

How to Easily Perform a Paired Samples t-Test in SAS

Authored by
stats writer

December 1, 2025

RECOMMENDED CITATION

stats writer (2025). *How to Easily Perform a Paired Samples t-Test in SAS*.

PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=103394>

A Paired Samples t-Test is a fundamental statistical test utilized when researchers need to compare the means of two variables that are inherently related or dependent. Unlike the independent samples t-test, this methodology requires that each observation in the first group corresponds directly to an observation in the second group, often representing measurements taken on the same subjects under different conditions (e.g., pre-test and post-test scores, or measurements taken before and after an intervention). Performing this analysis efficiently and reliably requires powerful software, and the SAS (Statistical Analysis System) platform is widely recognized for its capabilities in this domain. This comprehensive guide details the necessary steps and commands within SAS to execute a valid Paired Samples t-Test, ensuring accurate interpretation of the results.

To execute the test within the SAS environment, the primary tool is the **PROC TTEST** statement. This procedure is specifically designed for t-tests and facilitates the comparison of means. To define the relationship between the two variables, the **PAIRED** statement must be employed, explicitly telling SAS which two variables form the related pair. While the core test only requires the **PROC TTEST** and **PAIRED** statements, leveraging auxiliary statements like **VAR** (to list general variables) or **MEANS** (to generate preliminary descriptive statistics) can significantly enhance the initial exploratory data analysis.

Upon successful execution of the SAS code, the output will yield critical results necessary for hypothesis testing. Key statistical elements provided include the calculated **t-statistic**, the relevant degrees of freedom, and, most importantly for decision-making, the associated p-value (probability value). Understanding how to correctly locate and interpret these metrics is essential for drawing a statistically sound conclusion regarding the differences, or lack thereof, between the paired measurements.

Understanding the Paired Samples t-Test

The Paired Samples t-Test, sometimes referred to as the dependent t-test or paired-difference test, is fundamentally concerned with the mean difference between two sets of scores. The defining characteristic is the dependency between the samples. This usually arises in experimental designs where measurements are taken repeatedly on the same participants, such as longitudinal studies, crossover trials, or simple before-and-after designs. The primary goal is not to compare the absolute means of the two groups, but rather to assess whether the mean of the differences between the paired observations is significantly different from zero.

The statistical theory behind the Paired Samples t-Test simplifies the problem into a single-sample t-test applied to the difference scores ($D = X_2 - X_1$). By converting the two related variables into a single variable representing the change, the test effectively minimizes the impact of inter-subject variability. This makes the paired design often more powerful than an independent samples design.

when addressing within-subject effects, as it controls for individual differences in baseline performance or characteristics.

Choosing the correct statistical tool is crucial for valid research. If the data consists of two measurements where each observation in the first measurement corresponds uniquely to an observation in the second, the paired t-test is appropriate. If the samples were drawn independently (e.g., comparing two completely separate groups who received different treatments), an independent samples t-test would be required. Misclassifying the relationship between samples can lead to incorrect calculations of standard error and ultimately, flawed conclusions about the intervention's effect.

When to Use a Paired Samples t-Test

Determining the appropriate context for utilizing a Paired Samples t-Test hinges on recognizing the underlying structure of the data collection process. This test is suitable primarily for within-subjects designs. Common scenarios include situations where an experimental manipulation is applied, and the outcome is measured both before and after the manipulation on the same set of individuals. For instance, testing the effectiveness of a new medication by measuring patients' blood pressure before and after treatment, or assessing the impact of specialized training on employee productivity measured pre- and post-training.

Beyond intervention studies, the test is also frequently used in situations where natural pairing occurs. Examples include comparing married couples (husband vs. wife scores on a satisfaction scale), twins (comparing performance under two different conditions), or comparing machine performance using two different methods on the exact same material batch. The consistency of the pairing ensures that the variance attributed to extraneous factors (individual differences) is reduced, thereby increasing the statistical power to detect true effects related to the variable of interest.

It is important to ensure that the assumptions underlying the Paired Samples t-Test are met before proceeding with interpretation. The two main assumptions are that the paired differences must be approximately normally distributed, and that the observations within each pair must be independent of the other pairs. While the t-test is relatively robust to minor deviations from normality, particularly with larger sample sizes, extreme skewness or heavy tails in the difference scores may necessitate non-parametric alternatives, such as the Wilcoxon Signed-Rank Test.

Case Study: Evaluating a Study Program

To illustrate the application of this test in SAS, let us consider a practical research scenario. A university professor seeks to evaluate the efficacy of a newly developed, intensive study program intended to boost student performance in a challenging statistics course. To rigorously assess the

program's impact, the professor employs a classic pre-test/post-test design. Fifteen randomly selected students are subjected to this process.

