Introduction to ANOVA

Education comes from within; you get it by struggle and effort and thought. - Napoleon Hill

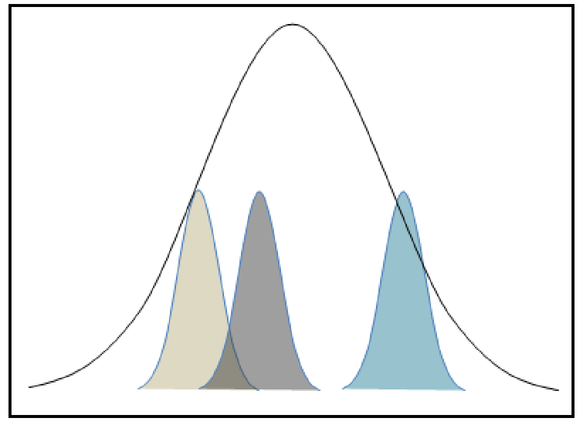

Are you ready to **think** and **struggle** through the next topic in the CQE Body of Knowledge - **ANOVA**.

ANOVA stands for **AN**alysis **O**f **VA**riance and it is a **hypohesis test** used to compare the means of 2 or more groups.

I know what you're thinking - WHY would we test MEAN VALUES using VARIANCE.

I thought the same thing. . .

Don't worry though.

It'll all make sense soon!

Ok, so here's what you'll learn in this chapter:

First we cover the **assumptions** associated with ANOVA.

Second are the **common terms & definitions** within ANOVA (and **DOE**).

Third, I start very high level to answer the question - **Why Does ANOVA use Variance to Test Mean Values?**

Variation Source	Sum of Squares (SS)	Degrees of freedom (DF)	Mean Squares (MS)	F-Value
Treatment (Between)	Sum of Squares of the Treatment **(SS_t)**	DF of the Treatment **(DF_t)**	Mean Squares of the Treatment **(MST)**	**= MST / MSE**
Error (Within)	Sum of Squares of the Error **(SS_e)**	DF of the Error **(DF_e)**	Mean Squares of the Error **(MSE)**	
Total	Total Sum of Squares **(SS_{total})**	Total DF **(DF_{Total})**		

Fourth, we go through the **basics of ANOVA** including the **Sum of Squares, Degrees of Freedom, Mean Squares**, **F-value**.

Then we go into an example of a **One Way ANOVA** and use all of the basics we just learned.

Last, but not least I'll introduce the idea of a **Two Way ANOVA**, including new terms and concepts.

Ready to get started?

Assumptions of ANOVA

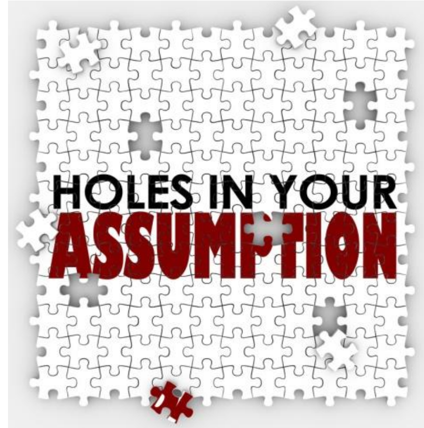

Ok, so first things first - **ANOVA** is a type of **Hypothesis Test** used to test hypothesis about Mean values.

So, similar to the other Hypothesis Tests we studied, ANOVA analysis is based on the *starting assumption that the null hypothesis is true*.

Within ANOVA, the **Null Hypothesis** is always the same - that all of our sample mean values are equal.

$$H_o: \mu_a = \mu_b = \mu_c = \ldots\ldots \mu_k$$

The **Alternative Hypothesis** is that at least one mean value is different than the rest.

H_a: Not all means are equal

Make Sense?

The other **3 major assumptions** for ANOVA are identical to the assumptions associated with the **t-test**.

Remember that the **t-test** is used to test the hypothesis that two means are equal, while **ANOVA** is used to test the hypothesis that 3+ means are equal. So it's logical that they would share common assumptions, which include:

- The Population being studied is Normally Distributed
- The Variance is the same between the various treatments (Homogeneity of Variance)
- There is Independence between sample observations

So your population must be normally distributed, and the variances between your treatments should be equal and your samples should be independent from each other.

Terminology in ANOVA

Alright, let' s go deeper with ANOVA and review the **common terms** and their **definitions**.

What you'll often find is that **ANOVA Analysis** is paired with a **Designed Experiment** to measure the affect that an **independent variable** has on a **dependent variable**.

For example, let's say you wanted to study the effect that different octane gasses have on the horsepower of your car.

You would **design an experiment** where you would vary the octane of gas and measure horsepower.

In this experiment, the **independent variable** would be the octane of gas (87, 89, 91, etc), and the **dependent variable** would be the horsepower.

The independent variable in ANOVA is called a **Factor**.

In the Horsepower example there is only one factor (Octane), which makes this experiment a **one way ANOVA**. If there had been two factors (Octane and Fuel Injector Size) it would be a **two way ANOVA**.

You'll notice above that I listed different octane's of gas - 87, 89, 91. These are the different **levels or treatments** associated with your factor.

So our **Factor** (Octane) can have multiple **Levels** (87, 89, 91) in a **One Way ANOVA** analysis.

The **Response** that we're measuring in this experiment is Horsepower, which is our **dependent variable**.

Make Sense? Let's move on from here to review the basics of ANOVA.

Why Does ANOVA use Variance to Test Mean Values?

