

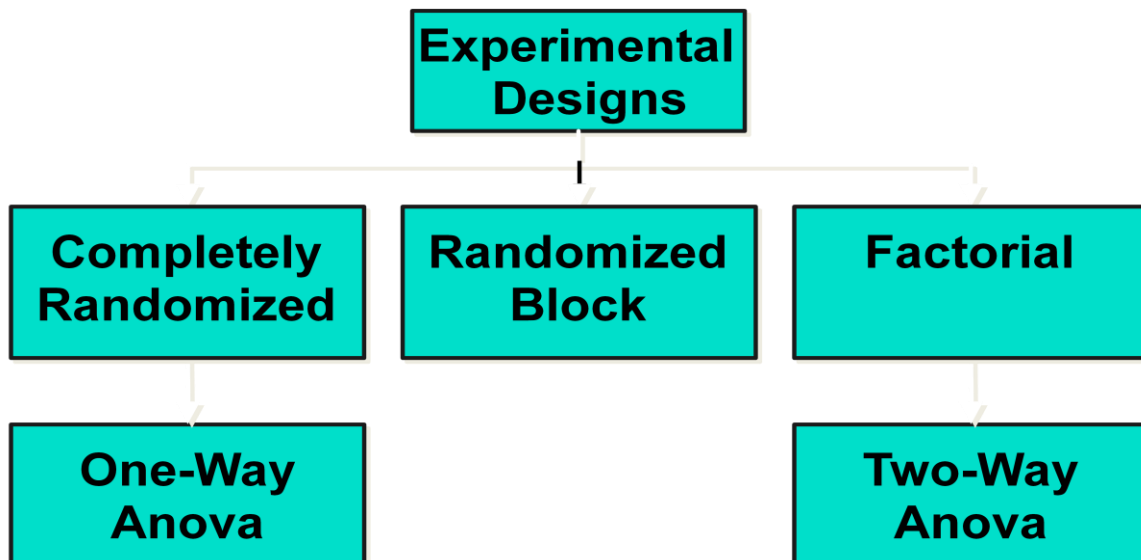
Analysis of Variance

ANOVA stands for analysis of variance, the generic name given to a set of techniques for studying cause-and-effect of one or more factors (independent variables) on a single dependent variable. The ANOVA is used when the independent variables are of nominal scale (categorical) and dependent variable is metric (continuous). The application areas in marketing research for experiments using ANOVA as the analytical method are wide. Whenever a marketing mix variable (independent variable) such as price, a specific promotion or type of distribution, even specific elements like shelf space, or color of packaging and so on is changed, we would want to know its effect. Under proper conditions, an experiment can tell us the effects of specific variation in one or more elements of the marketing mix.

Assumptions of ANOVA:

- (i) All populations involved follow a normal distribution.
- (ii) All populations have the same variance (or standard deviation).
- (iii) The samples are randomly selected and independent of one another.

Types of Experimental Designs



Completely Randomized Design

One-way ANOVA: This particular design is used when there is only one categorical independent variable, and one dependent (metric) variable. Each category of an independent variable is called a level. The independent variable may be different levels of prices, or different pack sizes, or different product colors, and the effect (dependent variable) could be the sale of the product. In this type of design, we randomly allocate the various sampling elements to the different levels of the independent variable, and measure the resulting dependent variable. Then, we conduct an F-test under the ANOVA to test the null hypothesis that the mean values of the dependent variable are not significantly different from each other, at different levels of the independent variable. If the computer output from the F-test shows a significance level (p-value)

of less than 0.05 on the ANOVA table, the null hypothesis is rejected. If the p-value from the F-test is greater than or equal to 0.05, the null hypothesis is accepted. Thus, if the p-value from the F-test is less than 0.05, it proves at the 95% confidence level that variation in the independent variable is able to cause significant variation in the dependent variable. That is, one variable depends on the other, assuming other variables are not in the picture.

