CHAPTER 15 Analysis of Variance

Chapter Outline

15.1 The One-Way ANOVA Test

Learning Objectives

- Understand the shortcomings of comparing multiple means as pairs of hypotheses.
- Understand the steps of the ANOVA method and the method's advantages.
- Compare the means of three or more populations using the ANOVA method.
- Calculate pooled standard deviations and confidence intervals as estimates of standard deviations of populations.

Introduction

Previously, we have discussed analyses that allow us to test if the means and variances of two populations are equal. Suppose a salesperson wants to compare the level of satisfaction of customers for four difference insurance companies. Our question is:

"Is there a difference in satisfaction scores across the four difference insurance companies?”

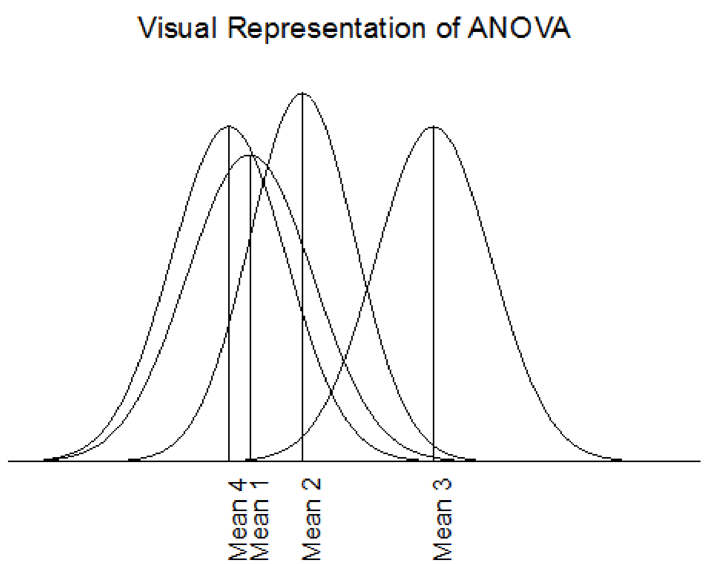

FIGURE 15.1

The satisfaction scores for a sample of customers for each insurance company are recorded:

TABLE 15.1:

Company 1	Company 2	Company 3	Company 4
3.2	4.2	5.4	4.5
3.5	3.7	4.6	3.8
2.7	3.4	4	4.1

TABLE 15.1: (continued)

4.1	4.3	5.3	3.1
3.1	3.9	4.7	4.2
3.7	4.1	4.2	3.4
4.2	3.1	4.9	4.2
3.6	4.5	4.7	4.5

We could conduct a series of t-tests to determine if any of the sample means differ. However, this would be tedious and has a major flaw, which we will discuss shortly. Instead, we use something called the **Analysis of Variance (ANOVA)**, which allows us to test the hypothesis that multiple population means and variances of scores are equal.

The Null and Alternative hypotheses for a one-way ANOVA can be written as:

H_0 : Means of all factor levels are equal

H_A : At least one factor level has a different mean

The ANOVA can be used when we want to test the means of three or more populations aat once. Theoretically, we could test hundreds of population means using this procedure. This ANOVA is technically called "one-way" as it has just one main grouping factor: company. In the next chapter we'll see how we can have an ANOVA with more than one factor.

Shortcomings of Comparing Multiple Means Using Previously Explained Methods

Why should we learn a new test called ANOVA, when we could just conduct a series of t-tests? To answer our question, we could just run six different independent samples t-tests (company 1 vs. company 2; company 1 vs. company 3; company 1 vs. company 4; company 2 vs. company 3; company 2 vs. company 4; and company 3 vs. company 4). This would be tedious, but we could use a computer to compute these quickly and easily. It turns out this is a very bad idea, and has a major flaw:

When more than one t-test is run, each at its own level of significance, the probability of making one or more Type I errors multiplies exponentially. Recall that a Type I error occurs when we reject the null hypothesis when we should not. The level of significance, α, is the probability of a Type I error in a single test. So, for a single t-test in our example, with an α of 0.05, we have a Type I error probability of 5%.