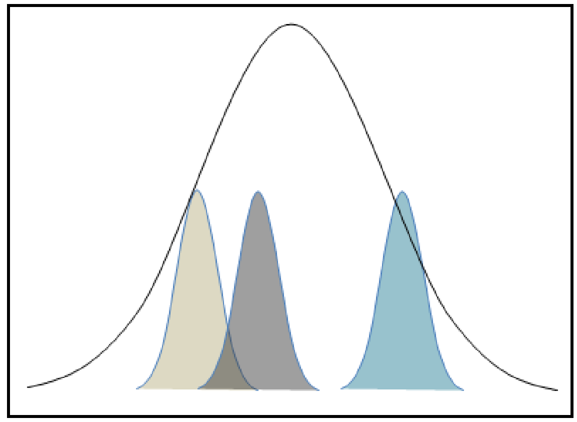

ANOVA Takes a very unique and interesting approach to determining if 3 or more sample means all have an **equal population mean**.

What **ANOVA Analysis** does - and this is truly brilliant - is it breaks down your data set into **two different sources of variation.**

When you look at the standard One Way ANOVA table (Below), you'll see the two "**sources**" of variation are called the **Treatment** and the **Error**.

The **Treatment Variation** is the variability within the data set that can be attributed to the difference **between the different treatment groups** (difference in sample means).

The **Error Variation** is the variability within the data set that can be attributed to the random error associated with the response variable. This is the **variability within the different treatment groups**.

Recall that our null hypothesis within ANOVA is that all population means are equal.

So if our null hypothesis were true, then we would expect that the **treatment variation** (difference in sample means) could be fully explained by the random nature of the data.

So **the treatment variation** should be nearly equal to the **error variation**. This statement is the key to explaining **why ANOVA uses Variance to test the difference between Means**.

Variation Source	Sum of Squares (SS)	Degrees of freedom (DF)	Mean Squares (MS)	F-Value
Treatment (Between)	Sum of Squares of the Treatment (SS_t)	DF of the Treatment (DF_t)	Mean Squares of the Treatment (MST)	= MST / MSE
Error (Within)	Sum of Squares of the Error (SS_e)	DF of the Error (DF_e)	Mean Squares of the Error (MSE)	
Total	Total Sum of Squares (SS_{total})	Total DF (DF_{Total})		

By breaking down our data into the two sources of variance, we can then compare those variances against each of to see if they are statistically significantly different (**the alternate hypothesis**).

What I'd like to do next is to quickly go through each of the columns of the ANOVA Table and go over each topic specifically.

Including **the Sum of Squares and Degrees of Freedom** which are combined to calculate the **Mean Square Values**, which are further combined to calculate your **F-Statistic or F-value.**

Sum of Squares in ANOVA

Step 1 in the ANOVA process is **calculating the Sum of Squares** for each of the sources of variation (**Treatment & Error**), and then adding them up to the **total variation** within the data set.

The **Variation** for each of these sources is calculated using the **"Sum of Squares"** calculation. So we calculate the **Sum of Squares of the Treatment (SS_t)**, and the **Sum of Squares of the Error (SS_e)**.

These different sources of variation combine to add up to the **Total Sum of Squares (SS_{total})** within your data set.

Total Sum of Squares (SS_{total}) = Sum of Square of the Treatment (SS_t) + Sum of Squares of the Error (SS_e)

Sum of Squares Example

Ok, let's go over the formulas to calculate the sum of squares for the treatment, error and then the total sum of squares.

Let's say we're back on the horsepower/octane example and we've designed an experiment where we're testing 3 treatment levels (octane levels), and we're measuring horsepower 3 times per treatment.

	Treatment Group #1 87 Octane	Treatment Group #2 89 Octane	Treatment Group #3 91 Octane
Sample #1 Horsepower	223 hp	224 hp	225 hp
Sample #2 Horsepower	224 hp	225 hp	226 hp
Sample #3 Horsepower	225 hp	226 hp	227 hp
Sample Mean	224 hp	225 hp	226 hp

So, Treatment group #1 captures our 3 measurements at 87 Octane, and it had a sample mean of 224 hp. Make Sense?

Before we calculate the sum of squares, let's introduce a new topic. **The Grand Mean**.

So, recall that our null hypothesis is that all three sample means come from the same population mean. If that's true, we can calculate the **Grand Mean**, which is an estimate of the population mean.

The Grand Mean for this example is 225 hp. You can calculate the grand mean by averaging all of the individual values within the data set.

Sum of Squares of the Treatment

Now we can use the **grand mean** to calculate the **Sum of Squares of the Treatment (SS_t).**

$$SS_t = \sum \boldsymbol{n(\bar{X}_1 - GM)^2 + n(\bar{X}_2 - GM)^2 + n(\bar{X}_3 - GM)^2 + \ldots n(\bar{X}_i - GM)^2}$$

Recall that the **Treatment Variation** is the variability within the data set that can be attributed to the difference **between the different treatment groups**.

We do this by comparing the treatment sample means (X-bar of Treatment Group 1, etc) against the grand mean. We're also multiplying this difference by "n", which is the number of samples per treatment group.

$$SS_t = \sum \boldsymbol{3(224-225)^2 + 3(225-225)^2 + 3(226-225)^2} = \sum \boldsymbol{3+0+3=6}$$

So the **sum of squares of the treatment** is equal to **6**.

Sum of Squares of the Error

Now let's move on to the **Sum of Squares of the Error (SS_e)** which can be done using the equation below:

$$SS_e = \sum (X_i - \bar{X})^2$$

By the way, do you see why it's called the **sum of squares calculation**? We're squaring the difference between two values, then summing up those differences. Hence, the **sum of squares**.

Remember that the **Error Variation** is the variability within the data set that can be attributed to the random error associated with the response variable. This is the **variability within the different treatment groups**.

So we're comparing the individual values (X_i) against that values sample mean (X-bar).

$$SS_e = \sum (223-224)^2 + (224-224)^2 + (225-224)^2 + (224-225)^2 + (225-225)^2 + (226-225)^2 + (225-226)^2 + (226-226)^2 + (227-226)^2$$

I've color coded this equation so you can see how the different treatment groups are considered within this equation.

