

# Learning Least Squares Means (LSMEANS) in SAS for ANOVA: A Step-by-Step Guide

Authored by  
**Mohammed looti**

November 14, 2025

## RECOMMENDED CITATION

Mohammed looti (2025). *Learning Least Squares Means (LSMEANS) in SAS for ANOVA: A Step-by-Step Guide*. PSYCHOLOGICAL STATISTICS. Retrieved from <https://statistics.arabpsychology.com/?p=1475>

## The Role of Least Squares Means and ANOVA in Comparative Analysis

When statistical practitioners aim to compare the average outcomes across three or more independent groups, the analysis typically begins with the [One-Way Analysis of Variance \(ANOVA\)](#). This robust foundational tool determines the overall probability that the population means of all groups under study are statistically equal. ANOVA serves as a critical first step, establishing whether the factor being investigated has any significant general effect. However, a crucial limitation of the ANOVA test is that while it confirms the existence of a statistically significant difference among the group means, it deliberately does not specify precisely which pairs of groups are distinct from one another.

The core mechanism of ANOVA involves partitioning the total variance observed in the data. It contrasts the variability found between the different groups with the inherent variability found within the groups themselves. If this comparison yields a sufficiently large F-statistic, and the corresponding p-value falls below a predefined [significance level](#) (conventionally set at  $\alpha = 0.05$ ), we are compelled to reject the [null hypothesis](#). Rejecting the null hypothesis confirms that the factor has a significant influence, implying that at least one group mean is different from the others. This conclusion, however, merely signals the need to transition to a more detailed, granular phase of inspection to locate the exact source of these differences.

This necessity for pinpointing specific differences mandates the use of [post-hoc tests](#), also known as multiple comparison procedures. These procedures are indispensable because they allow for systematic pairwise comparisons while simultaneously controlling the critical risk of inflating the family-wise error rate--specifically, the likelihood of committing [Type I errors](#) (false positives) across the entire set of comparisons. Within the widely utilized [SAS](#) statistical software environment, the powerful **LSMEANS** statement provides an essential, flexible, and robust mechanism for executing these controlled post-hoc comparisons, enabling analysts to accurately identify the specific loci of significant group distinctions.

## Understanding the Need for Controlled Multiple Comparisons

When an initial ANOVA result indicates an overall significant effect, researchers must proceed with caution. It is statistically unsound to jump to conclusions about specific group comparisons (e.g., comparing Group A versus Group B) solely based on the general significant result. The ANOVA model is fundamentally designed to test the global hypothesis of equality across all groups, not the individual differences between pairs. Without proper adjustment, performing multiple independent t-tests substantially increases the probability of mistakenly declaring a non-existent difference as significant. This accumulated risk is referred to as the family-wise error rate; for instance, conducting three separate comparisons at the standard  $\alpha = 0.05$  level results in a family-wise error rate significantly exceeding 5%.

The **LSMEANS** statement in SAS is specifically engineered to address this challenge by estimating and comparing [Least Squares Means \(LSMEANS\)](#). LSMEANS, often referred to as population marginal means, represent the predicted means for the levels of a factor within a statistical model. Crucially, these means are adjusted for the presence of any other factors or covariates included in the model. In the simplest scenario of a basic one-way ANOVA, where no other factors are present, LSMEANS are numerically equivalent to the standard arithmetic means. However, their true statistical value becomes apparent in more complex experimental designs, such as ANCOVA or factorial ANOVA, where they ensure unbiased comparisons by standardizing and accounting for potential confounding variables.

By deploying the **LSMEANS** statement for post-hoc analysis, researchers ensure that comparisons are conducted between these standardized, adjusted means. Furthermore, the statement is designed to integrate various statistical methods for rigorously controlling the family-wise error rate. Utilizing options like Tukey's adjustment guarantees that the probability of making at least one false significant finding across the entire battery of comparisons remains precisely controlled at the designated alpha level, thereby safeguarding the statistical integrity and reliability of the research findings.