The methodology involves two key stages. First, all 15 students complete a standardized pre-test designed to assess their baseline knowledge. Then, he has each student use the study program for one month, followed by a post-test of similar difficulty. By comparing the scores of the pre-test (measurement 1) against the scores of the post-test (measurement 2) for each individual student, we can determine if the study program induced a significant change in performance.

Since each student generates a pair of scores (Pre, Post), the analysis demands a Paired Samples t-Test. The data collected summarizes the scores for each of the 15 participants, allowing for a direct comparison of the mean performance before and after the intervention. The goal is to statistically determine if the observed increase in scores post-program is large enough to be considered meaningful rather than simply due to random variation. The test scores for each of the 15 students are shown below:

Student	Pre-Test Score	Post-Test Score
1	88	91
2	82	84
3	84	88
4	93	90
5	75	79
6	78	80
7	84	88
8	87	90
9	95	90
10	91	96
11	83	88
12	89	89
13	77	81
14	68	74
15	91	92

The professor's rationale is clear: if the study program is effective, the mean post-test score should be statistically higher than the mean pre-test score. By calculating the difference score for every student (Post minus Pre) and performing the t-test on those differences, we can isolate the effect of the program, controlling for the inherent variation in student ability.

Step 1: Defining and Inputting the Data in SAS

The initial and most crucial step in any SAS analysis is the accurate creation and input of the dataset. For this case study, we define a dataset named **test_scores**, which will contain two essential variables: **pre** (pre-test score) and **post** (post-test score). Data creation in SAS typically involves the **DATA** statement followed by the **INPUT** statement to define the variables and type of data being read, and finally the **DATALINES** statement to enter the raw observations directly into the program.

The following syntax demonstrates how to structure and input the scores for the 15 students. Note the use of comments (preceded by `/*`) to explain the code blocks, a best practice for readability and reproducibility. Each line following `datalines;` represents one student's paired scores.

```
/*create dataset*/
```

```
data test_scores;
```

```
input pre post;
```

```
datalines;
```

```
88 91
```

```
82 84
```

```
84 88
```

```
93 90
```

```
75 79
```

```
78 80
```

```
84 88
```

```
87 90
```

```
95 90
```

```
91 96
```

```
83 88
```

```
89 89
```

```
77 81
```

```
68 74
```

```
91 92
```

```
;
```

```
run;
```

```
/*view dataset*/
```

```
proc print data=test_scores;
```

Once the data is successfully entered using the **DATALINES** block and terminated by the **RUN** statement, it is prudent to verify the input. The use of **PROC PRINT** ensures that the data table is

correctly loaded and structured within the SAS environment, confirming that the pairing (row by row) is accurate before proceeding to the actual statistical computation.

Reviewing the Initial Dataset Structure

Verification is a mandatory step in statistical programming to prevent errors arising from faulty data entry. Running the **PROC PRINT** command on the newly created `test_scores` dataset provides a clean visual confirmation of the data structure. This output confirms that we have 15 observations, each with a corresponding 'pre' and 'post' score, maintaining the integrity of the paired design necessary for the t-test.

The output table generated by SAS ensures that the data is ready for analysis, demonstrating the two variables side-by-side:

Obs	pre	post
1	88	91
2	82	84
3	84	88
4	93	90
5	75	79
6	78	80
7	84	88
8	87	90
9	95	90
10	91	96
11	83	88
12	89	89
13	77	81
14	68	74
15	91	92

This structure is crucial because the subsequent analysis relies entirely on the assumption that the `pre` score in row N belongs to the same individual as the `post` score in row N. Having confirmed the structure, we can confidently move forward to the execution of the Paired Samples t-Test.

Step 2: Executing the Paired Samples t-Test using PROC TTEST

The core calculation for the Paired Samples t-Test is performed using the robust **PROC TTEST** procedure in SAS. This procedure is designed to handle various types of t-tests, and the specific

pairing is established using the **PAIRED** statement. When using the **PAIRED** statement, the syntax requires specifying the two variables separated by an asterisk (e.g., `pre*post`), indicating that SAS should calculate the difference scores and test if the mean of these differences is zero.