So treatment group 1 is shown in green, where the 3 individual values (223, 224 & 225) are compared against the sample mean of treatment group #1 (224). This continues with Treatment Group 2 & 3.

$$SS_e = \sum 1 + 0 + 1 + 1 + 0 + 1 + 1 + 0 + 1 = 6$$

So the **sum of squares of the error** is equal to **6**.

Total Sum of Squares

Now let's move on to the **Total Sum of Squares (SS_{total}).**

We could simply calculate the **total sum of squares** by adding up the **Sum of Square of the Treatment (SS_t)** and the **Sum of Squares of the Error (SS_e).** Or we could calculate it using the following equation:

$$SS_{total} = \sum (X_i - GM)^2$$

This looks similar to the sum of squares of the error calculation; however, we're now comparing all individual values (X_i) against the grand mean (**225**).

$$SS_{total} = \sum (223-225)^2 + (224-225)^2 + (225-225)^2 + (224-225)^2 + (225-225)^2 + (226-225)^2 + (225-225)^2 + (226-225)^2 + (227-225)^2$$

$$SS_{total} = \sum 4 + 1 + 0 + 1 + 0 + 1 + 0 + 1 + 4 = 12$$

Let's compare this to the Sum of Squares of the Treatment and Sum of Squares of the Error:

Total Sum of Squares (SS_{total}) = Sum of Square of the Treatment (SS_t) + Sum of Squares of the Error (SS_e)

12 = 6 + 6

See how that reconciles? The **total sum of squares captures the total variability within the data set** by comparing all values against the grand mean, while **the Treatment & Error sum of squares are a unique component of that overall variability**.

Alright, so we've knocked out the **Sum of Squares** portion of the ANOVA Table, let's move on to the **Degrees of Freedom column**.

Degrees of Freedom in ANOVA

Step 2 in the ANOVA process is to calculate the **degrees of freedom (DF)** for each **source of variation (Treatment & error),** which add up to the **total degrees of freedom**.

We will use the degrees of freedom to "normalize" the **sum of squares data** to convert the raw "variation" into estimations of the **population variance** in step 3 (Mean Squares).

So we will start by calculating the **total degrees of freedom (DF_{total})** associated with the entire sample data set.

$$DF_{total} = DF_{error} + DF_{treatment}$$

Then we will break this down into the **degrees of freedom for the treatment ($DF_{treatment}$),** and the **degrees of freedom for the error (DF_{error}).**

The **Total Degrees of Freedom** is the easiest to calculate - It's the **total number of observations within your data set, minus 1**. The letter N is often used to represent the total number of observations within a data set.

$$DF_{total} = N - 1$$

If we go back to the example above, there are 9 total observations within our experiment, so N = 9. Then, the **total degrees of freedom is 8**.

$$DF_{total} = N - 1 = 9 - 1 = 8$$

The **DF of the Treatment** is the next easiest to calculate - It's the number of treatments minus 1. The letter "a" is often used to denote the number of treatments, but this can vary between textbooks.

$$DF_{treatment} = a - 1$$

If we go back to the example above, there are 3 treatment levels (87, 89 & 91 octane) within our experiment, so a = 3. Then, the **degrees of freedom of the treatment is 2**.

$$DF_{treatment} = a - 1 = 3 - 1 = 2$$

The DF of the Error is often calculated by taking the total degrees of freedom, and subtracting the treatment degrees of freedom.

$$DF_{error} = DF_{total} - DF_{treatment} = (N-1) - (a-1) = (9 - 1) - (3 - 1) = 6$$

Where N equals the total number of observations, which can be calculated as n(a), which is the number of treatment groups (a) times the number of samples per treatment group (n). Using this transformation of **N=n*a**, we can re-arrange the equation to this:

$$DF_{error} = a(n-1) = 3(3-1) = 6$$

Ok, before we move on to the next section, let's take a look at our ANOVA table to see how we've incorporated all we've learned so far.

Variation Source	Sum of Squares (SS)	Degrees of freedom (DF)	Mean Squares (MS)	F-Value
Treatment (Between)	6	2		
Error (Within)	6	6		
Total	12	8		

So for the Treatment and Error we've calculated the **Sum of Squares** for each and the **Degrees of Freedom** for each. Now it's time to use that information to move to step 3 in the ANOVA process, **calculating the Mean Squares**.

Mean Squares in ANOVA

Alright, on to step 3 of the ANOVA analysis, which is the calculation of our **Mean Squares**, where again we'll calculate a **Mean Square of the Treatment (MST)**, and a **Mean Square of the Error (MSE)**.

Mean Squares are calculated using the equation below:

$$Mean\ Square = \frac{Sum\ of\ Squares}{Degrees\ of\ Freedom}$$

We take our **Sum of Squares** and we divide it by our **degrees of freedom**. We will do this for both our treatment and error terms.

To help explain what the Mean Squares represents I'll jump right to the calculation for the **Mean Square of the Error:**

$$\boldsymbol{Mean\ Square\ of\ the\ Error} = \frac{Sum\ of\ Squares\ of\ the\ Error}{Degrees\ of\ Freedom\ of\ the\ Error} = \frac{\sum(X_i - \bar{X})^2}{(N-1) - (a-1)}$$

This might look super familiar to how we calculate **sample variance**:

$$\boldsymbol{Sample\ Variance} = s^2 = \frac{\sum(x-\bar{x})^2}{n-1} = \frac{sum\ of\ squares}{degrees\ of\ freedom}$$

I wanted to show you this to make the point that the **Mean Square calculation** is a calculation of variance!

Specifically, **MST & MSE** are both **unique & different estimates** of the **population variance** associated with our data set.