Randomized Block Design

A **randomized block design** is used, if there is an additional variable (called the block) which has an impact on the relationship between the independent and dependent variables. This variable is accounted for in the design of a randomized block design by explicitly changing the levels of the block and testing if that has an impact on the relationship between the independent and dependent variables. An example of this is the **day of the week** having an impact on the relationship between type of display in a store (independent variable) and sales (dependent variable). Another example is the magazine in which an advertisement appears (block) affecting the impact (dependent variable) of different versions of the ad (independent variable).

Factorial Design

If two or more independent variables are to be tested through an ANOVA, we use a **factorial design**, because each independent variable in ANOVA is also known as a factor. The factorial design can accommodate several factors (independent variables) at several levels (categories) each. The major difference in analyzing factorial design with two or more factors is that interactions of two or three factors among themselves form a separate effect. These **interaction effects** need to be tested along with the **main effects** of the individual factors. This is quite easy to do on a computer, but care should be taken to assign sampling units or elements to each combination of the factors being tested for their effect. The interaction effects should first be looked at to see if they are significant (p-value less than 0.05 on the ANOVA table), which would show if there is a significant combined effect of the factors. If yes, the analysis must stop there. If not, we must also look at the results of the individual factor F-tests (called the main effects) to check if the factors individually have an effect on the dependent variable.

If any of the main effects are statistically significant, there are further pairwise tests to find out which levels of each factor have a significant impact on the dependent. These are called pairwise tests, and the common options available in most packages are Tukey's test, Duncan's test, and Scheffe's test.

One-Way Analysis of Variance (ANOVA) Example Problem

Introduction

Analysis of Variance (ANOVA) is a hypothesis-testing technique used to test the equality of two or more population (or treatment) means by examining the variances of samples that are taken. ANOVA allows one to determine whether the differences between the samples are simply due to random error (sampling errors) or whether there are systematic treatment effects that causes the mean in one group to differ from the mean in another.

Most of the time ANOVA is used to compare the equality of three or more means, however when the means from two samples are compared using ANOVA it is equivalent to using a t-test to compare the means of independent samples.

ANOVA is based on comparing the variance (or variation) *between* the data samples to variation *within* each particular sample. If the between variation is much larger than the within variation, the means of different samples will not be equal. If the between and within variations are approximately the same size, then there will be no significant difference between sample means.

Assumptions of ANOVA:

- (i) All populations involved follow a normal distribution.
- (ii) All populations have the same variance (or standard deviation).
- (iii) The samples are randomly selected and independent of one another.

Since ANOVA assumes the populations involved follow a normal distribution, ANOVA falls into a category of hypothesis tests known as parametric tests. If the populations involved did not follow a normal distribution, an ANOVA test could not be used to examine the equality of the sample means. Instead, one would have to use a non-parametric test (or distribution-free test), which is a more general form of hypothesis testing that does not rely on distributional assumptions.

Example

Consider this example:

Suppose the National Transportation Safety Board (NTSB) wants to examine the safety of compact cars, midsize cars, and full-size cars. It collects a sample of three for each of the treatments (cars types). Using the hypothetical data provided below, test whether the mean pressure applied to the driver’s head during a crash test is equal for each types of car. Use $\alpha = 5\%$.

Table ANOVA.1

	Compact cars	Midsize cars	Full-size cars
	643	469	484
	655	427	456
	702	525	402
\bar{X}	666.67	473.67	447.33
S	31.18	49.17	41.68

(1.) State the null and alternative hypotheses

The null hypothesis for an ANOVA always assumes the population means are equal. Hence, we may write the null hypothesis as:

$H_0: \mu_1 = \mu_2 = \mu_3$ - The mean head pressure is statistically equal across the three types of cars.

Since the null hypothesis assumes all the means are equal, we could reject the null hypothesis if only mean is not equal. Thus, the alternative hypothesis is:

H_a : At least one mean pressure is not statistically equal.