When testing more than one pair of samples, the probability of making at least one Type I error is:

$$1-(1-\alpha)^c$$

where α is the level of significance for each t-test and c is the number of independent t-tests. Using the example from the introduction, if our salesperson conducted separate t-tests to examine the means of the populations, they would have to conduct six separate t-tests. If they performed these tests with $\alpha = 0.05$, the probability of committing a Type I error is not 0.05 as one would initially expect. Instead, it would be 0.265 or 26.5%. So the salesperson has a 26.5% chance of committing at least one Type I error.

Assumptions of the ANOVA test

Before we can use the one-way ANOVA, we must see if we satisfy some assumptions, just like we had in our previous hypothesis tests:

1. All observations are independent of one another and randomly selected from the population which they represent.
2. The population at each factor level is approximately normal.
3. The variances for each factor level are approximately equal to one another.

The Steps of the ANOVA Method

With the ANOVA method, we are actually analyzing the total variation of the scores, including the variation of the scores within the groups and the variation between the group means. Since we are interested in two different types of variation, we first calculate each type of variation independently and then calculate the ratio between the two –called an F-value.

Just like our z-score, our t-test, and chi-square tests, ANOVA has its own distribution that we need to use, called an F-distribution to set our critical values and test our hypothesis. Just like the t-distribution and the chi-square distribution which use degrees of freedom, the F-distribution also relies on degrees of freedom. Since the F-value is actually a ratio of two different sources of variance, we'll need two different degrees of freedom, which we'll talk about when we get to our example.

When using the ANOVA method, we are testing the null hypothesis that the means and the variances of our samples are equal. When we conduct a hypothesis test, we are testing the probability of obtaining an extreme F-statistic by chance. If we reject the null hypothesis that the means and variances of the samples are equal, and then we are saying that the difference that we see could not have happened just by chance.

To test a hypothesis using the ANOVA method, there are several steps that we need to take. These steps are to simply conduct the ANOVA –not the steps of Hypothesis Testing that we've seen previously. To help us in completing those steps, we need to employ a nice little tool called the ANOVA table:

TABLE 15.2:

Source	**SS**	**df**	**MS**	**F**	**F critical**
Between					
Within					
Total					

Notice on the left side, there's a column called "source." This column lists where the variation in the test is coming from: Between the groups, within the groups, or all the variance for all the observations (Total). The columns may also be familiar to you as well: SS is the Sums of Squares (hint: we used Sums of Squares to calculate Standard Deviation and Variance) and df is the Degrees of Freedom. We'll explain the other columns shortly.

Working through ANOVA

TABLE 15.3:

Company 1	**Company 2**	**Company 3**	**Company 4**
3.2	4.2	5.4	4.5
3.5	3.7	4.6	3.8
2.7	3.4	4	4.1
4.1	4.3	5.3	3.1
3.1	3.9	4.7	4.2
3.7	4.1	4.2	3.4
4.2	3.1	4.9	4.2
3.6	4.5	4.7	4.5

Let's use the ANOVA table with our ANOVA calculation steps and some data from our company satisfaction example from above.

1. **Calculate the total sum of squares (SS_T).** This is the difference between each score and the grand mean. We use the following formula:

   $SS_T = \sum(y - \bar{y}^2) = \sum y^2 - \frac{(\sum y)^2}{N} = 12.965$

   Where:
   y = each observation
   N = total number of scores
   $\bar{y}$ = grand mean (mean of all scores)

2. **Calculate the sum of squares between (SS_B).** We use the following formula:

   $SS_B = \sum n_k \left(\bar{y}_k - \bar{y}\right)^2 = 6.166$

   Where:
   k = the number of groups
   n_k = the number of scores in group k
   $\bar{y}_k$ = the mean of group k

3. **Find the sum of squares within groups (SS_W) by subtracting:**

   $SS_W = SS_T - SS_B = 6.799$

4. **Next solve for degrees of freedom for the test:**

   $df_{Total} = N - 1 = 31$
   $df_{Between} = k - 1 = 3$
   $df_{Within} = N - k = 28$