If our null hypothesis is true then these two estimates of the population variance will be approximately equal.

So, the **Mean Square of the Error (MSE)** is an estimate of the population variance that's based solely on the **variability within each treatment group**.

The **Mean Square Between Treatments (MST)** is an estimate of the population variance that's based solely on the **variability between the treatment group sample means and the grand mean**.

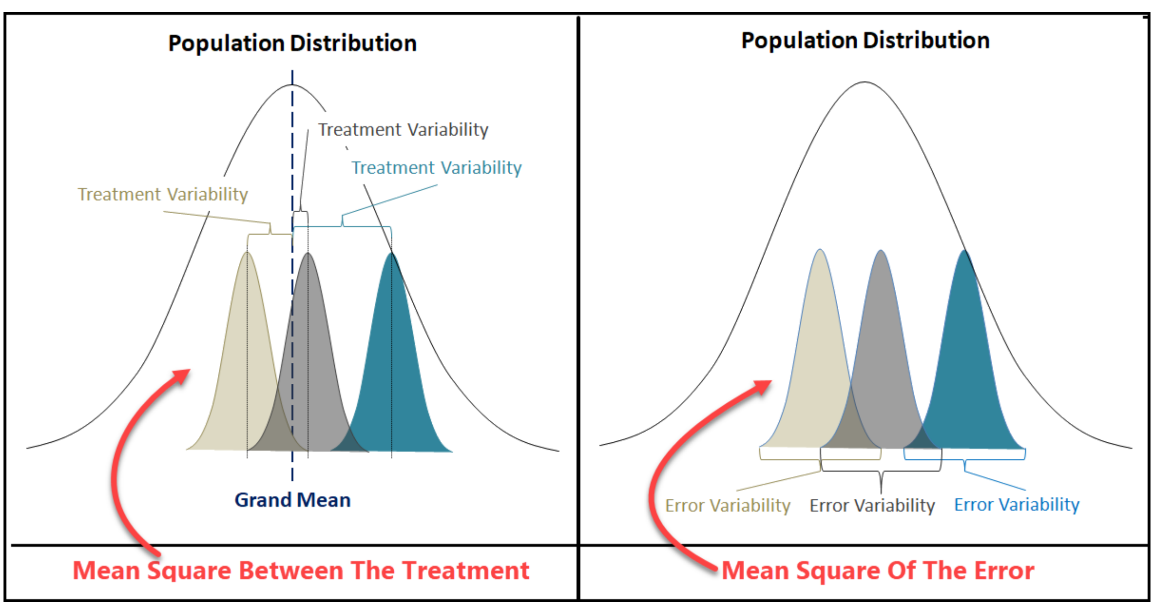

Can you see on the left of this image how **MST (Mean Square Between the Treatment)** is an estimate of the **variability between each treatment group and the grand mean**?

It compares the sample mean of each treatment to the overall grand mean of the population.

Can you see on the right of this image how **MSE (Mean Square of the Error)** is an estimate of the **variability within each treatment group**?

It compares the individual observations within each treatment group, to the sample mean of that treatment group.

Let's review these two **Mean Squares** individually.

MSE - The Variability ***Within*** the Treatment Groups

The first type of variation measured within ANOVA is the **Mean Square of the Error** which represents the **random variability** that is normal to the **response variable** which.

You'll see other textbooks call this the **Within Treatment Variability** because it's a reflection of the variability within each individual treatment group.

If we go back to our original example of octane & horsepower, you can see that within the 87 octane group, there is some slight variability WITHIN each treatment group.

Variability within this treatment group

	Treatment Group #1 87 Octane	Treatment Group #2 89 Octane	Treatment Group #3 91 Octane
Sample #1 Horsepower	223 hp	224 hp	225 hp
Sample #2 Horsepower	224 hp	225 hp	226 hp
Sample #3 Horsepower	225 hp	226 hp	227 hp
Sample Mean	224 hp	225 hp	226 hp

This random variation in horsepower that's inherent to each treatment group is the **Error Variance**.

This next statement is **important!**

Whether the null hypothesis is true or not, the **MSE (Mean Square of the Error) is a good approximation of the population variance.**

To calculate the MSE, we use the following formula:

$$\boldsymbol{Mean\ Square\ of\ the\ Error} = \frac{Sum\ of\ Squares\ of\ the\ Error}{Degrees\ of\ Freedom\ of\ the\ Error}$$

The **Mean Square of the Error (MSE)** is the **Sum of Squares of the Error** (SS_{error}) divided by the **Degrees of Freedom of the Error (DF_e)**. In our specific example for octane and horsepower, the MSE is:

$$\boldsymbol{Mean\ Square\ of\ the\ Error} = \frac{Sum\ of\ Squares\ of\ the\ Error}{Degrees\ of\ Freedom\ of\ the\ Error} = \frac{6}{6} = \mathbf{1}$$

Let's move on to discuss the next estimate of the population variance - **the MST.**

Treatment - The Variability *Between* the Treatment Groups

The second type of variability within ANOVA is called the **MST or Mean Square Between the Treatments** and it's a estimate of the population variance that's based on the difference between the sample means and the grand mean.

You might also see textbooks call this **MSB** for Mean Square Between (the treatments), or the **Variance Between Treatments** because it reflects the variance caused by the different treatments (levels).

This measure of variability compares the sample mean of each treatment (X-bar), against the **Grand Mean** of the entire sample space (all sample observations).

