BONUS CHAPTER

1

Analysis of Variance

In This Chapter

- Comparing three or more population means using analysis of variance (ANOVA)
- Using the *F*-distribution to perform a hypothesis test for ANOVA
- Using Excel to perform a one-way ANOVA test
- Comparing pairs of sample means using the Scheffé test

In Chapter 16, you learned about a hypothesis test where you could compare the means of two different populations to see whether they were different. But what if you want to compare the means of three or more populations? Well, you've come to the right place because that's what this chapter is all about.

To perform this new hypothesis test, we need to introduce one more probability distribution, known as the *F*-distribution. The test that we will perform has a very impressive name associated with it—the analysis of variance. This test is so special that it even has its own acronym: ANOVA. Sounds like something from outer space … keep reading to find out.

One-Way Analysis of Variance

If you want to compare the means for three or more populations, ANOVA is the test for you. Let's say I'm interested in determining whether there is a difference in consumer satisfaction ratings between three fast-food chains. I would collect a sample of satisfaction ratings from each chain and test to see whether there is a significant difference between the sample means. Suppose my data are as follows:

Population	Fast-Food Chain	Sample Mean Rating
1	McDoogles	7.8
2	Burger Queen	8.2
3	Windy's	8.3

My hypothesis statement would be:

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

$$H_1 : \text{not all } \mu\text{'s are equal}$$

Essentially, I'm testing to see whether the variations in customer ratings from the previous table are due to the fast-food chains or whether the variations are purely random. In other words, do customers perceive any differences in satisfaction between the three chains? If I reject the null hypothesis, however, my only conclusion is that a difference does exist. One-way analysis of variance tells me that a difference exists but does not allow me to compare population means to one another to determine where the difference is. That task requires further analysis, such as the Scheffé test.

BOB'S BASICS

To use one-way ANOVA, the following conditions must be satisfied:

- The populations of interest must be normally distributed.
- The samples must be independent of each other.
- Each population must have the same variance.
- The observations must be measured on an interval or ratio level scale.

A *factor* in ANOVA describes the cause of the variation in the data, like the independent variable. In the previous example, the factor would be the fast-food chain. This would be considered a *one-way ANOVA* because we are considering only one factor. More complex types of ANOVA can examine multiple factors, but that topic goes beyond the scope of this book.

A *level* in ANOVA describes the number of categories within the factor of interest. For our example, we have three levels based on the three different fast-food chains we're examining.

DEFINITION

A **factor** in ANOVA describes the cause of the variation in the data. When only one factor is being considered, the procedure is known as **one-way ANOVA**. A **level** in ANOVA describes the number of categories within the factor of interest.

To demonstrate one-way ANOVA, let's use the following example. Bob admits, much to Debbie's chagrin, that he is clueless when it comes to lawn care. His motto is, "If it's green, it's good." Debbie, on the other hand, knows exactly what type of fertilizer to get and when to apply it during the year. Bob hates spreading this stuff on the lawn because it makes the grass grow faster, which means he has to cut it more often.

To make matters worse, Bob has a neighbor, Bill, whose yard puts Bob's yard to shame. Mr. Perfect Lawn is out every weekend, meticulously manicuring his domain until it looks like the home field for the National Lawn Bowling Association. This gives Debbie a serious case of "lawn envy."

Anyway, there are several different types of analysis of variance, and covering them all would take a book unto itself. So throughout the remainder of this chapter, we'll use the lawn-care topic to describe two basic ANOVA procedures.

Completely Randomized ANOVA

The simplest type of ANOVA is known as *completely randomized one-way ANOVA,* which involves an independent random selection of observations for each level of one factor. Now that's a mouthful! To help explain this, let's say we're interested in comparing the effectiveness of three lawn fertilizers. Suppose we select 18 random patches of Bob's lawn and apply either Fertilizer 1, 2, or 3 to each of them. After a week, we mow the patches and weigh the grass clippings.

DEFINITION

The simplest type of ANOVA is known as **completely randomized one-way ANOVA**, which involves an independent random selection of observations for each level of one factor.

The factor in this example is fertilizer. There are three levels, representing the three types of fertilizer we are testing. The table that follows indicates the weight of the clippings in pounds from each patch. The mean and variance of each level are also shown.