(2.) Calculate the appropriate test statistic

The test statistic in ANOVA is the ratio of the *between* and *within* variation in the data. It follows an F distribution.

Total Sum of Squares – the total variation in the data. It is the sum of the between and within variation.

Total Sum of Squares (SST) = $\sum_{i=1}^r \sum_{j=1}^c (X_{ij} - \bar{\bar{X}})^2$, where r is the number of rows in the table, c is the number of columns, $\bar{\bar{X}}$ is the grand mean, and X_{ij} is the i th observation in the j th column.

Using the data in Table ANOVA.1 we may find the grand mean:

$$\bar{\bar{X}} = \frac{\sum X_{ij}}{N} = \frac{(643 + 655 + 702 + 469 + 427 + 525 + 484 + 456 + 402)}{9} = 529.22$$

SST =

$$(643 - 529.22)^2 + (655 - 529.22)^2 + (702 - 529.22)^2 + (469 - 529.22)^2 + \dots + (402 - 529.22)^2 = 96303.55$$

Between Sum of Squares (or Treatment Sum of Squares) – variation in the data between the different samples (or treatments).

Treatment Sum of Squares (SSTR) = $\sum r_j (\bar{X}_j - \bar{\bar{X}})^2$, where r_j is the number of rows in the j th treatment and \bar{X}_j is the mean of the j th treatment.

Using the data in Table ANOVA.1,

$$\text{SSTR} = [3 * (666.67 - 529.22)^2] + [3 * (473.67 - 529.22)^2] + [3 * (447.33 - 529.22)^2] = 86049.55$$

Within variation (or Error Sum of Squares) – variation in the data from each individual treatment.

$$\text{Error Sum of Squares (SSE)} = \sum \sum (X_{ij} - \bar{X}_j)^2$$

From Table ANOVA.1,

$$\begin{aligned} \text{SSE} = & [(643 - 666.67)^2 + (655 - 666.67)^2 + (702 - 666.67)^2] + \\ & [(469 - 473.67)^2 + (427 - 473.67)^2 + (525 - 473.67)^2] + \\ & [(484 - 447.33)^2 + (456 - 447.33)^2 + (402 - 447.33)^2] = 10254. \end{aligned}$$

Note that SST = SSTR + SSE (96303.55 = 86049.55 + 10254).

Hence, you only need to compute any two of three sources of variation to conduct an ANOVA. Especially for the first few problems you work out, you should calculate all three for practice.

The next step in an ANOVA is to compute the “average” sources of variation in the data using SST, SSTR, and SSE.

Total Mean Squares (MST) = $\frac{SST}{N-1}$ → “average total variation in the data” (N is the total number of observations)

$$MST = \frac{96303.55}{(9-1)} = 12037.94$$

Mean Square Treatment (MSTR) = $\frac{SSTR}{c-1}$ → “average between variation” (c is the number of columns in the data table)

$$MSTR = \frac{86049.55}{(3-1)} = 43024.78$$

Mean Square Error (MSE) = $\frac{SSE}{N-c}$ → “average within variation”

$$MSE = \frac{10254}{(9-3)} = 1709$$

Note: $MST \neq MSTR + MSE$

The test statistic may now be calculated. For a one-way ANOVA the test statistic is equal to the ratio of MSTR and MSE. This is the ratio of the “average between variation” to the “average within variation.” In addition, this ratio is known to follow an F distribution. Hence,

$F = \frac{MSTR}{MSE} = \frac{43024.78}{1709} = 25.17$. The intuition here is relatively straightforward. If the average between variation rises relative to the average within variation, the F statistic will rise and so will our chance of rejecting the null hypothesis.

(3.) Obtain the Critical Value

To find the critical value from an F distribution you must know the numerator (MSTR) and denominator (MSE) degrees of freedom, along with the significance level.