MST uses the Variability in the sample means to estimate the Population Variance

	Treatment Group #1 87 Octane	Treatment Group #2 89 Octane	Treatment Group #3 91 Octane
Sample #1 Horsepower	223 hp	224 hp	225 hp
Sample #2 Horsepower	224 hp	225 hp	226 hp
Sample #3 Horsepower	225 hp	226 hp	227 hp
Sample Mean	224 hp	225 hp	226 hp

This next **statement is important:**

If the null hypothesis is false, then the *MST will not be an accurate measure of the population variance.*

To calculate the MSE, we use the following formula:

$$\boldsymbol{Mean\ Square\ of\ the\ Treatment} = \frac{Sum\ of\ Squares\ of\ the\ Treatment}{Degrees\ of\ Freedom\ of\ the\ Treatment}$$

The **Mean Square of the Treatment (MST)** is the **Sum of Squares of the Treatment ($SS_{treatment}$)** divided by the **Degrees of Freedom of the Treatment** (**DF_t**). In our specific example for octane and horsepower, the MST is:

$$\boldsymbol{Mean\ Square\ of\ the\ Treatment} = \frac{Sum\ of\ Squares\ of\ the\ Treatment}{Degrees\ of\ Freedom\ of\ the\ Treatment} = \frac{6}{2} = \mathbf{3}$$

Now that we've calculated both **MST & MSE**, let's see what our **ANOVA table** looks like before we move onto the **final step, calculating the F-value**.

Variation Source	Sum of Squares (SS)	Degrees of freedom (DF)	Mean Squares (MS)	F-Value
Treatment (Between)	6	2	3	
Error (Within)	6	6	1	
Total	12	8		

Comparing The Two Variation Types Against Each Other

Alright, now we're on to the final step of ANOVA, which is to **calculate the F-value** associated with our data set, so that we can make an **accept/reject decision** for our hypothesis test.

Recall that if the null hypothesis is true both MST & MSE will both approximate the population variance, and MST ≈ MSE.

We can **compare MST & MSE against each other** using the **F-test**, which is a ratio of the two variances in order to make an accept/reject decision on our null hypothesis.

$$F - value = \frac{MST}{MSE} = \frac{Mean\ Square\ of\ the\ Treatment}{Mean\ Square\ of\ the\ Error}$$

*If **the null hypothesis is true** and all sample means are equal, then MSE & MST will be approximate equal, and our **F-value will equal ~1.***

*If **the null hypothesis is false**, and one or more sample groups are not equal to the other sample groups, then MST >> MSE and our **F-value will be >>1**.*

Now, there might be slight differences between MST & MSE which makes MST > MSE, simply due to random chance. How do we make sure this random variability doesn't cause us to reject the null hypothesis incorrectly?

This is why we're using the **F-Test**, which is a test used to determine the **equality of two variances**, such as **MST & MSE.** The F-Test allows us to determine if our differences between MST & MSE are due to chance or if they are statistically significant.

Many ANOVA tables also include the **P-value** associated with the F-value. The P-value represents the **probability** of getting that big of a difference between MST & MSE (or bigger).

This P-value can be compared to your alpha risk to make an **accept/reject decision**.

If you're not using a statistical software that gives you a p-value, you can also look up the **critical F-statistics in the F-table**.

The **critical F-statistic** is based on your **Alpha risk and degrees of freedom** for the numerator (**MST**) & denominator (**MSE**). Using the example above, the DF of the treatment is 2, and the degrees of freedom of the Error is 6.

The F-Value can then be compared against the critical F_{crit} to determine if the calculated F_{stat} is a statistically significant result.

Below is our ANOVA table with the F-Value.

Variation Source	Sum of Squares	Degrees of freedom	Mean Squares	F-Value
Treatment (Between)	6	2	3	3
Error (Within)	6	6	1	
Total	12	8		

The F-Value = (MST / MSE)

Below is a quick visual example two different situations to demonstrate how MST & MSE can be used to determine if one sample has a different population mean.

On the left, **MST & MSE will be approximately equal**, as the 3 sample groups are roughly evenly distributed around the grand mean. In this scenario MST would be approximately equal to MSE, and our null hypothesis would likely be true.

On the right, one of our sample means is far out to the right implying that it is not equal to the other sample means. This means **MST will be much greater than MSE**, and our null hypothesis will likely be false.