5. **Using the values, you can now calculate the Mean Squares Between (MS_B) and Mean Squares Within (MS_W) using the relationships below:**

   $MS_B = \frac{SS_B}{df_{Between}} = 2.055$

   $MS_W = \frac{SS_W}{df_{Within}} = 0.243$

6. **Finally, calculate the F statistic using the following ratio:**

   $F = \frac{MS_B}{MS_W} = 8.457$

7. **It is easy to fill in the Table from here – and also to see that once the SS and df are filled in, the remaining values in the table for MS and F are simple calculations.**

TABLE 15.4:

Source	SS	df	MS	F	F critical
Between	6.166	3	2.055	8.457	

TABLE 15.4: (continued)

Within	6.799	28	0.243		
Total	12.965	31			

8. **Find F critical**. For our example, we have 3,28 degrees of freedom, so our F-critical value is 2.947.

TABLE 15.5:

Source	**SS**	**df**	**MS**	**F**	**F critical**
Between	6.166	3	2.055	8.457	2.947
Within	6.799	28	0.243		
Total	12.965	31			

Interpret the results of the hypothesis test. In ANOVA, the last step is to decide whether to reject the null hypothesis and then provide clarification about what that decision means. In our example of the insurance companies, our F-value from the ANOVA test is greater than the F-critical value, so we would reject our Null Hypothesis. We can conclude that the average customer satisfaction scores of the four insurance companies are not equal to one another – at least one of them is different from the others.

Apply our Knowledge

Example A

Let's take an example of a teacher who is testing multiple reading programs to determine the impact on student achievement. There are five different reading programs, and her 31 students are randomly assigned to one of the five programs. She collects the following data.

TABLE 15.6:

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

- **Hypothesis Step 1:** Clearly state the Null and Alternative Hypotheses.

  H_0 : Mean student achievement is the same across all five reading programs

  H_A : At least one reading program has a mean student achievement level that is different from the others.

- **Hypothesis Step 2:** Identify the appropriate significance level and confirm test assumptions.

  We'll choose the default significant level of 0.05.

  Check our assumptions:

1. All observations are independent of one another and randomly selected from the population which they represent.

   Yes

2. The population at each factor level is approximately normal.

   We assume student achievement level (in the population) is normally distributed. We can check boxplots here to confirm none of the samples are highly skewed.

3. The variances for each factor level are approximately equal to one another.

   *We check the variance of each group, and they range from 2.5 to 4.3. Since 4.3<2*2.5, we can proceed.*

- **Hypothesis Step 3:** Analyze the data.

We'll use the one-way ANOVA test here because we want to compare five different independent methods of reading programs.

Let's use the steps of the ANOVA to analyze the data:

$SS_T = 237.935$

$SS_B = 150.503$

$SS_W = 237.935 - 150.503 = 87.432$

We know that there are five reading methods, so (5-1) = 4 degrees of freedom.

So our MS_B solution:

$MS_B = \frac{SS_B}{m-1} = \frac{150.503}{4} = 37.626$

Remember the degrees of freedom for the Within source is the total observations minus the number of groups. For this example we have (31-5) = 26.

So our MS_W solution:

$MS_W = \frac{SS_W}{n-m} = \frac{87.432}{26} = 3.363$

Now we are ready to calculate the F-value test statistic:

$F = \frac{MS_B}{MS_W} = \frac{37.626}{3.363} = 11.188$

We can find the critical value of F with (4,26) degrees of freedom to be 2.743.

Putting it all in our ANOVA table:

TABLE 15.7:

Source	SS	df	MS	F	F critical
Between	150.503	4	37.626	11.188	2.743
Within	87.432	26	3.363		
Total	237.935	30			

- **Hypothesis Step 4:** Interpret your results.

Since our calculated F-value is greater than the F-critical, we can reject our Null Hypothesis. We conclude the all five reading program means are not equal to one another, but that there is at least one reading method score mean that is not like the others.

Now What?