## Case Study Setup: Evaluating Studying Techniques

To demonstrate the practical application of the **LSMEANS** statement, we will analyze a hypothetical case study designed to evaluate the differential impact of various studying techniques on academic performance. The experiment involves a cohort of 24 students who are randomly assigned to one of three distinct studying methods: Method A, Method B, or Method C. After a controlled preparation period, each student takes a standardized exam, and their resulting score is recorded. The objective of this experimental design is to determine if the selection of a studying method leads to a statistically significant difference in the resulting mean exam scores.

The raw data collected from this experiment, detailing the individual exam scores categorized by their assigned studying method, is presented visually below. This serves as the foundational data structure for our subsequent SAS analysis:

Method A	Method B	Method C
78	81	84
81	83	88
82	83	88
82	85	89
85	86	90
88	88	93
88	90	95
90	91	98

The essential initial step in conducting this analysis within the SAS environment is the creation of a structured dataset. The following code snippet illustrates the process of defining and populating a dataset named `my_data`. This dataset requires two essential variables: the categorical independent variable, `Method` (defined using the dollar sign, \$, to indicate a character variable), and the continuous dependent variable, `Score` (representing the quantitative exam performance):

```
/*create dataset*/  
data my_data;  
input Method $ Score;  
datalines;  
A 78  
A 81  
A 82  
A 82  
A 85  
A 88  
A 88  
A 90  
B 81  
B 83  
B 83  
B 85  
B 86  
B 88  
B 90  
B 91  
C 84
```

```
C 88  
C 88  
C 89  
C 90  
C 93  
C 95  
C 98  
;  
run;
```

## Executing the Initial ANOVA and Interpreting Global Significance

Once the data is successfully loaded into the SAS environment, the primary analytical task is to perform the one-way ANOVA test. Given that our experimental design features an equal number of observations in each of the three groups (eight students per method), this is classified as a balanced design. Consequently, the most efficient and appropriate procedure in SAS for this analysis is **PROC ANOVA**. This procedure expertly partitions the total variance in the exam scores, separating the variance attributed to the differences between the study methods from the variance attributed to residual error.

The SAS syntax provided below initiates the ANOVA process. The `CLASS` statement is mandatory, explicitly designating `Method` as the classification (independent grouping) variable. Following this, the `MODEL` statement specifies the statistical relationship being tested: the dependent variable `Score` is modeled as a function of the categorical independent variable `Method`:

```
/*perform one-way ANOVA*/  
proc ANOVA data=my_data;  
class Method;  
model Score = Method;  
run;
```

The execution of this code yields the standard SAS output, the central component of which is the ANOVA Summary Table. This table is indispensable for evaluating the overall statistical significance of the model, clearly summarizing how the total variance is allocated across the various sources of variation:

## The ANOVA Procedure

Dependent Variable: Score

Source	DF	Sum of Squares	Mean Square	F Value	Pr > F
Model	2	175.5833333	87.7916667	5.26	0.0140
Error	21	350.2500000	16.6785714		
Corrected Total	23	525.8333333			

R-Square	Coeff Var	Root MSE	Score Mean
0.333914	4.698685	4.083941	86.91667

Source	DF	Anova SS	Mean Square	F Value	Pr > F
Method	2	175.5833333	87.7916667	5.26	0.0140

A careful examination of the summary table reveals two key metrics essential for our decision-making: the overall **F Value**, calculated in this case as **5.26**, and its corresponding p-value, reported as **0.0140**. Based on these statistics, we formally evaluate our hypotheses concerning the study methods:

**Null Hypothesis (H<sub>0</sub>):** The mean exam scores for Method A, Method B, and Method C are statistically equal.

**Alternative Hypothesis (H<sub>A</sub>):** At least one studying method produces a mean exam score that is significantly different from the others.