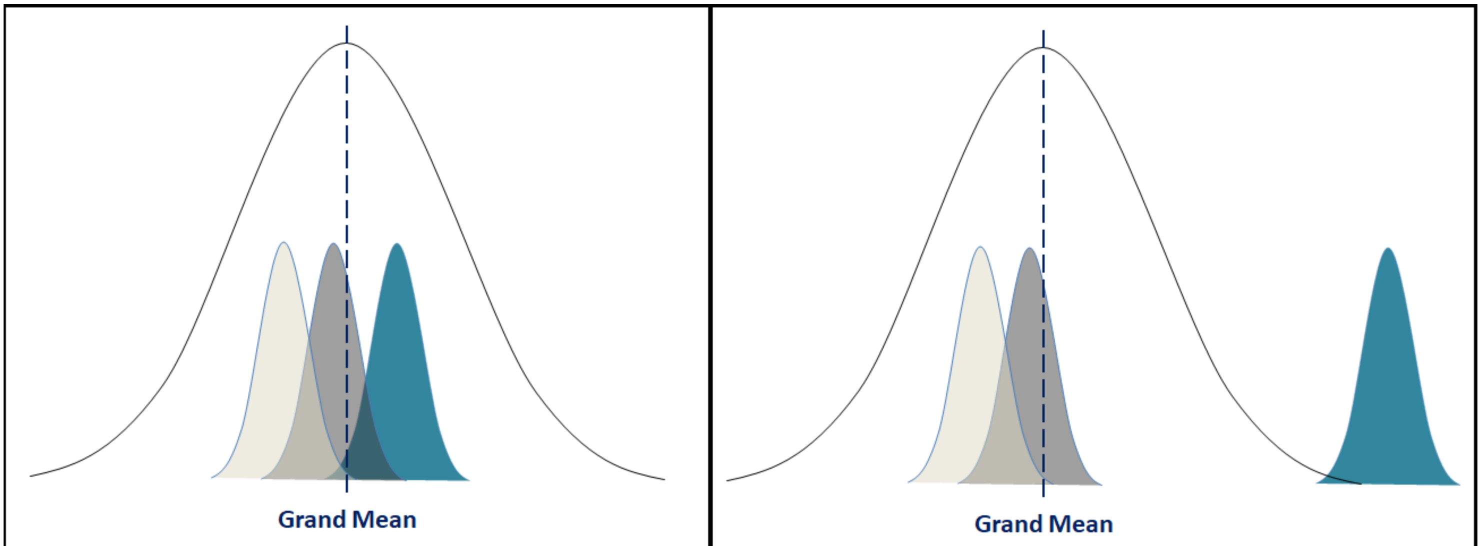

MST = MSE, Null is True	MST >> MSE, Null is False

One Way ANOVA Example

Ready for an example to practice?

Let's say you mold car bumpers, and you want to know if the injection temperature has an impact on the critical dimension of the car bumper - the length of the bumper.

In this experiment the **independent variable** (Factor) is the injection temperature and the **dependent variable** (response) is the length of the bumper.

Let's design an experiment where we'll **study our factor at four different levels** (temperatures) to see if this factor has an effect on the response variable. Let's measure four parts at each of the four levels and analyze the following results:

	Treatment Group #1	Treatment Group #2	Treatment Group #3	Treatment Group #4
Sample #1	58	61	61	60
Sample #2	56	62	57	60
Sample #3	57	59	59	57
Sample #4	58	59	60	56

Let's get started by reviewing the steps to an ANOVA Analysis.

Step 1 is to calculate the **Sum of Squares** for each of the sources of variation (**Treatment & Error**), and then adding them up to the **total variation** within the data set.

Step 2 is to calculate the **degrees of freedom (DF)** for each source of variation **(Treatment & error),** which add up to the **total degrees of freedom**.

Step 3 is to calculate the **Mean Squares** for each source of variation **(Treatment & Error)**, which is the **Mean Square of the Treatment (MST)**, and a **Mean Square of the Error (MSE)**.

Step 4 is to calculate the F-value associated with your **ANOVA** by taking the **ratio of MST to MSE**.

Step 5 is to **compare your F-value against the critical F-statistic**; or to compare your **P-value against your alpha risk.** Either of these comparisons will allow you to make an **Accept/Reject** decision for your null hypothesis.

You'll notice in these steps that we didn't start with a null hypothesis, because for ANOVA analysis, the null & alternative hypothesis are always the same.

The **Null Hypothesis** is that all of our sample mean values are equal, and the **Alternative Hypothesis** is that at least one mean value is different than the rest.

H_o: $\mu_a = \mu_b = \mu_c = \ldots\ldots \mu_k$ and H_a: **Not all means are equal**

Step 1 - The Sum of Squares

We're going to need to calculate the **Total Sum of Square (SS_{total})**, the **Sum of Squares of the Treatment (SS_t)**, and the **Sum of Squares of the Error (SS_e)**.

Before we calculate the sum of squares, we have to calculate the **sample mean for each group** and **the Grand Mean** of our data set.

	Treatment Group #1	Treatment Group #2	Treatment Group #3	Treatment Group #4
Sample #1	58	61	61	60
Sample #2	56	62	57	60
Sample #3	57	59	59	57
Sample #4	58	59	60	56
Sample Mean	57.25	60.25	59.25	58.25

The **grand mean** is the average of all data points which is **58.75.**

Now we can use the **grand mean** to calculate the **Sum of Squares of the Treatment (SS_t).**

$$SS_t = \sum n(\bar{X}_1 - GM)^2 + n(\bar{X}_2 - GM)^2 + n(\bar{X}_3 - GM)^2 + \ldots n(\bar{X}_i - GM)^2$$

$$SS_t = \sum 4(57.25 - 58.75)^2 + 4(60.25 - 58.75)^2 + 4(59.25 - 58.75)^2 + 4(58.25 - 58.75)^2 = 20$$

Now let's move on to the **Sum of Squares of the Error (SS_e)** which can be done using the equation below:

$$SS_e = \sum (X_i - \bar{X})^2$$

So we're comparing all 16 individual values (X_i) against the sample mean for that treatment.

$$\begin{aligned} SS_e = \sum &(58 - 57.25)^2 + (56 - 57.25)^2 + (57 - 57.25)^2 + (58 - 57.25)^2 \\ &+(61 - 60.25)^2 + (62 - 60.25)^2 + (59 - 60.25)^2 + (59 - 60.25)^2 \\ &+(61 - 59.25)^2 + (57 - 59.25)^2 + (59 - 59.25)^2 + (60 - 59.25)^2 \\ &+(60 - 58.25)^2 + (60 - 58.25)^2 + (57 - 58.25)^2 + (56 - 58.25)^2 = 31 \end{aligned}$$

I've color coded this equation so you can see how the different treatment groups are considered within this equation.

Now let's move on to the **Total Sum of Squares (SS_{total})**, which we could calculate by adding up the **Sum of Square of the Treatment (20)** and the **Sum of Squares of the Error (31).** Or we could calculate it using the following equation:

$$SS_{total} = \sum (X_i - GM)^2$$

This looks similar to the sum of squares of the error calculation; however, we're now comparing **all individual values (X_i)** against the **grand mean** (**58.75**).

$$\begin{aligned} SS_{total} = \sum &(58 - 58.75)^2 + (56 - 58.75)^2 + (57 - 58.75)^2 + (58 - 58.75)^2 \\ &+(61 - 58.75)^2 + (62 - 58.75)^2 + (59 - 58.75)^2 + (59 - 58.75)^2 \\ &+(61 - 58.75)^2 + (57 - 58.75)^2 + (59 - 58.75)^2 + (60 - 58.75)^2 \\ &+(60 - 58.75)^2 + (60 - 58.75)^2 + (57 - 58.75)^2 + (56 - 58.75)^2 \end{aligned}$$

$$SS_{total} = \sum 4 + 1 + 0 + 1 + 0 + 1 + 0 + 1 + 4 = 51$$

Alright, so we've knocked out the **Sum of Squares** portion of the ANOVA Table, let's move on to the **Degrees of Freedom column**.

