Learning Repeated Measures ANOVA in Stata: A Comprehensive Guide

Authored by

Mohammed looti

November 8, 2025

The Repeated Measures Analysis of Variance (RMANOVA) stands as a cornerstone technique in sophisticated statistical modeling, particularly within experimental and longitudinal research. This powerful method is specifically designed to determine if a statistically significant difference exists among the means of three or more measurements taken from the same group of subjects. The core distinction of RMANOVA lies in its handling of dependent data; by measuring the same individuals across all experimental conditions or time points, it effectively controls for individual differences, which are often the largest source of error variance in research. This intrinsic control not only lends itself perfectly to within-subjects research designs but also substantially enhances the statistical power of the analysis compared to traditional independent measures designs, such as the standard One-Way ANOVA.

Before diving into the practical implementation using **Stata**, it is crucial to appreciate the theoretical advantages offered by the repeated measures design. Because participants serve as their own controls, any observed variability between conditions is less likely to be attributed to pre-existing individual characteristics (like intelligence, background, or physical fitness) and more likely to reflect the true effect of the experimental manipulation or time progression. This efficiency allows researchers to use smaller sample sizes while maintaining robust conclusions, making RMANOVA an invaluable tool when studying behavioral changes, psychological responses, or physiological adaptation over time.

Theoretical Foundations of Repeated Measures ANOVA

Distinguishing the repeated measures design from its independent counterparts, such as the standard One-Way ANOVA, is vital for correct model selection. While the traditional ANOVA compares groups formed by different, unrelated participants (e.g., comparing three distinct groups receiving Drug A, Drug B, or a Placebo), the repeated measures model is explicitly structured for dependency. It is engineered for situations where a single cohort of participants undergoes multiple interventions, conditions, or assessments over a series of time intervals. This methodology's core benefit is its ability to partition variance more effectively. By isolating and removing the variance associated with the subject (individual differences) from the error term, researchers achieve a cleaner, more precise estimation of the experimental treatment effect, thus strengthening the internal validity of the study.

The decision to utilize a one-way repeated measures ANOVA hinges entirely on the underlying structure of the data, which must feature dependent observations. Two primary research contexts necessitate this statistical approach, both rooted in the principle of within-subjects measurement. Grasping these structural patterns is fundamental to applying the statistical model correctly and ensuring the subsequent output interpretation accurately reflects the research question. The following sections detail these contexts, illustrating the critical difference between time-based studies and condition-based comparisons.

The two major applications of the one-way repeated measures ANOVA are:

Longitudinal Measurement Across Time Points: This scenario involves tracking changes in a specific dependent variable over a series of time intervals. For instance, researchers conducting a longitudinal study on health and fitness might measure the resting heart rate of a cohort of participants at three distinct periods: initial baseline (T1), mid-intervention (T2), and post-intervention (T3). The primary objective is to statistically determine if there is a significant shift in the group's mean resting heart rate across these three related measurement periods. Since the structure of the data shows the same individuals contributing scores at T1, T2, and T3, the dependency of the observations is confirmed.

Subject	Resting Heart Rate 1 Month Before Training Program	Resting Heart Rate in Middle of Training Program	Resting Heart Rate 1 Month After Training Program
Michael	65	58	60
Dwight	55	48	49
Andy	58	55	55
Meredith	68	60	64
Angela	47	45	45

The defining feature here is the **repeated** measurement of the same subjects, making the one-way repeated measures ANOVA the statistically correct choice to analyze the potential impact of the intervention across time.

Comparing Response Scores Across Different Conditions: Alternatively, this design can be employed in experimental settings where subjects are exposed to multiple, distinct stimuli or conditions. A behavioral scientist might present participants with three unique visual stimuli (e.g., short films, images, or sounds) and require them to provide a subjective rating (such as enjoyment, arousal, or attention) for each one. The analysis then compares the average response ratings across the three different experimental conditions. Crucially, every participant must provide a score for every condition.