Data for Lawn Clippings

	Fertilizer 1	Fertilizer 2	Fertilizer 3	Total
	10.2	11.6	8.1	
	8.5	12.0	9.0	
	8.4	9.2	10.7	
	10.5	10.3	9.1	
	9.0	9.9	10.5	
	8.1	12.5	9.5	
Total	54.7	65.5	56.9	177.1
Sample Size	6	6	6	18
Mean	9.12	10.92	9.48	$= {}^{177.1}/_{18} = 9.84$
Variance	1.01	1.70	0.96	

We'll refer to the data for each type of fertilizer as a sample. From the previous table, we have three samples, each consisting of six observations. The hypotheses statement can be stated as:

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

$$H_1 : \text{not all } \mu\text{'s are equal}$$

where μ_1, μ_2, and μ_3 are the true population means for the pounds of grass clippings for each type of fertilizer.

Partitioning the Sum of Squares

The hypothesis testing for ANOVA compares two types of variations from the samples: variations between samples and variations within each sample. The total variations in the data from our samples can be divided, or as statisticians like to say, "partitioned," into these two parts.

The first part is the variation within each sample, which is officially known as the sum of squares within (*SSW*). This measures the deviation of each observation from the sample mean and can be found using the following equation:

$$SSW = \sum_{i=1}^{k} (n_i - 1)s_i^2$$

where k = the number of samples (or levels). For the fertilizer example, $k = 3$ and:

$$s_1^2 = 1.01 \qquad s_2^2 = 1.70 \qquad s_3^2 = 0.96$$

$$n_1 = 6 \qquad n_2 = 6 \qquad n_3 = 6$$

The sum of squares within can now be calculated as:

$$SSW = (6 - 1)1.01 + (6 - 1)1.70 + (6 - 1)0.96 = 18.35$$

Some textbooks will also refer to this value as the error sum of squares (*SSE*).

The second part is the variation among the samples, which is known as the sum of squares between (*SSB*). This measures the deviation of each sample mean from the grand mean and can be found by:

$$SSB = \sum_{i=1}^{k} n_i(\overline{x}_i - \overline{\overline{x}})^2$$

where $\overline{\overline{x}}$ is the grand mean or the average value of all the observations. For the fertilizer example

$$\overline{x}_1 = 9.12 \qquad \overline{x}_2 = 10.92 \qquad \overline{x}_3 = 9.48$$

We find $\overline{\overline{x}}$, the grand mean, using:

$$\overline{\overline{x}} = \frac{\sum_{i=1}^{N} x_i}{N}$$

where N = the total number of observations from all samples.

RANDOM THOUGHTS

Some textbooks refer to this *SSB* value as the treatment sum of squares (*SSTR*).

For the fertilizer example:

$$\bar{\bar{x}} = \frac{177.1}{18} = 9.84$$

We can now calculate the sum of squares between:

$$SSB = \sum_{i=1}^{k} n_i (\bar{x}_i - \bar{\bar{x}})^2$$

$$SSB = 6(9.12 - 9.84)^2 + 6(10.92 - 9.84)^2 + 6(9.48 - 9.84)^2 = 10.86$$

Finally, the total variation of all the observations is known as the total sum of squares (SST) and can be found by:

$$SST = \sum_{i=1}^{k} \sum_{j=1}^{b} (x_{ij} - \bar{\bar{x}})^2$$

RANDOM THOUGHTS

ANOVA does *not* require that all the sample sizes are equal, as they are in the fertilizer example. See Problem 1 in the Practice Problems section as an example of unequal sample sizes.

This equation may look nasty, but it is just the difference between each observation and the grand mean squared and then totaled over all of the observations. This is clarified more in the following table.

x_{ij}	$\bar{\bar{x}}$	$(x_{ij} - \bar{\bar{x}})$	$(x_{ij} - \bar{\bar{x}})^2$
10.2	9.84	0.36	0.1296
8.5	9.84	-1.34	1.7956
8.4	9.84	-1.44	2.0736
10.5	9.84	0.66	0.4356
9.0	9.84	-0.84	0.7056
8.1	9.84	-1.74	3.0276
11.6	9.84	1.76	3.0976
12.0	9.84	2.16	4.6656
9.2	9.84	-0.64	0.4096
10.3	9.84	0.46	0.2116

x_{ij}	$\bar{\bar{x}}$	$(x_{ij} - \bar{\bar{x}})$	$(x_{ij} - \bar{\bar{x}})^2$
9.9	9.84	0.06	0.0036
12.5	9.84	2.66	7.0756
8.1	9.84	-1.74	3.0276
9.0	9.84	-0.84	0.7056
10.7	9.84	0.86	0.7396
9.1	9.84	-0.74	0.5476
10.5	9.84	0.66	0.4356
9.5	9.84	-0.34	0.1156
			Total = 29.2028

$$SST = \sum_{i=1}^{k} \sum_{j=1}^{b} (x_{ij} - \bar{\bar{x}})^2 = 29.21$$

This total sum of squares calculation can be confirmed recognizing that:

$$SST = SSW + SSB$$

$$SST = 18.35 + 10.86 = 29.21$$

TEST YOUR KNOWLEDGE

Do you need to calculate all three sum of squares: *SSB*, *SSW*, and *SST*? If you said "no," then you are correct. Since $SST = SSW + SSB$, you can just calculate *SSB* and *SSW* and add them to find SST instead of having to go the long way and calculate *SST* using its own formula, as we did above.