Step 2 - The Degrees of Freedom

Step 2 in the ANOVA process is to calculate the **degrees of freedom (DF)** for each **source of variation (Treatment & error),** which add up to the **total degrees of freedom**.

So we will start by calculating the **total degrees of freedom (DF_{total})** associated with the entire sample data set.

$$DF_{total} = DF_{error} + DF_{treatment}$$

Then we will break this down into the **degrees of freedom for the treatment ($DF_{treatment}$),** and the **degrees of freedom for the error (DF_{error}).**

$$DF_{total} = N - 1 = 16 - 1 = 15$$

The **DF of the Treatment** is the next easiest to calculate - It's the number of treatments minus 1. The letter "a" is often used to denote the number of treatments, but this can vary between textbooks.

$$DF_{treatment} = 4 - 1 = 3$$

The DF of the Error is often calculated by taking the total degrees of freedom, and subtracting the treatment degrees of freedom.

$$DF_{error} = DF_{total} - DF_{treatment} = (N-1) - (a-1)$$

$$DF_{error} = (16 - 1) - (4 - 1) = 12$$

So for the Treatment and Error we've calculated the **Sum of Squares** for each and the **Degrees of Freedom** for each, let's see what that looks like in our ANOVA table.

Variation Source	Sum of Squares	Degrees of freedom	Mean Squares	F-Value
Treatment (Between)	20	3		
Error (Within)	31	12		
Total	51	15		

Now it's time to use that information to move to step 3 in the ANOVA process, **calculating the Mean Squares**.

Step 3 - The Mean Squares

Alright, on to step 3 of the ANOVA analysis, which is the calculation of our **Mean Squares**, where again we'll calculate a **Mean Square of the Treatment (MST)**, and a **Mean Square of the Error (MSE)**.

$$\boldsymbol{Mean\ Square\ of\ the\ Error} = \frac{Sum\ of\ Squares\ of\ the\ Error}{Degrees\ of\ Freedom\ of\ the\ Error} = \frac{31}{12} = 2.583$$

$$\boldsymbol{Mean\ Square\ of\ the\ Treatment} = \frac{Sum\ of\ Squares\ of\ the\ Treatment}{Degrees\ of\ Freedom\ of\ the\ Treatment} = \frac{20}{3} = 6.667$$

Step 4 - The F - Statistic

Alright, now we're on to one of the last steps in ANOVA, which is to **calculate the F-statistic** so that we can make an **accept/reject decision** for our hypothesis test.

Recall that if the null hypothesis is true both MST & MSE will both approximate the population variance, and MST ≈ MSE.

$$F-statistic = \frac{MST}{MSE} = \frac{Mean\ Square\ of\ the\ Treatment}{Mean\ Square\ of\ the\ Error} = \frac{6.667}{2.583} = 2.581$$

Here's what our ANOVA table looks like with all of the fields filed in.

Variation Source	Sum of Squares	Degrees of freedom	Mean Squares	F-Value
Treatment (Between)	20	3	6.667	2.581
Error (Within)	31	12	2.583	
Total	51	15		

Step 5 - Accept/Reject the Null Hypothesis

Step 5 is to **compare your F-value against the critical F-statistic**; or to compare your **P-value against your alpha risk.**

Either of these comparisons will allow you to make an **Accept/Reject** decision for your null hypothesis.

*If **the null hypothesis is true** and all sample means are equal, then MSE & MST will be approximate equal, and our **F-value will equal ~1.***

*If **the null hypothesis is false**, and one or more sample groups are not equal to the other sample groups, then MST >> MSE and our **F-value will be much greater than 1**.*

What about this F-value?

Should we reject the null hypothesis?

For this example, I'll show you have to look up the **critical F-value in the F-distribution**, which can be looked up using the **degrees of freedom** for **MST (3)** & **MSE (12)**, and our **Alpha (α) risk of 10%.**

$$F_{crit} = F_{.05(3,12)} = 3.490$$

Because our calculate F-statistic (2.581) is less than our critical F-value (3.490), we must fail to reject the null hypothesis.

Two Way ANOVA

How do you feel about One Way ANOVA?

Are you ready to move on to Two Way ANOVA?

Before we jump into the differences between One Way and Two Way ANOVA, it's important to note that all of the assumptions of One Way ANOVA apply to Two Way ANOVA.

On to the differences!

So in One Way ANOVA we were analyzing data that was associated with one single **Factor** or **Independent Variable**.

In Two Way ANOVA, we're going to analyze data where **two factors** are varied and studied.

In some textbooks you'll see this as **Factor A and Factor B**.

Sometimes the second factor is called a **"Block"**, because you can use it within your experiment to "block" out a certain factor that might be contributing to the **error variation**.

Similar to One Way ANOVA, Two way ANOVA allows you to **analyze the "main effects" of each factor** being varied, by carving out the variation associated with each factor.

This is a new term we haven't used yet - **Main Effects** - and it's meant to denote the variation associated with a single factor. For example, you can have the **Main Effects of Factor A**, and the **Main Effects of Factor B**.

When you've got **two factors**, you can also now study the **interaction effect between factors.** We will talk more about this in the chapter on Designed Experiments.

You can see how the **ANOVA table** grows in complexity when we move to two factors.

