample size ($text{n}$), which is 74 cars in this study.
Mean: This is the calculated sample mean ($bar{x}$) for the $text{mpg}$ variable. The average fuel efficiency observed in our sample is **21.2973** miles per gallon.
Std. Err: This crucial value is the Standard Error of the mean, reflecting the variability expected in the sample means if we were to repeat the sampling process multiple times. It is calculated as the Standard Deviation divided by the square root of $text{n}$, yielding approximately **0.6725511**. This metric is critical for assessing the precision of the sample mean as an estimate of the true population mean.
Std. Dev: This is the sample standard deviation ($sigma$), which quantifies the spread or average deviation of individual data points from the sample mean, found to be **5.785503**.
95% Conf. Interval: This represents the 95% Confidence Interval for the true population mean (mu). The interval (19.9569, 22.63769) suggests that we are 95% confident that the true average mpg for the entire population of these cars falls within this range.
t: This is the calculated test statistic. It measures how many standard errors the sample mean (21.2973) is away from the hypothesized mean (20). The calculated value is **1.9289**.
degrees of freedom (df): This is calculated as $text{n} - 1 = 74 - 1 = mathbf{73}$. The degrees of freedom are essential for correctly referencing the critical threshold within the t-distribution.

The conclusive decision in hypothesis testing relies on the bottom section of the Stata output, which provides the p-values corresponding to the three possible alternative hypotheses (left-tailed, two-tailed, and right-tailed). Given that our research question was non-directional--asking if the mean is **different** from 20--we must focus on the middle result, labeled $text{"Ha}: text{mean} neq 20$." The associated $text{p-value}$ is $mathbf{0.0576}$.

We now compare this derived p-value (0.0576) against our predefined significance level ($alpha = 0.05$). Since 0.0576 is greater than 0.05, the result is deemed statistically insignificant.

Consequently, we are compelled to **fail to reject the null hypothesis**. This critical finding indicates that the data does not provide sufficient statistical evidence to confidently conclude that the true mean mpg for this vehicle population is anything other than 20 mpg.

Drawing Conclusions and Formal Reporting Standards

Translating the statistical outcome--the decision to fail to reject the null hypothesis--back into the context of the initial research question is the final step in the analytical process. Failing to reject the null hypothesis implies that the observed discrepancy between the sample mean ($\bar{x} = 21.2973$ mpg) and the hypothesized mean ($\mu_0 = 20$ mpg) is small enough that it can be reasonably explained by inherent random sampling variability rather than a true population effect.

The 95% Confidence Interval further substantiates this conclusion. Stata calculated this interval as (19.9569, 22.63769). Since the hypothesized value of 20 mpg is securely contained within this estimated range, it strongly reinforces the decision to retain the null hypothesis. Had the hypothesized value fallen outside of this confidence interval, we would have been statistically justified in rejecting the null hypothesis.

Step 4: Standardized Formal Reporting of Findings

In academic literature or professional reports, statistical results must be communicated using clear, standardized notation. A complete report of a t-test must include the type of test, the sample size (n), the calculated test statistic (t), the associated degrees of freedom (df), and the resulting p-value.

The following template illustrates the precise format for formally reporting the findings of this One Sample t-test:

A One Sample t-test was executed on a sample of 74 automobiles to evaluate whether the true population mean miles per gallon differed statistically from the hypothesized value of 20 mpg.

The analysis indicated that the true population mean was **not** significantly different from 20 mpg ($t = 1.9289$ with $\text{df} = 73$, $p = 0.0576$) when using the $\alpha = 0.05$ significance level. Based on this finding, we fail to reject the null hypothesis.

Furthermore, the 95% Confidence Interval calculated for the true population mean ranged from 19.9569 to 22.63769. The fact that the hypothesized value of 20 mpg falls within this interval confirms the decision to retain the null hypothesis.