

Chapter 13: Comparing Three or More Means

13.1 Comparing Three or More Means (One-Way Analysis of Variance)

In [Section 11.3](#), we compared two means from independent populations. What if we have more than two means to compare?

Suppose you own a chain of four boutique resale clothing shops. All four have been open for at least three years, and you want to do some analysis regarding their performance. You suspect that the managers at the four shops are not hiring staff of equal quality, so you take a sample of the weekly sales amounts for the employees at each location. You want to determine if any location is performing statistically lower than any of the others.

Using the strategies from [Section 11.3](#), this would mean **six hypothesis tests** - one for each comparison (shop 1 vs. shop 2, shop 1 vs. shop 3, etc.). There has to be a better way! This is what One-Way Analysis of Variance is.

If you're ready to begin, just click on the "start" link below, or one of the section links on the left.



Source: stock.xchng

[:: start ::](#)



Section 13.1: Comparing Three or More Means (One-Way ANOVA)

13.1 Comparing Three or More Means (One-Way Analysis of Variance)

Objectives


By the end of this lesson, you will be able to...

1. verify the requirements to perform a one-way ANOVA
2. test a hypothesis regarding three or more means using one-way ANOVA

What is One-Way ANOVA?


Suppose you're an instructor teaching three sections of the same course. You suspect that the three sections are not equivalent, but you can't be sure. Since the students won't always earn the same score on every exam, we can treat an individual exam given to all the students in each section as a "sample" from each section.

Suppose the results from the first turn out as follows:



We can clearly see that these appear to be from very different populations - the means are very different, and there's very little variation within each group. With that description, we can be pretty sure they're different.

What about a different situation?



In this case, we can see that the means do appear to be different, but not by much. And within each sample, there is a lot of variation, so the difference in the means could just be due to the wide variation within each group.

The point here is that we can't just consider the differences in the means - whether those differences are significant or not depends on the standard deviations (and sample sizes, of course).

This is the basic idea behind One-Way ANOVA. It's **one-way**, because we're focusing on a single characteristic (time of class period, in our example above). And the ANOVA stands for **analysis of variance**. While it may seem odd that the title refers to *variance* when we're actually comparing *means*, it's actually because of a significant assumption - to perform the statistical analysis.

To perform the test, we focus on the *between-sample variation* (between the means) and the *within-sample variation* (i.e. the standard deviation). If the former is large in comparison to the latter, we can say that one of the means must be different. The test statistic that we use is another F-statistic, and it's the ratio of these two variations:

$$F_0 = \frac{\text{between-sample variability}}{\text{within-sample variability}}$$

In order to analyze this statistic, though, there are several requirements that need to be met.

Verifying the Requirements

In order to perform one-way ANOVA, the following requirements must be met:

Requirements to Perform a One- Way ANOVA Test

1. There must be k simple random samples, one from each of k populations or a randomized experiment with k treatments.
2. The k samples must be independent of each other; that is, the subjects in one group cannot be related in any way to subjects in a second group.
3. The populations must be normally distributed.
4. The populations must have the same variance; that is, each treatment group has population variance σ^2 .

Luckily, the procedure is *robust*, so slight variations from these criteria are OK. In particular, a good rule of thumb is that as long as the largest variance is no more than double the smallest, we can assume point #4 above has been satisfied. If not, there are other tests that can be performed, but they're beyond the scope of this course.

Example 1


Problem: Referring back to the example above, does it appear that the conditions are met to perform one-way ANOVA?

Solution: To answer this, let's consider the box plots we were provided:

9am: $\bar{x} = 78.8, s = 9.6$

10am: $\bar{x} = 80.4, s = 10.1$

1pm: $\bar{x} = 81.1, s = 8.7$



Granted, the information is a little imprecise here, but we do have 3 samples that are independent (3 different sections). Based on the boxplots, the scores do appear to be normally distributed, though without a histogram that's difficult to tell precisely. And the largest standard deviation of 10.1 is not more than double the smallest of 8.7.

So yes, it does appear that the conditions have been met and one-way ANOVA could be performed.

Performing a Hypothesis Test Regarding Three or More Means Using One-Way ANOVA

Step 1: State the null and alternative hypotheses.

$$H_0: \mu_1 = \mu_2 = \mu_3 = \dots$$

H_1 : At least one of the means is different

Step 2: Decide on a level of significance, α .

Step 3: Compute the test statistic, $F_0 = \frac{\text{between-sample variability}}{\text{within-sample variability}}$ (using StatCrunch).

Step 4: Determine the P -value.

Step 5: Reject the null hypothesis if the P -value is less than the level of significance, α .

Step 6: State the conclusion.

Hypothesis Testing Regarding Three or More Means Using One-Way ANOVA with StatCrunch

1. Either enter the raw data in separate columns for each sample or treatment, or enter the value of the variable in a single column with indicator variables for each sample or treatment in a second column.
2. Go to **Stat > ANOVA > One Way**.
3. If the raw data are in separate columns, select "Compare selected columns" and then click the columns you wish to compare. If the raw data are in a single column, select "Compare values in a single column" and then choose the column that contains the value of the variables and the column that indicates the treatment or sample.
4. Click **Calculate**.

The results should appear.

Example 2

Problem: Referring back to the example above, use the data provided in the link below.

[test scores by section \(CSV\)](#)

Is there enough evidence at the 5% level of significance to support the claim that one of the classes is performing differently from the others?

Solution:

Assuming again the conditions have been met, we have the following results:

Step 1:

$$H_0: \mu_1 = \mu_2 = \mu_3 = \dots$$

H_1 : At least one of the means is different

Step 2: $\alpha = 0.05$ (given)

Step 3: (we'll use StatCrunch)

Step 4: Using StatCrunch:

Analysis of Variance results:

Responses stored in var2.

Factors stored in var1.

Factor means

var1	n	Mean	Std. Dev.	Std. Error
10am	24	80.375	10.124496	2.0666542
1pm	24	81.083336	8.722368	1.7804459
9am	26	78.84615	9.611212	1.8849137

ANOVA table

Source	df	SS	MS	F-Stat	P-value
Treatments	2	65.819214	32.909607	0.36413267	0.6961
Error	71	6416.843	90.37807		
Total	73	6482.662			

Step 5: Since the P -value $> \alpha$, we do not reject the null hypothesis.

Step 6: There is not enough evidence at the 5% level of significance to support the claim that the mean score from one of these sections is different from the others.

<< previous section | next section >>

1 2 3 4 5 6 7 8 9 10 11 12 13



This work is licensed under a Creative Commons License.