Variation Source	Sum of Squares (SS)	Degrees of freedom (DF)	Mean Squares (MS)	F-Value
Factor A	Sum of Squares of the Factor A (SS_a)	DF of the Factor A (DF_a)	Mean Squares of the Factor A (MS_a)	$= MS_a / MSE$
Factor B	Sum of Squares of the Factor B (SS_b)	DF of the Factor B (DF_b)	Mean Squares of the Factor A (MS_b)	$= MS_b / MSE$
Interactions AxB	Sum of Squares of the Interaction (SS_{axb})	DF of the Interaction (DF_{axb})	Mean Squares of the Interaction (MST_{axb})	$= MS_{axb} / MSE$
Error	Sum of Squares of the Error (SS_e)	DF of the Error (DF_e)	Mean Squares of the Error (MSE)	
Total	Total Sum of Squares (SS_{total})	Total DF (DF_{Total})		

I've shown the interactions above, however it's important to note that if you're interested in studying the interactions, you must have multiple **replicates** for each treatment. A replicate is the number of observations per factor.

If you only have 1x replicate (observation) per factor, then there won't be enough degrees of freedom left over to calculate the interactions effect, and any variability due to the interaction between factors will fall into the error term.

Conclusion

Alright!!! Are you glad that's done??

Let's recap quickly.

Ok, so ANOVA stands for **AN**alysis **Of** **VA**riance and it is a **hypothesis test** used to compare the means of 2 or more groups.

The first we covered are the **assumptions** associated with ANOVA, which include:

- The Population being studied is Normally Distributed
- The Variance is the same between the various treatments (Homogeneity of Variance)
- There is Independence between sample observations

Second was the **common terms & definitions** within ANOVA which included **Independent variable, dependent variable, factors and treatments (levels).**

Third, we answered - **Why Does ANOVA use Variance to Test Mean Values?**

Variation Source	Sum of Squares (SS)	Degrees of freedom (DF)	Mean Squares (MS)	F-Value
Treatment (Between)	Sum of Squares of the Treatment (SS_t)	DF of the Treatment (DF_t)	Mean Squares of the Treatment (MST)	= MST / MSE
Error (Within)	Sum of Squares of the Error (SS_e)	DF of the Error (DF_e)	Mean Squares of the Error (MSE)	
Total	Total Sum of Squares (SS_{total})	Total DF (DF_{Total})		

Then we jumped into the **basics of ANOVA** including the **Sum of Squares, Degrees of Freedom, Mean Squares**, **F-value**.

For the **Sum of Squares** we reviewed the following equations:

$$\textit{Total Sum of Squares}: SS_{total} = \sum (X_i - GM)^2$$

$$\textit{Sum of Squares of the Error}: SS_e = \sum (X_i - \bar{X})^2$$

$$\textit{Sum of Squares of the Treatment} = SS_t = \sum n(\bar{X}_i - GM)^2$$

Total Sum of Squares (SS_{total}) = Sum of Square of the Treatment (SS_t) + Sum of Squares of the Error (SS_e)

For the **Degrees of Freedom** we reviewed the following equations:

$$DF_{total} = DF_{error} + DF_{treatment}$$

$$DF_{total} = N - 1$$

$$DF_{treatment} = a - 1$$

$$DF_{error} = (N-1) - (a-1) = a(n-1)$$

We then combined the **Sum of Squares** and **Degrees of Freedom** to calculate our **Mean Squares.**

$$\boldsymbol{Mean\ Square\ of\ the\ Error} = \frac{Sum\ of\ Squares\ of\ the\ Error}{Degrees\ of\ Freedom\ of\ the\ Error}$$

$$\boldsymbol{Mean\ Square\ of\ the\ Treatment} = \frac{Sum\ of\ Squares\ of\ the\ Treatment}{Degrees\ of\ Freedom\ of\ the\ Treatment}$$

We then combined MSE & MST into our final calculation of an F-statistic:

$$\boldsymbol{F - statistic = \frac{MST}{MSE} = \frac{Mean\ Square\ of\ the\ Treatment}{Mean\ Square\ of\ the\ Error}}$$

To determine if our null hypothesis is true or false, we can compare this F-statistic against a [critical f-value from the NIST Tables].

Looking up this critical value will require you to have an **Alpha risk** for your hypothesis test and combine that with the **Degrees of Freedom for your MST & MSE**.

If your f-stat is less than your critical f-value, then you must fail to reject the null hypothesis.

If your f-stat is greater than your critical f-value, then you reject the null hypothesis, and accept the alternative hypothesis that one of the sample means is not equal to the rest.

Lastly, we covered the nuances of going from a **One Way ANOVA to a Two Way ANOVA**.

In Two Way ANOVA we introduced new concepts like **Main Effects of factors** and **Interactions between factors**. You can see how the **ANOVA table** grows in complexity when we move to two factors.

Variation Source	Sum of Squares (SS)	Degrees of freedom (DF)	Mean Squares (MS)	F-Value
Factor A	Sum of Squares of the Factor A **(SS_a)**	DF of the Factor A **(DF_a)**	Mean Squares of the Factor A **(MS_a)**	**$= MS_a / MSE$**
Factor B	Sum of Squares of the Factor B **(SS_b)**	DF of the Factor B **(DF_b)**	Mean Squares of the Factor A **(MS_b)**	**$= MS_b / MSE$**
Interactions AxB	Sum of Squares of the Interaction **(SS_{axb})**	DF of the Interaction **(DF_{axb})**	Mean Squares of the Interaction **(MST_{axb})**	**$= MS_{axb} / MSE$**
Error	Sum of Squares of the Error **(SS_e)**	DF of the Error **(DF_e)**	Mean Squares of the Error **(MSE)**	
Total	Total Sum of Squares **(SS_{total})**	Total DF **(DF_{Total})**		