Subject	Movie 1 Rating	Movie 2 Rating	Movie 3 Rating
Michael	88	84	92
Dwight	76	78	90
Andy	78	94	95
Meredith	80	83	88
Angela	82	90	99

Because the data points are linked by the individual subject, testing for a disparity in means across the three related conditions mandates the application of the one-way repeated measures ANOVA. Having established the conceptual framework, the subsequent sections detail the precise requirements and procedural steps for executing this analysis using the Stata statistical software package.

Preparing Data for Stata Implementation

To transition from theory to practice, we will utilize a classic pharmacological research example. Imagine a study investigating the neurocognitive effects of different treatments. Researchers measure the reaction time (the dependent variable) of five unique patients after administering four distinct drug compounds: Drug A, Drug B, Drug C, and Drug D. Because each of the five patients is tested under all four drug conditions, the resulting data points are inherently correlated, confirming the necessity of the repeated measures ANOVA. The central statistical question is whether the average reaction time exhibits a statistically significant difference across the administration of the four drugs.

Effective implementation within the **Stata** environment is contingent upon proper data structuring. Unlike many other statistical packages that prefer a long (stacked) format for RMANOVA, **Stata's** built-in `anova` command--when using the repeated measures options--often requires the data to be in a **wide format**. In this required format, each row represents a single subject (patient), and each level of the within-subjects factor (Drug A, Drug B, Drug C, Drug D) is represented by its own dedicated column. This specific structure allows Stata to correctly identify the dependency between the measurements and accurately calculate the sums of squares related to the within-subjects factor.

The subsequent guide outlines the exact procedural sequence required within **Stata**. These steps ensure the software correctly maps the variables, designating the drug variable as the within-subject factor (the variable being repeated) and the measured reaction time score as the dependent variable. Adhering to this sequence guarantees a reliable and accurate execution of the repeated measures model.

Executing the Repeated Measures Analysis in Stata

The execution of the RMANOVA in the **Stata** environment is highly procedural and requires careful attention to command syntax and variable specification. Although **Stata** permits analysis via command line, using the graphical user interface (GUI) is often preferred for initial clarity, especially when defining complex within-subjects factors. The following steps guide the user through loading the sample data, inspecting its structure, and finally running the statistical test.

Step 1: Load the Data into the Working Environment.

The initial requirement is to acquire the necessary dataset. For instructional purposes, we load a publicly available example dataset frequently used in **Stata** documentation. This command is entered directly into the command window and executed, instantly importing the necessary variables (representing the patients and the drug measurements) into the active memory.

