What is a Nested ANOVA? (Definition & Example)

Authored by

Mohammed looti

November 5, 2025

RECOMMENDED CITATION

Mohammed looti (2025). *What is a Nested ANOVA? (Definition & Example)*. PSYCHOLOGICAL STATISTICS. Retrieved from https://statistics.arabpsychology.com/?p=10438

The Fundamentals of Nested Analysis of Variance

A **nested ANOVA**, often interchangeably called a "hierarchical ANOVA," represents a specialized and powerful form of Analysis of Variance (ANOVA). This statistical modeling technique is essential when analyzing experimental data where the relationship between factors is not entirely independent, but rather, hierarchical. Specifically, a nested design is required when the levels of one experimental factor are specific to, or contained entirely within, the levels of another factor.

In standard experimental designs, researchers typically assume factors are crossed, meaning every combination of treatment levels is observed. However, the core identity of the nested model is its acknowledgment of dependency. If Factor B is nested within Factor A, the specific instances of B are unique to--and defined by--the specific level of A they occur under. This means that a comparison of Factor B levels across different levels of Factor A is statistically meaningless because those B levels are fundamentally distinct entities.

Understanding this structure is paramount for correct statistical inference. The primary function of a nested design is to ensure that the total variation observed in the dependent variable is accurately partitioned among the various sources of variability. By correctly assigning variability to the nested factor, we avoid inflating the overall error variance, thus yielding more precise and reliable tests for the main effects.

Why Structure Matters: Nested vs. Factorial Designs

The defining characteristic that separates a **nested ANOVA** from a standard Two-Way ANOVA lies in the concept of non-crossing factors. In conventional factorial designs, researchers investigate the potential for an interaction effect, which occurs when the effect of one factor changes depending on the level of the other factor. This investigation requires that all possible combinations of factor levels are present and observed.

In contrast, the structure of a nested design inherently precludes the calculation of interaction effects between the grouping factor and the nested factor. Because the levels of the nested factor (e.g., Technician 1 under Treatment A) never appear with any other grouping factor level (e.g., Treatment B), the necessary data points for evaluating a multiplicative interaction simply do not exist. Therefore, the effects in a nested model are purely additive, focusing on isolating the variability contributed by each factor independently within the existing hierarchy.

The statistical necessity of employing a nested design arises when experimental constraints or the natural organization of the study units dictate dependency. For example, if we are analyzing student performance, and we measure samples within Classrooms, and Classrooms are uniquely assigned to different Schools, then the Classroom factor is nested within the School factor. Using a nested model in this scenario ensures that variability unique to specific classrooms is properly isolated and accounted for before testing the broader effect of the School factor.

Case Study: Modeling Variability in a Fertilizer Experiment

To grasp the practical application of this design, consider an experiment aimed at evaluating the efficacy of three distinct fertilizer types (A, B, and C) on **plant growth**, which serves as our dependent variable. **Fertilizer Type** is the primary independent factor of interest.

To conduct the study, nine different laboratory technicians are employed to administer the fertilizers. Crucially, the experimental design mandates a strict hierarchy: three unique technicians are exclusively assigned to Fertilizer A, three others exclusively to Fertilizer B, and the final three exclusively to Fertilizer C. Each technician then applies their assigned fertilizer to a set number of plants.

In this setup, the **Technician factor is nested within the Fertilizer factor**. Technician 1 (assigned to A) and Technician 4 (assigned to B) are distinct individuals who never interact with the other fertilizer types. Their performance variability contributes to the variance within their specific fertilizer group, but they cannot be compared directly across groups in the context of the experiment's structure. This hierarchical relationship clearly demonstrates the dependency of the sub-units (Technicians) on the larger groups (Fertilizers):

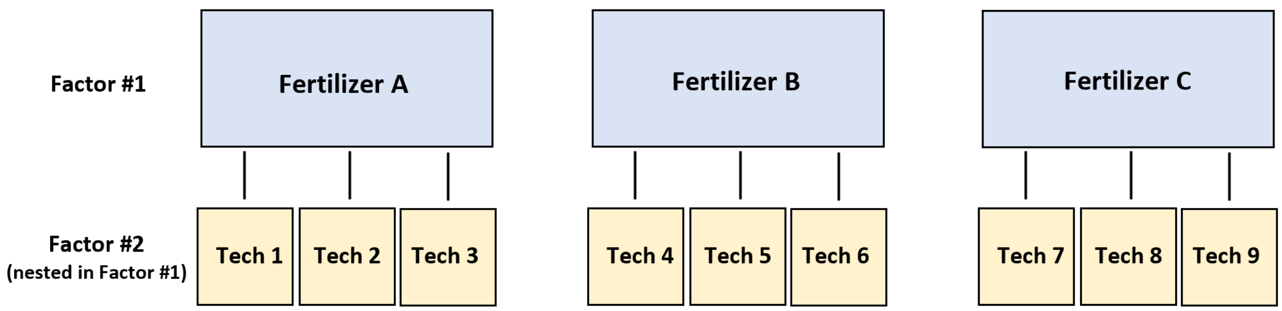

The resultant dataset mirrors this structure, where the identifying label for the nested factor (Technician ID) is only meaningful within the context of the grouping factor (Fertilizer Type). Analyzing this data requires a model that acknowledges this non-crossing structure to accurately attribute variation:

Plant Growth (inches)	Fertilizer	Technician
13	A	1
8	A	1
12	A	1
12	A	1
15	A	2
16	A	2
19	A	2
16	A	2
15	A	3
15	A	3
12	A	3
15	A	3
19	B	4
19	B	4
…	…	…

Hypothesis Testing and Goals of the Nested Model

The primary goal of performing a **nested ANOVA** is to rigorously test specific null hypotheses concerning the population means, effectively determining whether the primary factor or the nested factor significantly explains the variability observed in the dependent variable (plant growth).