Note that we can determine the variance of the original 18 observations, s^2, by:

$$s^2 = \frac{SST}{N-1} = \frac{29.21}{18-1} = 1.72$$

This result can be confirmed by using the variance equation that we discussed in Chapter 5 or by using Excel.

Determining the Calculated F-Statistic

To test the hypothesis for ANOVA, we need to compare the calculated test statistic to a critical test statistic using the *F*-distribution. The calculated *F*-test statistic can be found using the equation:

$$F = \frac{MSB}{MSW}$$

where MSB is the *mean square between (MSB)*, found by:

$$MSB = \frac{SSB}{k-1}$$

and MSW is the *mean square within (MSW)*, found by:

$$MSW = \frac{SSW}{N-k}$$

BOB'S BASICS

The **mean square between (*MSB*)** is a measure of variation between the sample means. The **mean square within (*MSW*)** is a measure of variation within each sample. A large *MSB* variation, relative to the *MSW* variation, indicates that the sample means are not very close to one another. This condition will result in a large value of *F*, the calculated *F*-test statistic. The larger the value of *F*, the more likely it will exceed the critical *F*-statistic (to be determined shortly), leading us to conclude there is a difference between population means.

Now, let's put these guys to work with our fertilizer example.

$$MSB = \frac{SSB}{k-1} = \frac{10.86}{3-1} = 5.43$$

$$MSW = \frac{SSW}{N-k} = \frac{18.35}{18-3} = 1.22$$

$$F = \frac{MSB}{MSW} = \frac{5.43}{1.22} = 4.45$$

If the variation *between* the samples (*MSB*) is much greater than the variation *within* the samples (*MSW*), we will tend to reject the null hypothesis and conclude that there is a difference between population means. To complete our test for this hypothesis, we need to introduce the *F*-distribution.

Determining the Critical F-Statistic

We use the F-distribution to determine the critical F-statistic, which is compared to the calculated F-test statistic for the ANOVA hypothesis test. The critical F-statistic, $F_{\alpha, k-1, N-k}$, depends on two different degrees of freedom—degrees of freedom for the numerator (v_1) and degrees of freedom for the denominator, (v_2)—which are determined by:

$$v_1 = k - 1 \text{ and } v_2 = N - k$$

For our fertilizer example:

$$v_1 = 3 - 1 = 2 \text{ and } v_2 = 18 - 3 = 15$$

The critical F-statistic is read from the F-distribution table found in Table 6 in Appendix B of this book. Here is an excerpt of this table.

Table of Critical *F*-Statistics

$\alpha = 0.05$

v_1	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
v_2										
1	161.448	199.500	215.707	224.583	230.162	233.986	236.768	238.882	240.543	241.882
2	18.513	19.000	19.164	19.247	19.296	19.330	19.353	19.371	19.385	19.396
3	10.128	9.552	9.277	9.117	9.013	8.941	8.887	8.845	8.812	8.786
4	7.709	6.944	6.591	6.388	6.256	6.163	6.094	6.041	5.999	5.964
5	6.608	5.786	5.409	5.192	5.050	4.950	4.876	4.818	4.772	4.735
6	5.987	5.143	4.757	4.534	4.387	4.284	4.207	4.147	4.099	4.060
7	5.591	4.737	4.347	4.120	3.972	3.866	3.787	3.726	3.677	3.637
8	5.318	4.459	4.066	3.838	3.687	3.581	3.500	3.438	3.388	3.347
9	5.117	4.256	3.863	3.633	3.482	3.374	3.293	3.230	3.179	3.137
10	4.965	4.103	3.708	3.478	3.326	3.217	3.135	3.072	3.020	2.978
11	4.844	3.982	3.587	3.357	3.204	3.095	3.012	2.948	2.896	2.854
12	4.747	3.885	3.490	3.259	3.106	2.996	2.913	2.849	2.796	2.753
13	4.667	3.806	3.411	3.179	3.025	2.915	2.832	2.767	2.714	2.671

continues

Table of Critical F-Statistics (continued)

v_1	1	2	3	4	5	6	7	8	9	10
v_2										
14	4.600	3.739	3.344	3.112	2.958	2.848	2.764	2.699	2.646	2.602
15	4.543	<u>3.682</u>	3.287	3.056	2.901	2.790	2.707	2.641	2.588	2.544
16	4.494	3.634	3.239	3.007	2.852	2.741	2.657	2.591	2.538	2.494

Note that this table is based only on $\alpha = 0.05$. Other values of α will require a different table. For $v_1 = 2$ and $v_2 = 15$, the critical F-statistic, $F_{.05,2,15} = 3.682$, is indicated in the underlined part of the table. Figure 1.1 shows the results of our hypothesis test.

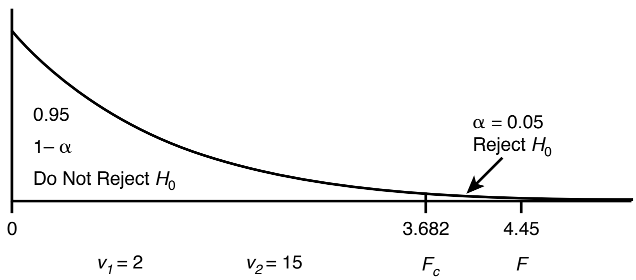

Figure 1.1
ANOVA test for the fertilizer example.

According to Figure 1.1, the calculated F-statistic of 4.45 is within the "Reject H_0" region, which leads us to the conclusion that the population means are not equal. We will always reject H_0 as long as $F_{\alpha, k-1, N-k} \leq F$.

BOB'S BASICS

The F-distribution has the following characteristics:

- It is not symmetrical, but rather has a positive skew.
- The shape of the F-distribution will change with the degrees of freedom specified by the values of v_1 and v_2.
- As v_1 and v_2 increase in size, the shape of the F-distribution becomes more symmetrical.
- The total area under the curve is equal to 1.
- The F-distribution mean is approximately equal to 1.

Our final conclusion is that one or more of those darn fertilizers is making the grass grow faster than the others.

WRONG NUMBER

Even though we have rejected H_0 and concluded that the population means are not all equal, ANOVA does not allow us to make comparisons between means. In other words, we do not have enough evidence to conclude that Fertilizer 2 produces more grass clippings than Fertilizer 1. This requires another test known as pairwise comparisons, which we'll address in the next section.

Pairwise Comparisons

Once we have rejected H_0 using ANOVA, we can determine which of the sample means are different using the Scheffé test. This test compares each pair of sample means from the ANOVA procedure. For the fertilizer example, we would compare $\bar{x}_1$ versus $\bar{x}_2$, $\bar{x}_1$ versus $\bar{x}_3$, and $\bar{x}_2$ versus $\bar{x}_3$ to see whether any differences exist.

BOB'S BASICS

After rejecting H_0 using ANOVA, we can determine which of the sample means are different using the Scheffé test.

First, the following test statistic for the Scheffé test, F_S, is calculated for each of the pairs of sample means:

$$F_s = \frac{(\bar{x}_a - \bar{x}_b)^2}{\frac{SSW}{\sum_{i=1}^{k}(n_i - 1)}\left[\frac{1}{n_a} + \frac{1}{n_b}\right]}$$

where:

$\bar{x}_a$, $\bar{x}_b$ = the sample means being compared

SSW = the sum of squares within from the ANOVA procedure

n_a, n_b = the samples sizes

k = the number of samples (or levels)

Comparing $\overline{x}_1$ and $\overline{x}_2$, we have:

$$F_s = \frac{(9.12-10.92)^2}{\frac{18.35}{5+5+5}\left[\frac{1}{6}+\frac{1}{6}\right]} = \frac{3.24}{1.22[0.33]} = 8.048$$

Comparing $\overline{x}_1$ and $\overline{x}_3$, we have:

$$F_s = \frac{(9.12-9.48)^2}{\frac{18.35}{5+5+5}\left[\frac{1}{6}+\frac{1}{6}\right]} = \frac{0.13}{1.22[0.33]} = 0.323$$

Comparing $\overline{x}_2$ and $\overline{x}_3$, we have:

$$F_s = \frac{(10.92-9.48)^2}{\frac{18.35}{5+5+5}\left[\frac{1}{6}+\frac{1}{6}\right]} = \frac{2.07}{1.22[0.33]} = 5.142$$

Next the critical value for the Scheffé test, F_{SC}, is determined by multiplying the critical F-statistic from the ANOVA test by $k - 1$ as follows:

$$F_{SC} = (k-1)F_{\alpha,\, k-1,\, N-k}$$

For the fertilizer example:

$$F_{.05,\, 2,\, 15} = 3.682$$

$$F_{SC} = (3-1)(3.682) = 7.364$$

If $F_S \le F_{SC}$, we conclude there is no difference between the sample means; otherwise there is a difference. The following table summarizes these results.

Summary of the Scheffé Test

Sample Pair	F_S	F_{SC}	**Conclusion**
$\overline{x}_1$ and $\overline{x}_2$	8.048	7.364	Difference
$\overline{x}_1$ and $\overline{x}_3$	0.323	7.364	No Difference
$\overline{x}_2$ and $\overline{x}_3$	5.142	7.364	No Difference

According to our results, the only statistically significant difference is between Fertilizer 1 and Fertilizer 2. It appears that Fertilizer 2 is more effective in making grass grow faster when compared to Fertilizer 1.

TEST YOUR KNOWLEDGE

There is another commonly used test for pairwise comparisons: the Tukey test. A main difference between the Scheffé and Tukey tests is that for the Tukey test, the sample sizes must be equal. The Scheffé test is more general and can be used with unequal sample sizes.

Completely Randomized Block ANOVA

Now let's modify the original fertilizer example: rather than selecting 18 random samples from Bob's lawn, we are going to collect 3 random samples from 6 different lawns. Using the original data, the samples look as follows:

Lawn	Fertilizer 1	Fertilizer 2	Fertilizer 3	Block Mean
1	10.2	11.6	8.1	9.97
2	8.5	12.0	9.0	9.83
3	8.4	9.2	10.7	9.43
4	10.5	10.3	9.1	9.97
5	9.0	9.9	10.5	9.80
6	8.1	12.5	9.5	10.03
Fertilizer Mean	9.12	10.92	9.48	

One concern in this scenario is that the variations in the lawns will account for some of the variation in the three fertilizers, which may interfere with our hypothesis test. We can control for this possibility by using a *completely randomized block ANOVA*, which is used in the previous table. The type of fertilizer is still the factor, and the lawns are called blocks.

DEFINITION

Completely randomized block ANOVA controls for variations from sources other than the factors of interest. This is accomplished by grouping the samples using a **blocking variable.**

There are two hypotheses for the completely randomized block ANOVA. The first (primary) hypothesis tests the equality of the population means, just like we did earlier with one-way ANOVA:

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

H_1 : not all μ's are equal

The secondary hypothesis tests the effectiveness of the *blocking variable* as follows:

H_0' : the block means are all equal

H_1' : the block means are not all equal

The blocking variable would be an effective contributor to our ANOVA model if we could reject H_0' and claim that the block means are not equal to each other.

Partitioning the Sum of Squares

For the completely randomized block ANOVA, the sum of squares total is partitioned into three parts according to the following equation:

$SST = SSW + SSB + SSBL$

where:

SSW = sum of squares within

SSB = sum of squares between

$SSBL$ = sum of squares for the blocking variable (lawns)

Fortunately for us, the calculations for SST and SSW are identical to the one-way ANOVA procedure that we've already discussed, so those values remain unchanged ($SST = 29.21$ and $SSB = 10.86$). We can find the sum of squares block ($SSBL$) by using the equation:

$$SSBL = \sum_{j=1}^{b} k(\bar{x}_j - \bar{\bar{x}})^2$$

where:

$\bar{x}_j$ = the average observation of each blocking level

b = the number of blocking levels ($b = 6$ for our example)

Using the values from the previous table, we have:

$$SSBL = 3(9.97 - 9.83)^2 + 3(9.83 - 9.83)^2 + 3(9.43 - 9.83)^2 + 3(9.97 - 9.83)^2 + 3\,(9.80 - 9.83)^2 + 3(10.03 - 9.83)^2$$

$$SSBL = 0.72$$

That leaves us with the sum of squares within (*SSW*), which we can find using:

$$SSW = SST - SSB - SSBL$$

$$SSW = 29.21 - 10.86 - 0.72 = 17.63$$

Almost done!

Determining the Calculated F-Statistic

Since we have two hypothesis tests for the completely randomized block ANOVA, we have two calculated *F*-test statistics. The *F*-test statistic to test the equality of the population means (the original hypothesis) is found using:

$$F = \frac{MSB}{MSW}$$

where *MSB* is the mean square between, found by:

$$MSB = \frac{SSB}{k-1}$$

and *MSW* is the mean square within, found by:

$$MSW = \frac{SSW}{(k-1)(b-1)}$$

Inserting our fertilizer values into these equations looks like this:

$$MSB = \frac{SSB}{k-1} = \frac{10.86}{3-1} = 5.43$$

$$MSW = \frac{SSW}{(k-1)(b-1)} = \frac{17.63}{(3-1)(6-1)} = 1.76$$

$$F = \frac{MSB}{MSW} = \frac{5.43}{1.76} = 3.09$$

The second *F*-test statistic will test the significance of the blocking variable (the second hypothesis) and will be denoted *F*'. We will determine this statistic using:

$$F' = \frac{MSBL}{MSW}$$

where *MSBL* is the (can you guess?) *mean square blocking*, found by:

$$MSBL = \frac{SSBL}{b-1}$$

Plugging our numbers into these guys results in:

$$MSBL = \frac{SSBL}{b-1} = \frac{0.72}{6-1} = 0.14$$

$$F' = \frac{MSBL}{MSW} = \frac{0.14}{1.76} = 0.08$$

We now need to sit back, catch our breath, and figure out what all these numbers mean.

To Block or Not to Block, That Is the Question

First, we will examine the primary hypothesis, H_0, that all population means are equal using $\alpha = 0.05$. The degrees of freedom for this critical F-statistic would be:

$$v_1 = k - 1 = 3 - 1 = 2$$

$$v_2 = (k - 1)(b - 1) = (3 - 1)(6 - 1) = 10$$

The critical F-statistic from Appendix B is $F_{0.05,\,2,\,10} = 4.103$. Since the calculated F-test statistic equals 3.09 and is less than this critical F-statistic, we fail to reject H_0 and cannot conclude that the fertilizer means are different.

We next examine the secondary hypothesis, H_0', concerning the effectiveness of the blocking variable, also using $\alpha = 0.05$. The degrees of freedom for this critical F-statistic would be:

$$v_1' = b - 1 = 6 - 1 = 5$$

$$v_2' = (k - 1)(b - 1) = (3 - 1)(6 - 1) = 10$$

The critical F-statistic from Appendix B is $F_{0.05,\,5,\,10} = 3.326$. Since the calculated F-test statistic, F', equals 0.08 and is less than this critical F-statistic, we fail to reject H_0' and cannot conclude that the block means are different.

What does all this mean? Since we failed to reject H_0', the hypothesis that states the blocking means are equal, the blocking variable (lawns) proved not to be effective and should not be included in the model. Including an ineffective blocking variable in the ANOVA increases the chance of a Type II error in the primary hypothesis, H_0. The conclusion of the primary hypothesis in this example would be more precise without the blocking variable. In fact, this is what essentially happened when we included the blocking variable with the randomized block design. With the blocking variable present in the model, we failed to discover a difference in the population means. Now go back to the beginning of the chapter. When we tested the population means using one-way ANOVA (without a blocking variable), we concluded that the population means were indeed different.

In summary (It's about time!), if you feel there is a variable present in your model that could contribute undesirable variation, such as taking samples from different lawns, use the randomized block ANOVA. First test H_0', the blocking hypothesis.

- If you reject H_0', the blocking procedure was effective. Proceed to test H_0, the primary hypothesis concerning the population means, and draw your conclusions.
- If you fail to reject H_0', the blocking procedure was not effective. Redo the analysis using one-way ANOVA (without blocking) and draw your conclusions.
- If all else fails, take two aspirin and call me in the morning.

Using Excel to Perform ANOVA

I'm sure you've come to the conclusion that calculating ANOVA manually is a lot of work, and I think you'll be amazed by how easy this procedure is when using Excel. Excel can do both types of ANOVA we discussed in this chapter. We will start with the example for the completely randomized ANOVA.

1. Start by placing the fertilizer data in Columns A, B, and C in a blank sheet.
2. From the toolbar menu at the top of the Excel window, click on the Data tab, and select Data Analysis. (Refer to the section "Installing the Data Analysis Add-In" from Chapter 2 if you don't see the Data Analysis command.)
3. From the Data Analysis dialog box, select Anova: Single Factor as shown in Figure 1.2, and click OK.

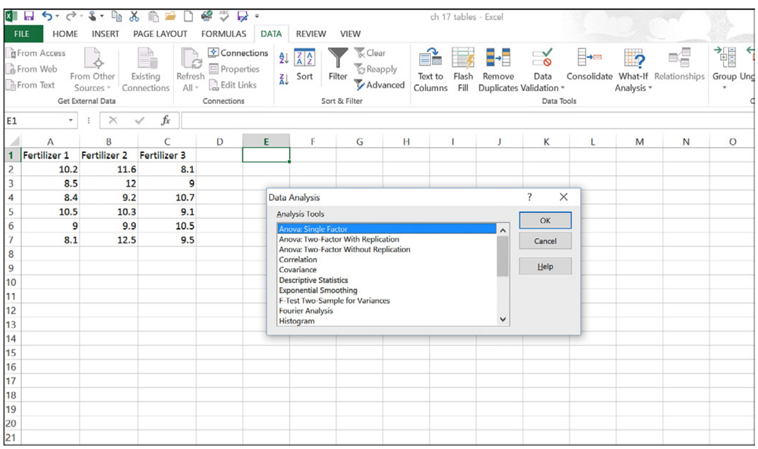

Figure 1.2
Setting up Anova: Single Factor in Excel.

4. Set up the Anova: Single Factor dialog box according to Figure 1.3.

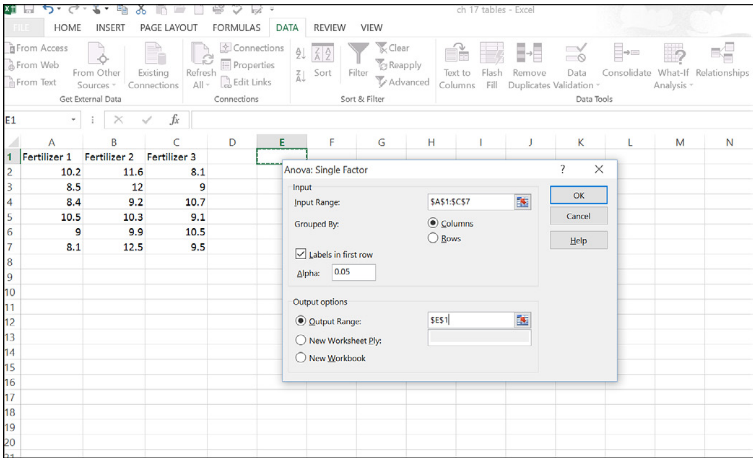

Figure 1.3

The Anova: Single Factor dialog box.

5. Click OK. Figure 1.1 shows the final ANOVA results.

Anova: Single Factor

SUMMARY

Groups	*Count*	*Sum*	*Average*	*Variance*
Fertilizer 1	6	54.7	9.116667	1.005667
Fertilizer 2	6	65.5	10.91667	1.701667
Fertilizer 3	6	56.9	9.483333	0.961667

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Between Groups	10.85778	2	5.428889	4.438993	0.030599	3.68232
Within Groups	18.345	15	1.223			
Total	29.20278	17				

Figure 1.4

Final results of Anova: Single Factor in Excel.

Under the ANOVA table, *SS* stands for the sum of squares, MS stands for the mean square, F is the calculated *F*-test statistic, and F crit is the critical value from the *F*-distribution. For example, 10.85778 is the *SSB* and 18.345 is the *SSW.* These results are consistent with what we found doing it the hard way in the previous sections. Notice that the p-value = 0.0305 for this test, meaning we can reject H_0, because this $p\text{-}value < \alpha$. Recall that we set $\alpha = 0.05$ when we stated the hypothesis test.

Now, let's do the example for the completely randomized block ANOVA.

1. Start by placing the data in Columns A, B, C, and D in a blank sheet.
2. From the toolbar menu at the top of the Excel window, click on the Data tab, and select Data Analysis.
3. From the Data Analysis dialog box, select Anova: Two-Factor Without Replication as shown in Figure 1.5 and click OK.

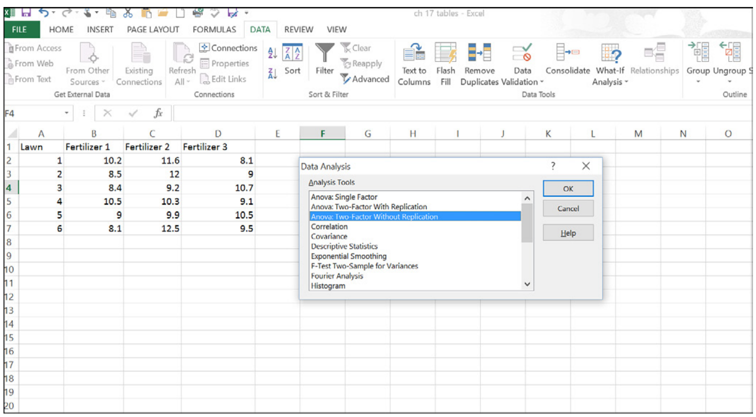

Figure 1.5
Setting up Anova: Two-Factor Without Replication in Excel.

4. Set up the Anova: Two-Factor Without Replication dialog box according to Figure 1.6.

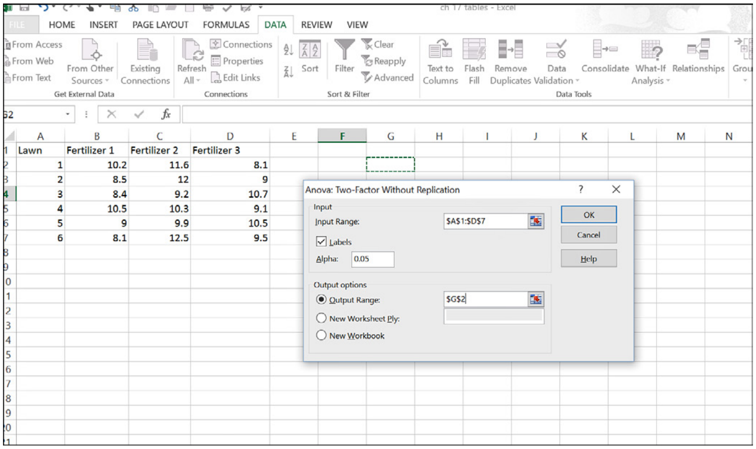

Figure 1.6
The Anova: Two-Factor Without Replication dialog box.

5. Click OK. Figure 1.7 shows the final ANOVA results.

Anova: Two-Factor Without Replication						
SUMMARY	*Count*	*Sum*	*Average*	*Variance*		
1	3	29.9	9.966667	3.103333		
2	3	29.5	9.833333	3.583333		
3	3	28.3	9.433333	1.363333		
4	3	29.9	9.966667	0.573333		
5	3	29.4	9.8	0.57		
6	3	30.1	10.03333	5.053333		
Fertilizer 1	6	54.7	9.116667	1.005667		
Fertilizer 2	6	65.5	10.91667	1.701667		
Fertilizer 3	6	56.9	9.483333	0.961667		
ANOVA						
Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Rows	0.709444	5	0.141889	0.080456	0.993829	3.325835
Columns	10.85778	2	5.428889	3.078377	0.09083	4.102821
Error	17.63556	10	1.763556			
Total	29.20278	17				

Figure 1.7
Final results of Anova: Two-Factor Without Replication in Excel.

These results are also consistent with what we found using the formulas in the previous sections. You can see how easily Excel can do the ANOVA analysis for us. What would we have done without you, Excel?

Practice Problems

1. A consumer group is testing the gas mileage of three different models of cars. Several cars of each model were driven 500 miles and the mileage was recorded as follows.

Car 1	Car 2	Car 3
22.5	18.7	17.2
20.8	19.8	18.0
22.0	20.4	21.1
23.6	18.0	19.8
21.3	21.4	18.6
22.5	19.7	

Note that the size of each sample does not have to be equal for ANOVA.

Test for a difference between the sample means using $\alpha = 0.05$.

2. Perform a pairwise comparison test on the sample means from Problem 1.

3. A vice president would like to determine whether there is a difference between the average number of customers per day among four different stores using the following data.

Store 1	Store 2	Store 3	Store 4
36	35	26	26
48	20	20	52
32	31	38	37
28	22	32	36
31	19	37	18
55	42	15	30
	29	21	

Note that the size of each sample does not have to be equal for ANOVA.

Test for a difference between sample means using $\alpha = 0.05$.

4. A certain unnamed statistics author and his two sons played golf at four different courses with the following scores:

	Dad	**Brian**	**John**
Course 1	93	85	80
Course 2	98	87	88
Course 3	89	82	84
Course 4	90	80	82

Using completely randomized block ANOVA, test for the difference of golf score means using $\alpha = 0.05$ and using the courses as the blocking variable.

The Least You Need to Know

- Analysis of variance, also known as ANOVA, compares the means of three or more populations.
- A factor in ANOVA describes the cause of the variation in the data. When only one factor is being considered, the procedure is known as a one-way ANOVA.
- A level in ANOVA describes the number of categories within the factor of interest.
- The simplest type of ANOVA is known as completely randomized one-way ANOVA, which involves an independent random selection of observations for each level of one factor.
- After rejecting H_0 using ANOVA, we can determine which of the sample means are different using the Scheffé test.
- Completely randomized block ANOVA controls for variations from sources other than the factors of interest. This is accomplished by grouping the samples using a blocking variable.