

3 *Central Tendency*

Numbers usually come in sets, also referred to as *samples*. A set of numbers by itself makes little or no sense if the numbers are not organized in some coherent manner. To understand the numbers and determine their meaning, they must be arranged in certain ways. In the previous chapter, methods for determining the shape, or distribution, of a set of numbers were presented. From the shape, the peak in the distribution and whether the distribution is symmetric can be determined.

This chapter discusses the typical or most representative value of a set of numbers. Of interest is the value around which the observations cluster. This value, the *central tendency* of a sample, estimates the typical value of an observation from the sample. A measure of central tendency represents all the numbers in the sample.

Central tendencies are a very common part of modern life. In sports, we hear about the average number of points a basketball player scores per game. In economics, we hear about the average cost of buying a home. In health, we hear of the average age at which children get a particular disease. Central tendencies are so common that they are not usually viewed as statistics.

There are two major reasons for knowing the central tendency: simplification and prediction. *Simplification* is needed because the whole sample of numbers often contains just too much information. Imagine that you are keeping a record of your monthly expenses for gasoline. Instead of remembering the dollar figures for each month over the past four years (48 numbers), it might be much more convenient to know only the average (one number). So for reasons of economy, it is often much more useful to record the mean, or the average of the numbers, instead of all the numbers. For the United States Census, it would be unthinkable not to compute a measure of central tendency. Imagine trying to report the incomes of 125 million households. There are too many numbers to keep track of. These numbers need to be boiled down to one measure of central tendency.

The second reason for computing the typical value is *prediction*. The knowledge of the average winter temperature in New York City for the past

ten years can assist in predicting how cold it will be there this year. This knowledge will be useful in the determination of how much energy will be needed to keep the house warm. Thus, by knowing the average temperature, one can make a prediction and make choices that are consistent with that prediction. An average over ten years is probably a better predictor of next year's temperature than that of any one year. Let us consider a second example. Say a city is faced with a steady stream of immigrants from a foreign country. If the city is to plan for its future, it will need to know the age of these immigrants. Thus, by learning the average age of the immigrants, the city can predict future demand for schooling.

There are many ways of determining the typical value of a sample. First, consider one seemingly reasonable procedure. If it were possible to determine what person or observation is typical, that observation and that value would provide a measure of central tendency. This strategy is often employed by journalists. To predict an election, the journalist travels to a typical town and interviews the typical person in that town. There is one major advantage of using the response of the typical person as the typical response: It seems so sensible. The response of the typical person seems more valid than some statistical amalgamation of numbers. However, a more careful examination reveals two major drawbacks to the "typical person" strategy. First, in order to measure the typical value, where the typical town is and who in it is a typical person must be determined. Thus, there is a definitional problem. How can a typical value be determined if first it must be determined who is typical? The "typical person" definition of central tendency is circular, so it cannot be used. Second, the definition is not a very efficient way of determining the central tendency. To see this, it is likely that the typical value would be quite different if a different "typical" person was chosen.

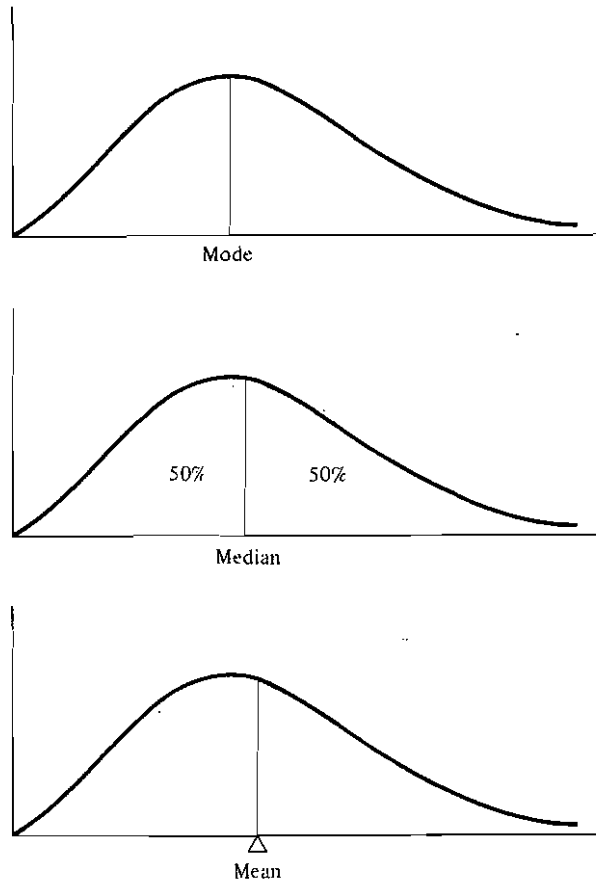
The strategy of locating the typical person must be abandoned and a new strategy must be developed to determine the central tendency. One way is to determine what value is most frequent in the sample. A second way is to determine what value is in the middle of the distribution. Finally, a third way is to determine what value is the closest to all the others. These three definitions of typical values—most frequent, in the middle, and closest—correspond to the three standard ways of determining the typical value: the mode, the median, and the mean.