use http://www.stata-press.com/data/r14/t43

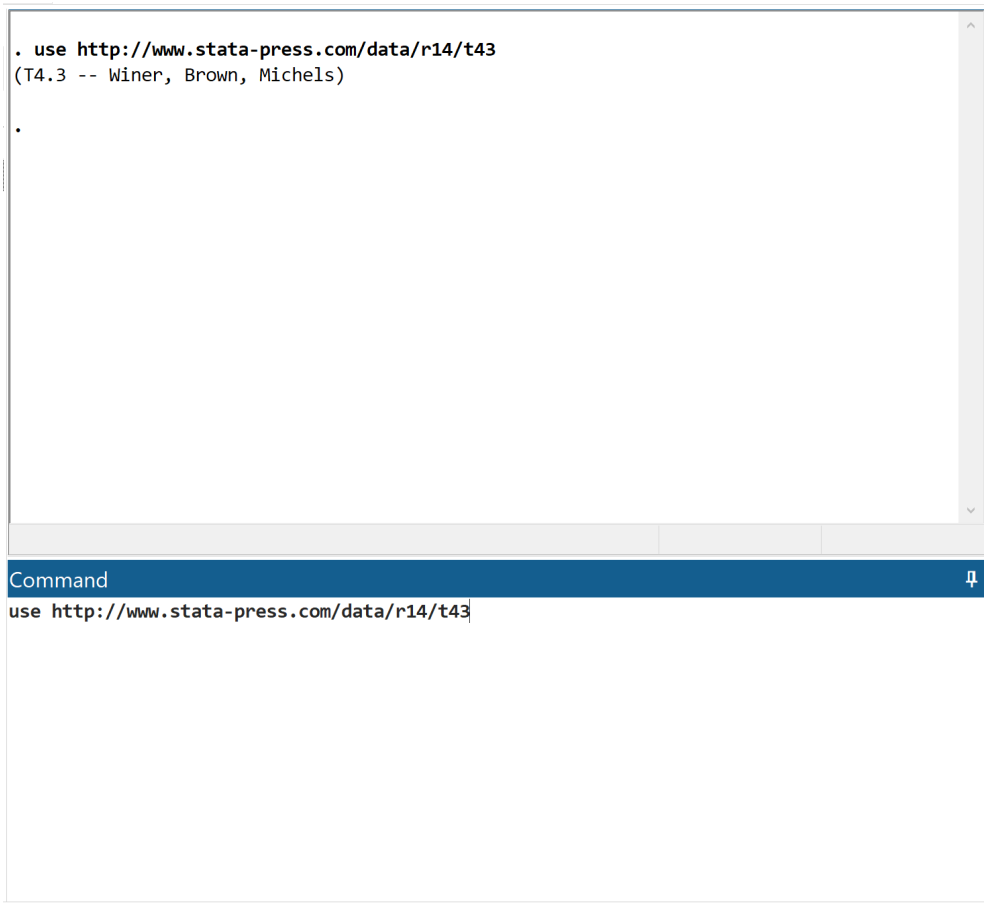

Verifying successful data loading is a necessary quality assurance step before proceeding to complex modeling. The data should now appear in the variables list, ready for manipulation and analysis.

Step 2: Inspect the Data Structure (Wide Format Verification).

Prior to analysis, it is standard methodological practice to visually inspect the raw data to confirm its integrity and adherence to the expected wide format. This is accomplished by navigating the GUI menu: selecting **Data > Data Editor > Data Editor (Browse)**. This action opens a read-only

view of the dataset. Confirm that each row corresponds to a unique subject (patient) and that the four columns--representing the reaction times under the four drugs--are present. This visual confirmation is crucial for ensuring the subsequent ANOVA command is correctly specified.

	person	drug	score
1	1	1	30
2	1	2	28
3	1	3	16
4	1	4	34
5	2	1	14
6	2	2	18
7	2	3	10
8	2	4	22
9	3	1	24
10	3	2	20
11	3	3	18
12	3	4	30
13	4	1	38
14	4	2	34
15	4	3	20
16	4	4	44
17	5	1	26
18	5	2	28
19	5	3	14
20	5	4	30

Step 3: Execute the Repeated Measures ANOVA Command.

The central statistical analysis is initiated using the ANOVA dialog box. Navigate to the appropriate analysis module (typically located under **Statistics > Linear models and related > ANOVA/MANOVA > Analysis of variance and covariance**). Within the ANOVA setup, the model must be precisely defined:

For the **Dependent variable**, specify **score** (representing reaction time).
In the **Model** specification, include both **person** (the subject identifier) and the within-subject factor, **drug**.
The critical step for RMANOVA is checking the box for **Repeated-measures variables**. Here, you must explicitly declare **drug** as the factor that repeats, which informs **Stata** about the dependency

structure of the observations.