Now that we've found a way to test the Null Hypothesis when we want to compare three or more population means, we need to learn one more step. What happens when we reject the Null Hypothesis? Are we done? Do we simply say: "There's a difference somewhere amongst these groups" and leave it at that? No.

If we have an ANOVA that rejects the Null Hypothesis we must find out where the difference lies –what group or groups are difference from one another. To do this we use a **post-hoc test**. Post-hoc is Latin for "after this" –so we literally run another analysis after the ANOVA shows a rejection of the Null Hypothesis.

The easiest post-hoc analysis (and there are several) to run is called the Bonferroni post-hoc analysis. This is really just all possible t-tests to compare the groups. Now, you might want to bring up the problem of inflating the Type I error with all those t-tests –and you would be right. That's where Bonferroni comes in: it's simply a correction to the level of significance, α, when we run these post-hoc tests.

Example B

Let's use our Insurance Company example:

We saw that we rejected the Null Hypothesis, so we are allowed to run the post-hoc tests to discover more about the difference in the group means. We have four groups, so we'll need to run: (4 * (4-1)) / 2) = 6 group comparison t-tests. The Bonferroni correction simply divides the significance level we used for the ANOVA of 0.05 by 6. So our new significance level –just for the post-hoc comparisons –is 0.008.

If we were to do these post-hoc comparisons by hand for our example, we would need to run 6 independent samples t-tests and find the critical value for each comparison, based on the degrees of freedom for the comparison and the new significance level of 0.008.

Most of the time, these comparisons are done using technology, so we can take advantage of the p-value reported for each t-test. Using the p-value reported, we would compare that value to the new Bonferroni corrected significance level. If the reported p-value of the comparison was *less than* the corrected significance level, then we reject the Null hypothesis that the comparison means are equal.

Here are the results of each comparison, along with the mean of the groups, the t-value, the comparison df, t-critical value, and the p-value:

TABLE 15.8:

Comparison	Mean A	Mean B	t-value	df	t-critical	p-value
Company 1 - Company 2	3.5	3.9	-1.58	7	3.16	0.137
Company 1 - Company 3	3.5	4.7	-4.90	7	3.16	0.000
Company 1 - Company 4	3.5	4.0	-1.82	7	3.16	0.089
Company 2 - Company 3	3.9	4.7	-3.44	7	3.16	0.004
Company 2 - Company 4	3.9	4.0	-0.31	7	3.16	0.765

TABLE 15.8: (continued)

Company 3 - Company 4	4.7	4.0	3.03	7	3.16	0.009

Our post-hoc comparisons from the above table show us that Company 3 was significantly higher in customer satisfaction than Company 1 and Company 2, but not Company 4. Also, no other Company was significantly different from the others.

Lesson Summary

When testing multiple independent samples to determine if they come from the same population, we could conduct a series of separate t-tests in order to compare all possible pairs of means. However, a more precise and accurate analysis is the Analysis of Variance (ANOVA).

In ANOVA, we analyze the total variation of the scores, including the variation of the scores within the groups, the variation between the group means, and the total mean of all the groups (also known as the grand mean).

In this analysis, we calculate the F-value, which is the ratio of mean of squares between groups divided by the mean of squares within groups.

If we are able to reject our Null Hypothesis, we continue on, conducting post-hoc analyses to discover where the difference in the sample means lies.

Review Questions

1. What does the ANOVA acronym stand for?
2. If we are testing whether pairs of sample means differ by more than we would expect due to chance using multiple t-tests, the probability of making a type I error would ___.
3. In the ANOVA method, we use the ___ distribution.

   (a) Student's t-

   (b) normal

   (c) F-

4. In the ANOVA method, we complete a series of steps to evaluate our hypothesis. Put the following steps in chronological order.
   a. Calculate the mean squares between groups and the mean squares within groups.
   b. Determine the critical values in the F-distribution.
   c. Evaluate the hypothesis.
   d. Calculate the test statistic.
   e. State the null hypothesis.
5. A school psychologist is interested in whether or not teachers affect the anxiety scores among students taking the AP Statistics exam. The data below are the scores on a standardized anxiety test for students with three different teachers.