F^{CV} has df_1 and df_2 degrees of freedom, where df_1 is the numerator degrees of freedom equal to $c-1$ and df_2 is the denominator degrees of freedom equal to $N-c$.

In our example, $df_1 = 3 - 1 = 2$ and $df_2 = 9 - 3 = 6$. Hence we need to find $F_{2,6}^{CV}$ corresponding to $\alpha = 5\%$. Using the F tables in your text we determine that $F_{2,6}^{CV} = 5.14$.

(4.) Decision Rule

You reject the null hypothesis if: F (observed value) $> F^{CV}$ (critical value). In our example $25.17 > 5.14$, so we reject the null hypothesis.

(5.) Interpretation

Since we rejected the null hypothesis, we are 95% confident ($1 - \alpha$) that the mean head pressure is not statistically equal for compact, midsize, and full size cars. However, since only one mean must be different to reject the null, we do not yet know which mean(s) is/are different. In short, an ANOVA test will test us that *at least one mean is different*, but an additional test must be conducted to determine which mean(s) is/are different.

Determining Which Mean(s) Is/Are Different

If you fail to reject the null hypothesis in an ANOVA then you are done. You know, with some level of confidence, that the treatment means are statistically equal. However, if you reject the null then you must conduct a separate test to determine which mean(s) is/are different.

There are several techniques for testing the differences between means, but the most common test is the Least Significant Difference Test.

Least Significant Difference (LSD) for a *balanced* sample: $\sqrt{\frac{2 * MSE * F_{1,N-c}}{r}}$, where MSE is the mean square error and r is the number of rows in each treatment.

In the example above, $LSD = \sqrt{\frac{(2)(1709)(5.99)}{3}} = 82.61$

Thus, if the absolute value of the difference between any two treatment means is greater than 82.61, we may conclude that they are not statistically equal.

Compact cars vs. Midsize cars:

$|666.67 - 473.67| = 193$. Since $193 > 82.61 \rightarrow$ mean head pressure is statistically different between compact and midsize cars.

Midsize cars vs. Full-size cars:

$|473.67 - 447.33| = 26.34$. Since $26.34 < 82.61 \rightarrow$ mean head pressure is statistically equal between midsize and full-size cars.

Compact vs. Full-size:
Work this on your own.

One-way ANOVA in Excel

You may conduct a one-way ANOVA using Excel.

(Preliminary step) First, make sure that the “Analysis ToolPak” is installed.

Under “Tools” is the option “Data Analysis” present?

If yes – ToolPak is installed.

If no – select “Add-ins.”

Check the boxes entitled “Analysis ToolPak” and “Analysis ToolPak – VBA” and click “OK”. This will install the “Data Analysis ToolPak.”

(1.) Under “Tools” select “Data Analysis”

In the window that appears select “ANOVA: One factor” and click “OK.”

(2.) Using your mouse highlight the cells containing the data.

(3.) Select “Columns” if each treatment is its own column or “Row” if each treatment is its own row.

(4.) Set your level of significance. (The default is 5% or 0.05.)

(5.) Click “OK” and the ANOVA output will appear on a new worksheet.

ANOVA Results from Excel:

SUMMARY

<i>Groups</i>	<i>Count</i>	<i>Sum</i>	<i>Average</i>	<i>Variance</i>
Column 1	3	2000	666.6667	972.3333
Column 2	3	1421	473.6667	2417.333
Column 3	3	1342	447.3333	1737.333

ANOVA

<i>Source of Variation</i>	<i>SS</i>	<i>df</i>	<i>MS</i>	<i>F</i>	<i>P-value</i>	<i>F crit</i>
Between Groups	86049.55556	2	43024.78	25.17541	0.001207	5.143249
Within Groups	10254	6	1709			
Total	96303.55556	8				

The results under the heading “SUMMARY” simply provides you with summary statistics for each of your samples. The results of the ANOVA test are provided under the heading “ANOVA.” Comparing these figures with the example above, it should be simple to determine the meaning of the Excel output.