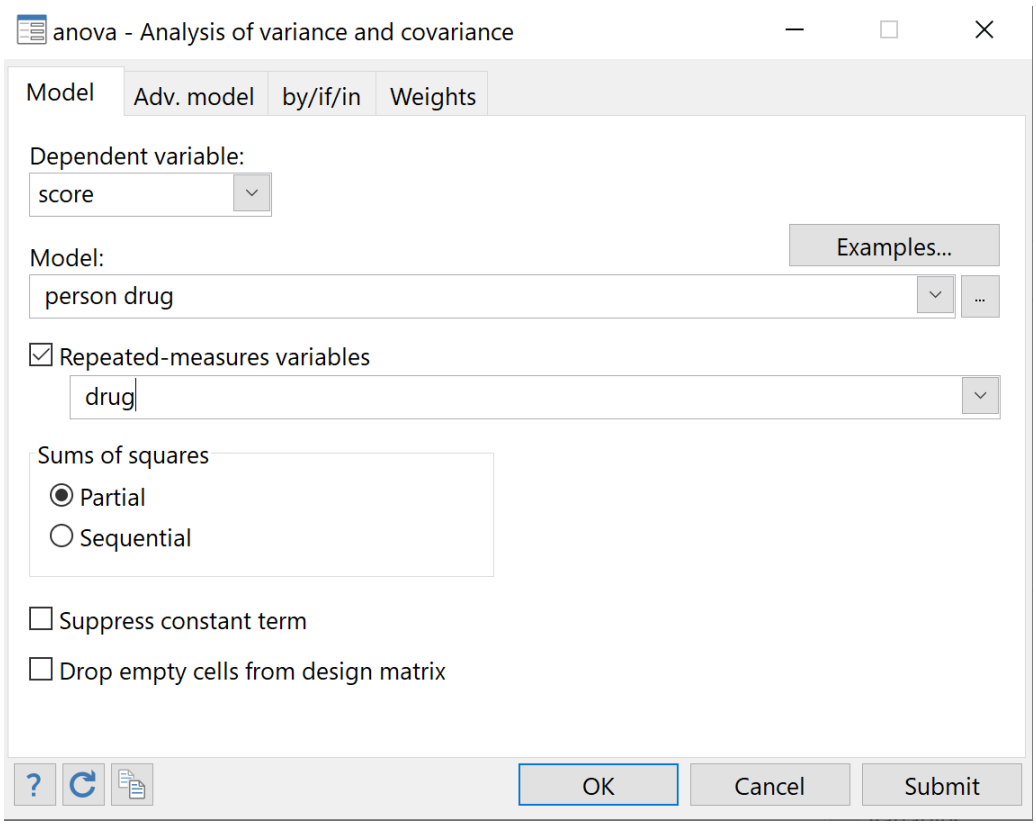

Upon clicking **OK**, **Stata** processes the command and outputs a series of results tables. These tables include the overall ANOVA summary and the critical tests related to model assumptions, which are vital for drawing valid conclusions about the drug effects.

Interpreting the Primary ANOVA Results

The output generated by Stata provides the necessary statistics to evaluate the primary research hypothesis--the null hypothesis--which posits that there is no variation in the population mean reaction time across the four administered drug treatments. The decision to reject or retain this null hypothesis hinges primarily on two crucial metrics presented in the main ANOVA table: the F value and its corresponding probability, or p-value.

```
. anova score person drug, repeated(drug)

                         Number of obs =         20    R-squared     =  0.9244
                         Root MSE      =    3.06594    Adj R-squared =  0.8803

                  Source | Partial SS         df         MS        F    Prob>F
              -----------+----------------------------------------------------
                   Model |       1379          7         197     20.96  0.0000
                         |
                  person |      680.8          4       170.2     18.11  0.0001
                    drug |      698.2          3   232.73333     24.76  0.0000
                         |
                Residual |      112.8         12         9.4
              -----------+----------------------------------------------------
                   Total |     1491.8         19   78.515789


Between-subjects error term:  person
                     Levels:  5          (4 df)
     Lowest b.s.e. variable:  person

Repeated variable: drug
                                          Huynh-Feldt epsilon        =  1.0789
                                          *Huynh-Feldt epsilon reset to 1.0000
                                          Greenhouse-Geisser epsilon =  0.6049
                                          Box's conservative epsilon =  0.3333

                                            ------------ Prob > F ------------
                  Source |     df      F    Regular    H-F      G-G      Box
              -----------+----------------------------------------------------
                    drug |      3    24.76   0.0000   0.0000   0.0006   0.0076
                Residual |     12
              -----------+----------------------------------------------------
```