A nested analysis typically tests two distinct null hypotheses simultaneously:

The mean plant growth is statistically equal across all levels of the primary factor (Fertilizer A = Fertilizer B = Fertilizer C). This tests the overall effect of the treatment itself.
The mean plant growth is statistically equal across all levels of the nested factor (Technician 1, 2, 3, etc.) within each level of the primary factor. This tests the variability introduced by the nested factor, such as human error or slight procedural differences among the technicians assigned to the same fertilizer.

The immense utility of the nested model lies in its ability to isolate these sources of variance. For instance, if the analysis reveals high variability attributable to the Technician factor, it suggests that inconsistencies in application or measurement protocols are a significant issue, even when the technicians are ostensibly performing the same task. Conversely, if the variability is primarily driven by the Fertilizer factor, the researcher can confidently conclude that the type of fertilizer is the dominant cause of the observed differences in plant growth.

Decoding the ANOVA Output Table

When statistical software executes the nested model, the results are synthesized into a standard ANOVA table format. This table meticulously partitions the total observed variance into the components attributed to the grouping factor (Fertilizer), the nested factor (Technician), and the unexplained residual error.

The output for our fertilizer example, structured to highlight the partitioning of variance, might appear as follows:

Source	Sum of Squares	df	Mean Square	F-Value	p-value
Fertilizer	372.7	2	186.3	53.238	< 0.0000
Technician(Fertilizer)	31.8	6	5.31	1.516	0.211
Residuals	94.5	27	3.5		

Interpreting this table requires a clear understanding of each column:

Source: Identifies the component of the model contributing to the total variance (e.g., Fertilizer, Technician, Residual).
Sum of Squares (SS): A measure of the total variability within that source. Larger SS values indicate greater variance attributable to that factor.
df (Degrees of Freedom): Represents the degrees of freedom associated with each source, reflecting the number of independent pieces of information used to estimate the variance.
Mean Square (MS): Calculated by dividing the Sum of Squares by the degrees_of_freedom (MS = SS/df). This value serves as an estimate of the population variance due to that specific factor.
F-Value: The test statistic, calculated by dividing the Mean Square of the factor by the appropriate error term. For nested designs, the choice of the correct denominator for the F-ratio is critical and depends on whether the factors are classified as fixed or random.
p-value: The probability value associated with the F-Value. This is the ultimate metric for determining statistical significance.

The final decision rests on the **p-value**. If the p-value is less than the chosen significance level (commonly 0.05), the null hypothesis for that factor is rejected, indicating a statistically significant effect. In the example provided, the Fertilizer factor demonstrates a significant effect (p-value < 0.05), whereas the Technician factor does not (p-value = 0.211). The practical conclusion is that the differences in plant growth are primarily driven by the **type of fertilizer**, and the variability introduced by individual technicians is not substantial enough to be considered a major source of variation.

Advanced Considerations: Fixed, Random, and Mixed Effects

While the structural definition of a **nested ANOVA** is straightforward, its statistical execution often requires advanced consideration regarding the nature of the factors involved. Factors are classified as either fixed or random, and this classification dictates the appropriate error term (the denominator in the F-ratio calculation) required for hypothesis testing.

A **fixed factor** is one whose levels were specifically chosen by the researcher, and the results are intended to generalize only to those specific levels (e.g., Fertilizer A, B, and C). Conversely, a **random factor** is one whose levels are assumed to be a random sample drawn from a larger population of potential levels, and the results are intended to generalize to that larger population (e.g., if the nine technicians were randomly selected from a pool of fifty available lab workers).

When the primary factor is fixed and the nested factor is random (a common scenario known as a mixed-effects nested ANOVA), the calculation of the F-ratio becomes more intricate. For instance, the effect of the primary (fixed) factor is often tested using the Mean Square of the nested factor as the denominator, rather than the Mean Square Residual. Incorrectly classifying the factors or using the wrong error term will lead to flawed F-statistics and inaccurate conclusions regarding statistical significance. Therefore, rigorous statistical planning is essential for any complex nested experiment.

Implementing these complex models often necessitates specific commands within statistical software packages (such as R or specialized modules in SPSS) to properly define the hierarchical nature and the fixed/random status of each factor. Standard ANOVA commands, which default to assuming fully crossed designs, will inevitably produce misleading results if applied to nested data.

How to perform a nested ANOVA in R.
How to perform a nested ANOVA in Excel or SPSS.