Since the observed p-value (**0.0140**) is unambiguously lower than the standard statistical threshold of  $\alpha = 0.05$ , we possess sufficient statistical evidence to confidently reject the null hypothesis. This rejection confirms that the choice of studying method significantly influences the mean exam score. However, as established earlier, the ANOVA result only validates the existence of an overall effect; it does not yet clarify which specific pairs of methods differ. We must now proceed to the crucial post-hoc stage utilizing the **LSMEANS** statement.

## Implementing Pairwise Comparisons with Tukey's HSD

To transition from the general finding of significance to specific, actionable insights, we must perform controlled pairwise comparisons. While **PROC ANOVA** can handle basic post-hoc requests, we typically transition to **PROC GLIMMIX** in SAS for greater statistical flexibility and access to advanced adjustment methods. Although designed for generalized linear mixed models, **PROC GLIMMIX** is fully capable of running standard linear model analyses while offering superior options, including comprehensive tools for handling multiple comparisons and complex models

involving covariates.

Within **PROC GLIMMIX**, the **LSMEANS** statement is deployed to calculate and compare the least squares means for the levels of the `Method` factor. A critical component here is the selection of the appropriate adjustment method to control the family-wise error rate. For comparing all possible pairs of means following a significant ANOVA, [Tukey's Honestly Significant Difference \(HSD\)](#) test is recognized as the gold standard, offering an optimal balance between statistical power and error control. We implement this adjustment using the `ADJUST=TUKEY` option.

The following SAS code demonstrates the necessary syntax to perform the Tukey post-hoc comparisons, setting the nominal significance level for these comparisons at 0.05:

```
/*perform Tukey post-hoc comparisons*/
proc glimmix data=my_data;
class Method;
model Score = Method;
lsmeans Method / adjust=tukey alpha=.05;
run;
```

Upon execution, SAS generates several tables, culminating in the "Differences of Least Squares Means" table. This output contains the results for the three possible pairwise comparisons (A vs. B, A vs. C, and B vs. C), which are essential for formulating final, specific conclusions about the relative effectiveness of the studying methods:

Differences of Method Least Squares Means Adjustment for Multiple Comparisons: Tukey												
Method	_Method	Estimate	Standard Error	DF	t Value	Pr >  t	Adj P	Alpha	Lower	Upper	Adj Lower	Adj Upper
A	B	-1.6250	2.0420	21	-0.80	0.4351	0.7096	0.05	-5.8715	2.6215	-6.7719	3.5219
A	C	-6.3750	2.0420	21	-3.12	0.0052	0.0137	0.05	-10.6215	-2.1285	-11.5219	-1.2281
B	C	-4.7500	2.0420	21	-2.33	0.0301	0.0739	0.05	-8.9965	-0.5035	-9.8969	0.3969

## Interpreting Adjusted P-Values and Drawing Conclusions

The interpretation of the **LSMEANS** output is entirely dependent on the column designated **Adj P**. This column provides the p-values that have been statistically adjusted using the Tukey method to account for the inherent inflation of error associated with multiple comparisons. It is paramount that analysts utilize these adjusted p-values, as they ensure that the overall probability of incorrectly rejecting the null hypothesis across the entire set of comparisons remains controlled at the designated family-wise error rate ( $\alpha = 0.05$ ). Any comparison where the **Adj P** value falls below 0.05 indicates a statistically significant difference between those two group means.

Analyzing the table above, we derive the following specific statistical outcomes for the three possible comparisons:

**Method A vs. Method B:** The adjusted p-value is 0.4578. As this value is substantially greater than 0.05, we must conclude there is no statistically significant difference in mean exam scores between students using these two studying methods.

**Method A vs. Method C:** The adjusted p-value is **0.0137**. Because this value is below the 0.05 threshold, this comparison is statistically significant. We can confidently conclude that students utilizing Method A performed significantly differently than students utilizing Method C.