TABLE 15.9: Teacher's Name and Anxiety Scores

Ms. Jones	Mr. Smith	Mrs. White

TABLE 15.9: (continued)

Ms. Jones	Mr. Smith	Mrs. White
8	23	21
6	11	21
4	17	22
12	16	18
16	6	14
17	14	21
12	15	9
10	19	11
11	10	
13		

a. State the null hypothesis.
b. Using the data above, fill out the missing values in the table below.

TABLE 15.10:

	Ms. Jones	Mr. Smith	Mrs. White	Totals
Number (n_k)			8	=
Total (T_k)		131		=
Mean ($\bar{x}$)		14.6		=
Sum of Squared Obs. ($\sum_{i=1}^{n_k} x_{ik}^2$)				=
Sum of Obs. Squared/Number of Obs. ($\frac{T_k^2}{n_k}$)				=

c. What is the value of the mean squares between groups, MS_B?
d. What is the value of the mean squares within groups, MS_W?
e. What is the F-ratio of these two values?
f. With $\alpha = 0.05$, use the F-distribution to set a critical value.
g. What decision would you make regarding the null hypothesis? Why?

15.2 The Two-Way ANOVA Test

Learning Objectives

- Understand the differences in situations that allow for one-way or two-way ANOVA methods.
- Know the procedure of two-way ANOVA and its application through technological tools.
- Understand completely randomized and randomized block methods of experimental design and their relation to appropriate ANOVA methods.

Introduction

In the previous section, we discussed the one-way ANOVA method, which is the procedure for testing the null hypothesis that the population means and variances of a single independent variable are equal. Sometimes, however, we are interested in testing the means and variances of more than one independent variable. Say, for example, that a researcher is interested in determining the effects of different dosages of a dietary supplement on the performance of both males and females on a physical endurance test. The three different dosages of the medicine are low, medium, and high, and the genders are male and female. Analyses of situations with two independent variables, like the one just described, are called two-way ANOVA tests.

TABLE 15.11: Mean Scores on a Physical Endurance Test for Varying Dosages and Genders

	Dietary Supplement Dosage	**Dietary Supplement Dosage**	**Dietary Supplement Dosage**	
	Low	Medium	High	Average
Female	35.6	49.4	71.8	52.3
Male	55.2	92.2	110.0	85.8
Average	45.4	70.8	90.9	

There are several questions that can be answered by a study like this, such as, "Does the medication improve physical endurance, as measured by the test?" and "Do males and females respond in the same way to the medication?"

While there are similar steps in performing one-way and two-way ANOVA tests, there are also some major differences. In the following sections, we will explore the differences in situations that allow for the one-way or two-way ANOVA methods, the procedure of two-way ANOVA, and the experimental designs associated with this method.

The Differences in Situations that Allow for One-way or Two-Way ANOVA

As mentioned in the previous lesson, ANOVA allows us to examine the effect of a single independent variable on a dependent variable (i.e., the effectiveness of a reading program on student achievement). With **two-way ANOVA**, we are not only able to study the effect of two independent variables (i.e., the effect of dosages and gender on the results of a physical endurance test), but also the interaction between these variables. An example of interaction between the two variables gender and medication is a finding that men and women respond differently to the medication.

We could conduct two separate one-way ANOVA tests to study the effect of two independent variables, but there are several advantages to conducting a two-way ANOVA test.

- *Efficiency.* With simultaneous analysis of two independent variables, the ANOVA test is really carrying out two separate research studies at once.
- *Control.* When including an additional independent variable in the study, we are able to control for that variable. For example, say that we included IQ in the earlier example about the effects of a reading program on student achievement. By including this variable, we are able to determine the effects of various reading programs, the effects of IQ, and the possible interaction between the two.
- *Interaction.* With a two-way ANOVA test, it is possible to investigate the interaction of two or more independent variables. In most real-life scenarios, variables do interact with one another. Therefore, the study of the interaction between independent variables may be just as important as studying the interaction between the independent and dependent variables.

When we perform two separate one-way ANOVA tests, we run the risk of losing these advantages.