When examining the initial output table (often organized by Source of variation), the researcher must direct their focus specifically to the row labeled **drug**, which represents the within-subjects factor. In this illustrative example, we observe an impressive **F value** of 24.76, accompanied by a probability value (labeled **Prob>F**) of 0.000. Applying the conventional alpha level (significance threshold) of 0.05, the observed **p-value** of 0.000 is substantially smaller. This highly significant finding compels us to reject the null hypothesis, leading to the conclusion that the four drug treatments do indeed result in a statistically significant difference in the mean reaction times experienced by the patients.

Addressing the Critical Sphericity Assumption

A secondary, yet profoundly important, component of the **Stata** output is the Sphericity Tests table, which addresses the foundational assumption of repeated measures ANOVA: the sphericity

assumption. Sphericity dictates that the variances of the differences between all possible pairs of the within-subjects factor (the drugs, in this case) must be approximately equal. Violation of this assumption is common and can severely distort the F-ratio, potentially inflating the test statistic and increasing the risk of a Type I error (falsely rejecting a true null hypothesis).

While the primary **Stata** output often omits Mauchly's test for sphericity, it proactively provides three crucial adjustments used when sphericity is questionable or violated. These correction factors--Greenhouse-Geisser, Huyhn-Feldt, and Box's Conservative--work by adjusting the degrees of freedom downward, thereby making the test more rigorous and conservative. These adjustments ensure the validity of the F-test, even when the underlying assumption is not met.

We review the adjusted **p-values** provided for the **drug** variable:

The **Hunyh-Feldt (H-F) p-value** is reported as 0.000.
The **Greenhouse-Geisser (G-G) p-value** is 0.0006.
The **Box's conservative (Box) p-value** is 0.0076.

In this specific example, since all three of the corrected **p-values** remain significantly below the 0.05 threshold, the powerful conclusion of a statistically significant difference among the drug conditions is maintained, demonstrating the robustness of the finding regardless of potential sphericity issues. If the primary F-test had been close to the significance boundary (e.g., p=0.04), relying on these conservative adjusted values would be essential for ethical and accurate reporting.

Standardized Reporting of RMANOVA Findings

The culmination of any statistical procedure is the standardized reporting of the results, typically following established style guidelines, such as those prescribed by the American Psychological Association (APA). A comprehensive report ensures both transparency and the reproducibility of the research. It is mandatory to include not only the primary inferential statistics but also the appropriate descriptive statistics (such as means and standard deviations for each condition) to contextualize the findings. The report must explicitly state the nature of the test, the F-statistic, the degrees of freedom used (both numerator and denominator), and the corresponding probability value.

When reporting the results derived from the **Stata** output, the researcher must clearly communicate the parameters of the study and the significance of the findings. If sphericity was violated, the corrected F-statistic and degrees of freedom (e.g., using Greenhouse-Geisser) should be cited instead of the uncorrected values. For this specific example, where the uncorrected result was highly significant and the corrected results confirmed the finding, we can report the primary result while noting the robust nature of the conclusion.

The following template illustrates the appropriate format for communicating the outcome of this repeated measures ANOVA executed in Stata:

A one-way repeated measures ANOVA was performed on the reaction times of five participants to assess the differential effects of four drug conditions.

The analysis revealed a statistically significant difference in mean reaction time dependent upon the drug administered ($F(3, 12) = 24.75$, $p < 0.001$). This finding robustly suggests that the pharmacological compounds have distinct effects on patient neurocognitive processing speed.