**Method B vs. Method C:** The adjusted p-value is 0.1340. This value is greater than 0.05, indicating that there is no statistically significant difference between Method B and Method C.

Focusing on the one significant comparison (A vs. C), the output provides further crucial detail. The estimated difference in means is **-6.375**, which indicates that students in Method A scored, on average, 6.375 points lower than those in Method C. The adjusted 95% [confidence interval](#) for this difference ranges from **-11.5219** to **-1.2281**. Since this interval does not encompass zero, it reinforces the statistical significance and provides a precise estimate of the magnitude of the performance difference. In summary, of the three methods tested, only Method C produced exam scores that were significantly higher than those produced by Method A.

## Advanced Options: Selecting the Appropriate Adjustment Method

While Tukey's method is the benchmark choice for comparing all possible pairs of means, the exceptional flexibility of the **LSMEANS** statement permits researchers to select from a diverse array of alternative adjustment methods. This allows the analysis to be precisely tailored to specific research hypotheses and the characteristics of the data. The choice of adjustment is crucial, as it dictates the inherent trade-off between rigorously controlling Type I errors and maximizing statistical power (the ability to correctly detect a true difference).

Researchers can replace `ADJUST=TUKEY` with one of the following widely recognized options, depending on their analytical objectives:

**BON (Bonferroni):** This procedure is statistically simple and highly conservative. It controls the family-wise error rate by dividing the nominal alpha level by the total number of comparisons. Although effective at error control, its conservatism often leads to reduced statistical power, increasing the risk of Type II errors (failing to detect a real finding).

**SIDAK:** Similar in principle to Bonferroni, the Sidak adjustment is mathematically based on a multiplicative formula that is generally less conservative. It often provides marginally greater statistical power than Bonferroni while maintaining tight control over the family-wise error rate.

**SCHEFFE:** [Scheffé's method](#) is notably robust and universally applicable, particularly valuable when the research involves testing complex, non-pairwise contrasts (linear combinations of

means). For simple pairwise comparisons, however, it is typically the most conservative choice, making it less powerful than the Tukey method for that specific task.

**DUNNETT:** This method is specifically engineered for designs where the primary research interest lies in comparing a set of treatment groups exclusively against a single, designated control group. When this specific goal aligns with the research design, Dunnett's procedure is more powerful than using generalized methods like Bonferroni or Tukey.

The selection criterion for the adjustment method must be guided by whether the comparisons were planned before data collection (a priori) or are exploratory (post hoc), and by the specific statistical penalty the researcher is willing to accept for rigorously controlling false positives. Regardless of the chosen adjustment, the consistent use of the **LSMEANS** statement ensures that the resulting statistical inferences are based on sound, controlled statistical principles.

## Conclusion and Resources for Advanced SAS Usage

The **LSMEANS** statement represents a cornerstone of advanced statistical analysis within the SAS environment, particularly after a significant overall effect has been identified through ANOVA. By providing adjusted estimates of population means and facilitating controlled pairwise comparisons, this function allows researchers to move beyond a generalized finding of group differences toward precise, statistically verified conclusions regarding specific group relationships.

The capacity of **LSMEANS** to seamlessly integrate various multiple comparison adjustments--such as the robust Tukey's HSD--ensures that the subsequent interpretation of results is both powerful and statistically reliable, effectively mitigating the pervasive risks associated with inflated Type I error rates in multiple testing scenarios. Mastery of this statement is indispensable for any statistical practitioner working with complex comparative data, as it ensures that the analytical rigor applied matches the complexity and sensitivity of the experimental design. This detailed example provides a practical, step-by-step template for applying and accurately interpreting this crucial statistical tool in future research endeavors.

## Additional Resources

For those interested in delving deeper into the theoretical and practical aspects of post-hoc tests and other sophisticated applications of ANOVA models, the following tutorial provides valuable supplementary information:

[A Guide to Using Post Hoc Tests with ANOVA](#)