Two-Way ANOVA Procedures

There are two kinds of variables in all ANOVA procedures-dependent and independent variables. In one-way ANOVA, we were working with one independent variable and one dependent variable. In two-way ANOVA, there are two independent variables and a single dependent variable. Changes in the dependent variables are assumed to be the result of changes in the independent variables.

In one-way ANOVA, we calculated a ratio that measured the variation between the two variables (dependent and independent). In two-way ANOVA, we need to calculate a ratio that measures not only the variation between the dependent and independent variables, but also the interaction between the two independent variables.

Before, when we performed the one-way ANOVA, we calculated the total variation by determining the variation within groups and the variation between groups. Calculating the total variation in two-way ANOVA is similar, but since we have an additional variable, we need to calculate two more types of variation. Determining the total variation in two-way ANOVA includes calculating: variation within the group (within-cell variation), variation in the dependent variable attributed to one independent variable (variation among the row means), variation in the dependent variable attributed to the other independent variable (variation among the column means), and variation between the independent variables (the interaction effect).

The formulas that we use to calculate these types of variations are very similar to the ones that we used in the one-way ANOVA. For each type of variation, we want to calculate the total sum of squared deviations (also known as the sum of squares) around the grand mean. After we find this total sum of squares, we want to divide it by the number of degrees of freedom to arrive at the mean of squares, which allows us to calculate our final ratio. We could do these calculations by hand, but we have technological tools, such as computer programs that can compute these figures much more quickly and accurately than we could manually.

The process for determining and evaluating the null hypothesis for the two-way ANOVA is very similar to the same process for the one-way ANOVA. However, for the two-way ANOVA, we have additional hypotheses, due to the additional variables. For two-way ANOVA, we have three null hypotheses:

1. In the population, the means for the rows equal each other. In the example above, we would say that the mean for males equals the mean for females.
2. In the population, the means for the columns equal each other. In the example above, we would say that the means for the three dosages are equal.
3. In the population, the null hypothesis would be that there is no interaction between the two variables. In the example above, we would say that there is no interaction between gender and amount of dosage, or that all effects equal 0.

Let's take a look at an example of a data set and see how we can interpret the summary tables produced by technological tools to test our hypotheses.

Example

Say that a gym teacher is interested in the effects of the length of an exercise program on the flexibility of male and female students. The teacher randomly selected 48 students (24 males and 24 females) and assigned them to exercise programs of varying lengths (1, 2, or 3 weeks). At the end of the programs, she measured the students' flexibility and recorded the following results. Each cell represents the score of a student:

TABLE 15.12:

		Length of Program	**Length of Program**	**Length of Program**
		1 Week	2 Weeks	3 Weeks
Gender	**Females**	32	28	36
		27	31	47
		22	24	42
		19	25	35
		28	26	46
		23	33	39
		25	27	43
		21	25	40
	Males	18	27	24
		22	31	27
		20	27	33
		25	25	25
		16	25	26
		19	32	30
		24	26	32
		31	24	29

Do gender and the length of an exercise program have an effect on the flexibility of students?

Solution

From these data, we can calculate the following summary statistics:

TABLE 15.13:

			Length of Program	**Length of Program**	**Length of Program**	
			1 Week	2 Weeks	3 Weeks	Total
Gender	**Females**	n	8	8	8	24
		Mean	24.6	27.4	41.0	31.0
		St. Dev.	4.24	3.16	4.34	8.23
	Males	n	8	8	8	24
		Mean	21.9	27.1	28.3	25.8
		St. Dev.	4.76	2.90	3.28	4.56
	Totals	n	16	16	16	48
		Mean	23.3	27.3	34.6	28.4
		St. Dev.	4.58	2.93	7.56	7.10

As we can see from the tables above, it appears that females have more flexibility than males and that the longer programs are associated with greater flexibility. Also, we can take a look at the standard deviation of each group to get an idea of the variance within groups. This information is helpful, but it is necessary to calculate the test statistic to more fully understand the effects of the independent variables and the interaction between these two variables.