For enhanced control and reporting, we can also specify the **ALPHA** level, which sets the significance threshold (commonly 0.05) for calculating the confidence interval. The following SAS code block executes the Paired Samples t-Test on our `test_scores` data:

```
/*perform paired samples t-test*/
proc ttest data=test_scores alpha=.05;
paired pre*post;
run;
```

The result of running this procedure is a detailed output table that includes descriptive statistics, the confidence interval for the mean difference, and the formal test results. The order specified in the **PAIRED** statement (`pre*post`) dictates the direction of the difference calculation (pre-test score minus post-test score). This is important for interpreting the sign of the mean difference.

Detailed Interpretation of the SAS Output

The output generated by **PROC TTEST** provides several crucial pieces of information necessary to evaluate the professor's study program. This output is divided into sections, offering statistics on the differences, the confidence interval, and the formal test results.

The TTEST Procedure					
Difference: pre - post					
N	Mean	Std Dev	Std Err	Minimum	Maximum
15	-2.3333	3.0394	0.7848	-6.0000	5.0000

Mean	95% CL Mean	Std Dev	95% CL Std Dev
-2.3333	-4.0165 -0.6502	3.0394	2.2252 4.7935

DF	t Value	Pr > t
14	-2.97	0.0101

The first set of results focuses on the descriptive statistics of the difference scores (Pre - Post):

Mean difference between pre-test and post-test score: **-2.3333**. Since the calculation was (Pre -

Post), a negative mean indicates that the Post scores were generally higher than the Pre scores, suggesting an improvement.

95% Confidence Interval for Mean difference: . Crucially, because the entire interval is negative and does not contain zero, this strongly suggests a statistically significant difference between the two measurement periods.

The subsequent section provides the formal results of the hypothesis test, including the calculated test statistic and the corresponding two-sided p-value:

t test statistic: **-2.97**. This statistic measures how many standard errors the mean difference is away from zero.

p-value: **.0101**. This is the probability of observing a mean difference as extreme as -2.3333 (or more extreme) if the true population mean difference were actually zero.

The low p-value is the key to making a statistical decision, providing the evidence needed to accept or reject the null hypothesis.

Drawing Statistical Conclusions and Reporting Results

To conclude the analysis, we must formalize the null and alternative hypotheses being tested by the Paired Samples t-Test. This procedure evaluates whether the average difference score across the population is truly zero. Using the defined difference (Pre - Post):

H₀: The true population mean difference is zero. This implies that the mean pre-test and post-test scores are equal, meaning the study program had no effect.

H_A: The true population mean difference is not zero. This implies that the mean pre-test and post-test scores are significantly different, meaning the study program had an effect.

The decision rule in statistical testing mandates that if the p-value (**0.0101**) is less than the predetermined significance level (\$alpha\$, set here at 0.05), we reject the null hypothesis. Since $0.0101 < 0.05$, we reject H₀.

This statistical decision translates into a meaningful conclusion: there is sufficient evidence, based on the sample data, to state that the true mean test score is significantly different for students after participating in the study program compared to their scores before the program. Furthermore, since the mean difference was negative (-2.3333), indicating that Post > Pre, we conclude specifically that the study program led to a statistically significant improvement in test scores.

Key Assumptions for the Paired Samples t-Test

While the execution of the SAS code is straightforward, the validity of the results hinges on certain underlying statistical assumptions. Failure to meet these assumptions can undermine the reliability

of the p-value and confidence intervals. Researchers must verify these assumptions before finalizing their conclusions.

The key assumptions are:

Dependent Observations: The primary assumption, already ensured by the study design, is that the observations are paired or dependent. Each data point in one group must correspond directly to a specific data point in the second group.

Independence of Pairs: While observations within a pair are dependent, the pairs themselves must be independent of each other. The change experienced by Student 1 must not influence the change experienced by Student 2.

Normality of Differences: The differences between the paired observations ($D = X_2 - X_1$) must be approximately normally distributed in the population. SAS provides output related to normality testing within the PROC TTEST results, which should be reviewed, especially for small sample sizes ($N < 30$). If the difference scores deviate significantly from normality, the t-test results may be unreliable, suggesting the use of non-parametric methods.

By carefully checking these conditions, the interpretation drawn from the Paired Samples t-Test output in SAS can be considered robust and trustworthy, supporting evidence-based decision-making regarding the efficacy of interventions or treatments.

Further Statistical Applications in SAS

Mastery of the Paired Samples t-Test in SAS opens the door to performing a wide variety of other common statistical test. The structure and principles utilized--data input via DATA steps and analysis via PROC procedures--are consistent across most analyses in the platform. Expanding your toolkit allows for comprehensive data exploration and hypothesis testing in diverse research contexts.

The following tutorials provide guidance on implementing other foundational statistical analyses commonly executed within the SAS environment, building upon the skills developed through mastering the Paired Samples t-Test procedure: