3 Vive la Différence

Understanding Variability

Difficulty Scale ☺☺☺☺ (moderately easy, but not a cinch)

What you'll learn about in this chapter

- Why variability is valuable as a descriptive tool
- How to compute the range, standard deviation, and variance
- How the standard deviation and variance are alike—and how they are different

WHY UNDERSTANDING VARIABILITY IS IMPORTANT

In Chapter 2, you learned about different types of averages, what they mean, how they are computed, and when to use them. But when it comes to descriptive statistics and describing the characteristics of a distribution, averages are only half the story. The other half is measures of variability.

In the most simple of terms, **variability** reflects how scores differ from one another. For example, the following set of scores shows some variability:

7, 6, 3, 3, 1

The following set of scores has the same mean (4) and has less variability than the previous set:

3, 4, 4, 5, 4

The next set has no variability at all—the scores do not differ from one another—but it also has the same mean as the other two sets we just showed you.

4, 4, 4, 4, 4

Variability (also called spread or dispersion) can be thought of as a measure of how different scores are from one another. It's even more accurate (and maybe even easier) to think of variability as how different scores are from one particular score. And what "score" do you think that might be? Well, instead of comparing each score to every other score in a distribution, the one score that could be used as a comparison is—that's right—the mean. So, variability becomes a measure of how much each score in a group of scores differs from the mean. More about this in a moment.

Remember what you already know about computing averages—that an average (whether it is the mean, the median, or the mode) is a representative score in a set of scores. Now, add your new knowledge about variability—that it reflects how different scores are from one another. Each is an important descriptive statistic. Together, these two (average and variability) can be used to describe the characteristics of a distribution and show how distributions differ from one another.

Three measures of variability are commonly used to reflect the degree of variability, spread, or dispersion in a group of scores. These are the range, the standard deviation, and the variance. Let's take a closer look at each one and how each one is used.

COMPUTING THE RANGE

The range is the most general measure of variability. It gives you an idea of how far apart scores are from one another. The **range** is computed simply by subtracting the lowest score in a distribution from the highest score in the distribution.

In general, the formula for the range is

$$r = h - l \tag{3.1}$$

where

r is the range
h is the highest score in the data set
l is the lowest score in the data set

Take the following set of scores, for example (shown here in descending order):

98, 86, 77, 56, 48

In this example, 98 – 48 = 50. The range is 50. In a set of 500 numbers, where the largest is 98 and the smallest is 37, then the range would be 61.

There really are two kinds of ranges. One is the *exclusive range,* which is the highest score minus the lowest score (or $h - l$) and the one we just defined. The second kind of range is the *inclusive range,* which is the highest score minus the lowest score plus 1 (or $h - l + 1$). You most commonly see the exclusive range in research articles, but the inclusive range is also used on occasion if the researcher prefers it.

The range is used almost exclusively to get a very *general* estimate of how wide or different scores are from one another—that is, the range shows how much spread there is from the lowest to the highest point in a distribution.

So, although the range is fine as a general indicator of variability, it should not be used to reach any conclusions regarding how individual scores differ from one another. And, you will usually never see it reported as the only measure of variability, but as one of several (which brings us to . . .)

COMPUTING THE STANDARD DEVIATION

Now we get to the most frequently used measure of variability, the standard deviation. Just think about what the term implies; it's a deviation from something (guess what?) that is standard. Actually, the **standard deviation** (abbreviated as *s* or *SD*) represents the average amount of variability in a set of scores. In practical terms, it's the average distance from the mean. The larger the standard deviation, the larger the average distance each data point is from the mean of the distribution.

So, what's the logic behind computing the standard deviation? Your initial thoughts may be to compute the mean of a set of scores and then subtract each individual score from the mean. Then, compute the average of that distance.

That's a good idea—you'll end up with the average distance of each score from the mean. But it won't work (see if you know why even though we'll show you why in a moment).

First, here's the formula for computing the standard deviation:

$$s = \sqrt{\frac{\Sigma(X - \overline{X})^2}{n - 1}} \tag{3.2}$$

where

s is the standard deviation

S is sigma, which tells you to find the sum of what follows

X is each individual score

$\overline{X}$ is the mean of all the scores

n is the sample size.

This formula finds the difference between each individual score and the mean ($X - \overline{X}$), squares each difference, and sums them all together. Then, it divides the sum by the size of the sample (minus 1) and takes the square root of the result. As you can see, and as we mentioned earlier, the standard deviation is an average deviation from the mean.

Here are the data we'll use in the following step-by-step explanation of how to compute the standard deviation.

5, 8, 5, 4, 6, 7, 8, 8, 3, 6

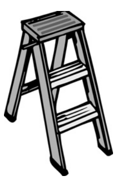

1. List each score. It doesn't matter whether the scores are in any particular order.

2. Compute the mean of the group.

3. Subtract the mean from each score.

Here's what we've done so far, where $X - \overline{X}$ represents the difference between the actual score and the mean of all the scores, which is 6.

X	$\bar{X}$	$X - \bar{X}$
8	6	8 − 6 = +2
8	6	8 − 6 = +2
8	6	8 − 6 = +2
7	6	7 − 6 = +1
6	6	6 − 6 = 0
6	6	6 − 6 = 0
5	6	5 − 6 = −1
5	6	5 − 6 = −1
4	6	4 − 6 = −2
3	6	3 − 6 = −3

4. Square each individual difference. The result is the column marked $(X - \bar{X})^2$.

X	$(X - \bar{X})$	$(X - \bar{X})^2$
8	+2	4
8	+2	4
8	+2	4
7	+1	1
6	0	0
6	0	0
5	–1	1
5	–1	1
4	–2	4
3	–3	9
Sum	0	28

5. Sum all the squared deviations about the mean. As you can see, the total is 28.

6. Divide the sum by n −1, or 10 − 1 = 9, so then 28/9 = 3.11.

7. Compute the square root of 3.11, which is 1.76 (after rounding). That is the standard deviation for this set of 10 scores.

What we now know from these results is that each score in this distribution differs from the mean by an average of 1.76 points.

Let's take a short step back and examine some of the operations in the standard deviation formula. They're important to review and will increase your understanding of what the standard deviation is.

First, why didn't we just add up the deviations from the mean? Because the sum of the deviations from the mean is always equal to 0. Try it by summing the deviations (2 + 2 + 2 + 1 + 0 + 0 − 1 − 1 − 2 − 3). In fact, that's the best way to check if you computed the mean correctly.

There's another type of deviation that you may read about, and you should know what it means. The **mean deviation** (also called the mean absolute deviation) is the sum of the absolute value of the deviations from the mean. You already know that the sum of the deviations from the mean must equal 0 (otherwise the mean is probably computed incorrectly). Instead, let's take the absolute value of each deviation (which is the value regardless of the sign). Sum them together and divide by the number of data points, and you have the mean deviation. So, if you have a set of scores such as 3, 4, 5, 5, 8, and the arithmetic mean is 5, the mean deviation is 2 (the absolute value of 5 − 3), 1, 0, 0, and 3, for a total of 6. (Note: The absolute value of a number is usually represented as that number with a vertical line on each side of it, such as $|5|$. For example, the absolute value of −6, or $|-6|$, is 6.)

Second, why do we square the deviations? Because we want to get rid of the negative sign so that when we do eventually sum them, they don't add up to 0.

And finally, why do we eventually end up taking the square root of the entire value in Step 7? Because we want to return to the same units with which we originally started. We squared the deviations from the mean in Step 4 (to get rid of negative values) and then took the square root of their total in Step 7. Pretty tidy.

Why n – 1? What's Wrong With Just n?

You might have guessed why we square the deviations about the mean and why we go back and take the square root of their sum. But how about subtracting the value of 1 from the denominator of the formula? Why do we divide by $n - 1$ rather than just plain ol' n? Good question.

The answer is that s (the standard deviation) is an estimate of the population standard deviation, and is an **unbiased estimate** at that, but only when we subtract 1 from n. By subtracting 1 from the denominator, we artificially force the standard deviation to be larger than it would be otherwise. Why would we want to do that? Because, as good scientists, we are conservative. Being conservative means that if we have to err, we will do so on the side of overestimating what the standard deviation of the population is. Dividing by a smaller denominator lets us do so. Thus, instead of dividing by 10, we divide by 9. Or instead of dividing by 100, we divide by 99.

TECH TALK Biased estimates are appropriate if your intent is only to describe the characteristics of the sample. But if you intend to use the sample as an estimate of a population parameter, then the unbiased statistic is best to calculate.

Take a look in the following table and see what happens as the size of the sample gets larger (and moves closer to the population in size). The $n - 1$ adjustment has far less of an impact on the difference between the biased and the unbiased estimates of the standard deviation (the bold column in the table). All other things being equal, then, the larger the size of the sample, the less of a difference there is between the biased and the unbiased estimates of the standard deviation. Check out the following table, and you'll see what we mean.

Sample Size	*Value of Numerator in Standard Deviation Formula*	*Biased Estimate of the Population Standard Deviation (dividing by n)*	*Unbiased Estimate of the Population Standard Deviation (dividing by n – 1)*	*Difference Between Biased and Unbiased Estimates*
10	500	7.07	7.45	.38
100	500	2.24	2.25	.01
1,000	500	0.7071	0.7075	.0004

The moral of the story? When you compute the standard deviation for a sample, which is an estimate of the population, the closer to the size of the population the sample is, the more accurate the estimate will be.

What's the Big Deal?

The computation of the standard deviation is very straightforward. But what does it mean? As a measure of variability, all it tells us is how much each score in a set of scores, on the average, varies from the mean. But it has some very practical applications, as you will find out in Chapter 4. Just to whet your appetite, consider this: The standard deviation can be used to help us compare scores from different distributions, *even when the means and standard deviations are different.* Amazing! This, as you will see, can be very cool.

THINGS TO REMEMBER

- The standard deviation is computed as the average distance from the mean. So, you will need to first compute the mean as a measure of central tendency. Don't fool around with the median or the mode in trying to compute the standard deviation.
- The larger the standard deviation, the more spread out the values are, and the more different they are from one another.
- Just like the mean, the standard deviation is sensitive to extreme scores. When you are computing the standard deviation of a sample and you have extreme scores, note that somewhere in your written report.
- If $s = 0$, there is absolutely no variability in the set of scores, and the scores are essentially identical in value. This will rarely happen.

COMPUTING THE VARIANCE

Here comes another measure of variability and a nice surprise. If you know the standard deviation of a set of scores and you can square a number, you can easily compute the variance of that same set of scores. This third measure of variability, the **variance**, is simply the standard deviation squared.

In other words, it's the same formula you saw earlier but without the square root bracket, like the one shown in Formula 3.3:

$$s^2 = \frac{\Sigma(X - \bar{X})^2}{n - 1} \tag{3.3}$$

If you take the standard deviation and never complete the last step (taking the square root), you have the variance. In other words, $s^2 = s \times s$, or the variance equals the standard deviation times itself (or squared). In our earlier example, where the standard deviation was equal to 1.76, the variance is equal to 1.76^2 or 3.11. As another example, let's say that the standard deviation of a set of 150 scores is 2.34. Then, the variance would be 2.34^2 or 5.48.

You are not likely to see the variance mentioned by itself in a journal article or see it used as a descriptive statistic. This is because the variance is a difficult number to interpret and apply to a set of data. After all, it is based on squared deviation scores.

But the variance is important because it is used both as a concept and as a practical measure of variability in many statistical formulas and techniques. You will learn about these later in *Statistics for People Who (Think They) Hate Statistics.*

The Standard Deviation Versus the Variance

How are standard deviation and the variance the same, and how are they different?

Well, they are both measures of variability, dispersion, or spread. The formulas used to compute them are very similar. You see them all over the place in the "Results" sections of journals.

They are also quite different.

First, and most important, the standard deviation (because we take the square root of the average summed squared deviation) is stated in the original units from which it was derived. The variance is stated in units that are squared (the square root is never taken).

What does this mean? Let's say that we need to know the variability of a group of production workers assembling circuit boards. Let's say that they average 8.6 boards per hour, and the standard deviation is 1.59. The value 1.59 means that the difference in the average number of boards assembled per hour is about 1.59 circuit boards from the mean.

Let's look at an interpretation of the variance, which is 1.59^2, or 2.53. This would be interpreted as meaning that the average difference between the workers is about 2.53 circuit boards *squared* from the mean. Which of these two makes more sense?

USING THE COMPUTER TO COMPUTE MEASURES OF VARIABILITY

Let's use SPSS to compute some measures of variability. We are using the file named Chapter 3 Data Set 1.

There is one variable in this data set:

Variable	*Definition*
ReactionTime	Reaction time on a tapping task

Here are the steps to compute the measures of variability that we discussed in this chapter.

1. Open the file named Chapter 3 Data Set 1.
2. Click Analyze → Descriptive Statistics → Frequencies.
3. Double-click on the ReactionTime variable to move it to the Variable(s) box.
4. Click Statistics, and you will see the Frequencies: Statistics dialog box. Use this dialog box to select the variables and procedures you want to perform.
5. Under Dispersion, click Std. Deviation.
6. Under Dispersion, click Variance.
7. Under Dispersion, click Range.
8. Click Continue.
9. Click OK.

The SPSS Output

Figure 3.1 shows selected output from the SPSS procedure for ReactionTime. There are 30 valid cases with no missing cases, and the standard deviation is .70255. The variance equals .494 (or s^2), and the range is 2.60.

Statistics

Reaction Time

N	Valid	30
	Missing	0
Std. Deviation		.70255
Variance		.494
Range		2.60

Figure 3.1 Output for the Variable ReactionTime

Let's try another one, titled Chapter 3 Data Set 2. There are two variables in this data set:

Variable	*Definition*
MathScore	Score on a mathematics test
ReadingScore	Score on a reading test

Follow the same set of instructions as given previously, only in Step 3, you select both variables. The SPSS output is shown in Figure 3.2, where you can see selected output from the SPSS procedure for these two variables. There are 30 valid cases with no missing cases, and the standard deviation for math scores is 12.36 with a variance of 152.7 and a range of 43. For reading scores, the standard deviation is 18.700, the variance is a whopping 349.689 (that's pretty big), and the range is 76 (which is large as well, reflecting the similarly large variance).

Statistics

		Math_Score	Reading_Score
N	Valid	30	30
	Missing	0	0
Std. Deviation		12.357	18.700
Variance		152.700	349.689
Range		43	76

Figure 3.2 Output for the Variables MathScore and ReadingScore

Summary

Measures of variability help us even more fully understand what a distribution of data points looks like. Along with a measure of central tendency, we can use these values to distinguish distributions from one another and effectively describe what a collection of test scores, heights, or measures of personality looks like. Now that we can think and talk about distributions, let's explore ways we can look at them.

Time to Practice

1. Why is the range the most convenient measure of dispersion, yet the most imprecise measure of variability? When would you use the range?

2. Compute the exclusive and inclusive ranges for the following items.

High Score	*Low Score*	*Inclusive Range*	*Exclusive Range*
7	6		
89	45		
34	17		
15	2		
1	1		

3. Why would you expect more variability on a measure of personality in college freshman graders than you would on a measure of height?
4. Why does the standard deviation get smaller as the individuals in a group score more similarly on a test?
5. For the following set of scores, compute the range, the unbiased and the biased standard deviations, and the variance. Do the exercise by hand.

31, 42, 35, 55, 54, 34, 25, 44, 35

Why is the unbiased estimate greater than the biased estimate?

6. Use SPSS to compute all the descriptive statistics for the following set of three test scores over the course of a semester. Which test had the highest average score? Which test had the lowest amount of variability?

Test 1	*Test 2*	*Test 3*
50	50	49
48	49	47
51	51	51
46	46	55
49	48	55
48	53	45
49	49	47
49	52	45
50	48	46
50	55	53

7. For the following set of scores, compute by hand the unbiased estimates of the standard deviation and variance.

4, 5, 6, 2, 5, 7, 5, 6, 8, 5

8. The variance for a set of scores is 25. What is the standard deviation and what is the range?

9. This practice problem uses the data contained in the file named Chapter 3 Data Set 3. There are two variables in this data set.

Variable	*Definition*
Height	height in inches
Weight	weight in pounds

Using SPSS, compute all of the measures of variability you can for height and weight.

10. How can you tell whether SPSS produces a biased or an unbiased estimate of the standard deviation?.