TABLE 15.14:

Source	SS	df	MS	F	F critical
Gender	330.75	1	330.75	22.37	4.07
Length	1065.50	2	532.75	36.03	3.21
Gender * Length	350.00	2	175	11.84	3.21
Error	621.00	42	14.79		
Total	2376.25	47			

Note that the computer finds the degrees of freedom for the interaction by multiplying together the degrees of freedom for each variable (rows and columns).

From this summary table, we can see that all three F-ratios exceed their respective critical values.

This means that we can reject all three null hypotheses and conclude that:

In the population, the mean for males differs from the mean of females.

In the population, the means for the three exercise programs differ.

There is an interaction between the length of the exercise program and the student's gender.

Experimental Design and its Relation to the ANOVA Methods

Experimental design is the process of taking the time and the effort to organize an experiment so that the data are readily available to answer the questions that are of most interest to the researcher. When conducting an experiment using the ANOVA method, there are several ways that we can design an experiment. The design that we choose depends on the nature of the questions that we are exploring.

In a totally randomized design, the subjects or objects are assigned to treatment groups completely at random. For example, a teacher might randomly assign students into one of three reading programs to examine the effects of the different reading programs on student achievement. Often, the person conducting the experiment will use a computer to randomly assign subjects.

In a randomized block design, subjects or objects are first divided into homogeneous categories before being randomly assigned to a treatment group. For example, if an athletic director was studying the effect of various physical fitness programs on males and females, he would first categorize the randomly selected students into homogeneous categories (males and females) before randomly assigning them to one of the physical fitness programs that he was trying to study.

In ANOVA, we use both randomized design and randomized block design experiments. In one-way ANOVA, we typically use a completely randomized design. By using this design, we can assume that the observed changes are caused by changes in the independent variable. In two-way ANOVA, since we are evaluating the effect of two independent variables, we typically use a randomized block design. Since the subjects are assigned to one group and then another, we are able to evaluate the effects of both variables and the interaction between the two.

Lesson Summary

With two-way ANOVA, we are not only able to study the effect of two independent variables, but also the interaction between these variables. There are several advantages to conducting a two-way ANOVA, including efficiency, control of variables, and the ability to study the interaction between variables. Determining the total variation in two-way ANOVA includes calculating the following:

Variation within the group (within-cell variation)

Variation in the dependent variable attributed to one independent variable (variation among the row means)

Variation in the dependent variable attributed to the other independent variable (variation among the column means)

Variation between the independent variables (the interaction effect)

It is easier and more accurate to use technological tools, such as computer programs like Microsoft Excel, to calculate the figures needed to evaluate our hypotheses tests.

Review Questions

1. In two-way ANOVA, we study not only the effect of two independent variables on the dependent variable, but also the ___ between the two independent variables.
2. We could conduct multiple t-tests between pairs of hypotheses, but there are several advantages when we conduct a two-way ANOVA. These include:

   (a) Efficiency

   (b) Control over additional variables

   (c) The study of interaction between variables

   (d) All of the above

3. Calculating the total variation in two-way ANOVA includes calculating ___ types of variation.

   (a) 1

   (b) 2

   (c) 3

   (d) 4

4. A researcher is interested in determining the effects of different doses of a dietary supplement on the performance of both males and females on a physical endurance test. The three different doses of the medicine are low, medium, and high, and again, the genders are male and female. He assigns 48 people, 24 males and 24 females, to one of the three levels of the supplement dosage and gives a standardized physical endurance test. Using technological tools, he generates the following summary ANOVA table:

TABLE 15.15:

Source	**SS**	**df**	**MS**	**F**	**F critical**
Rows (gender)	14.832	1	14.832	14.94	4.07
Columns (dosage)	17.120	2	8.560	8.62	3.23
Interaction	2.588	2	1.294	1.30	3.23
Within-cell	41.685	42	992		
Total	76,226	47			

$$^{*}\alpha = 0.05$$

a. What are the three hypotheses associated with the two-way ANOVA method?
b. What are the critical values for each of the three hypotheses? What do these tell us?
c. Would you reject the null hypotheses? Why or why not?
d. In your own words, describe what these results tell us about this experiment.