Measures of Central Tendency

Figure 3.1 illustrates a hypothetical distribution that is somewhat positively skewed. The *mode* is the value that occurs most often. As shown in Figure 3.1, the mode is the highest point or peak in the distribution.

The *median* is the value in the distribution above which 50% of the scores

FIGURE 3.1 Illustration of mode, median, and mean.



lie and below which 50% of the scores lie. The median divides the distribution in half in terms of frequency.

The *mean* is the arithmetic average of the set of numbers. It is the balance point in the distribution. If the distribution were a stack of toothpicks lying on a board, the mean would be the point of balance.

The mode requires only nominal data. It simply notes the most frequent occurrence and so the objects need only be differentiated. The median requires that the numbers be rank ordered and so only ordinal data are required. Finally, the mean requires interval data because the scores must be summed to calculate it. Note that the mode and median can be meaningfully calculated for interval data, but the mean should be used only with interval data.

Computation of Central Tendency

Now that three measures of central tendency have been defined, their computational methods are presented.

The mode is easy to determine. *Count the number of times each observation occurs. The observation that occurs most frequently is the mode.* There may be a tie for the mode and in that case there would be two modes. If the distribution has two peaks of unequal height, it may be useful to report both peaks.

For some variables, measurements are fine-grained, with the consequence that no two scores have exactly the same value. For instance, for a sample of 100 persons whose weight is measured in grams, it is very unlikely that two or more persons would have exactly the same weight. In such cases, one may create a frequency table, smooth the frequencies, and report the mode of the smoothed distribution (smoothing was described in Chapter 2).

The median is determined by the following procedure.¹ *Rank order the scores; the median is the value of the score that falls in the middle.* The middle observation is determined as follows: The numbers are rank ordered from the smallest to the largest, just as was done in the previous chapter in order to make a frequency table. If n , the sample size, is odd then the middle observation is the $(n + 1)/2$ th largest observation. So if there are eleven observations the median would be the sixth largest observation because there are five larger and five smaller scores. If n is even, the median is defined as the average of $(n/2)$ th and $(n/2 + 1)$ th observations. In words, if the sample size is even, the median is one-half the sum of the two middle scores. So if n is ten the median would be the average of the fifth, or $n/2$ th, score and the sixth, or $(n/2 + 1)$ th, score.

The mean is usually denoted in statistical work by the symbol \bar{X} which is read as "X-bar." (Less frequently M is used to denote the mean.) *The mean or \bar{X} is defined as the sum of the observations divided by the number of observations.* So, one simply adds up all the scores and divides by the total number of scores. The mean is the arithmetic average of the sample. In terms of a formula,

$$\bar{X} = \frac{\sum X}{n}$$

where $\sum X$ is the sum of the numbers in the sample and n is the sample size. The mean is the most common measure of central tendency. It is as commonly used in statistical work as it is in everyday life.

¹A more complicated formula is presented in some other texts. The formula presented here presumes that the numbers have not been rounded. If the scores have been rounded, a different formula must be used. The result using the more complicated formula differs only fractionally from the one presented in the text.

There are some procedures that can reduce the likelihood of computational errors in adding the scores. It may help to separate the positive and negative numbers and sum the two groups of numbers individually. Also if the numbers are very large, it may help to subtract a common number from all the scores and add that number back into the mean.

To illustrate the computation of measures of central tendency, three examples are provided. Given the following sample,

1, 1, 1, 1, 2, 2, 2, 3, 3

the mode is 1 because it is observed 4 times. The median is the fifth [(9 + 1)/2th] largest observation, which is 2. The mean, rounded to two decimal places, is

$$\frac{1 + 1 + 1 + 1 + 2 + 2 + 2 + 3 + 3}{9} = 1.78$$

As another example consider Smith's study of bias of psychotherapists on the basis of the client's sex, which was presented in Chapter 2. A negative score indicates that therapists have a profemale bias, and a positive score indicates that they have a promale bias. The rank-ordered data are presented again here:

-1.03	-.23	.00	.14
-.56	-.22	.00	.23
-.40	-.10	.00	.24
-.36	-.03	.01	.29
-.31	.00	.01	.35
-.31	.00	.02	.56
-.23	.00	.05	.56
-.23	.00	.11	.60

The mode is .00, which indicates that therapists are most often neutral. The median is also .00. The mean is equal to the sum of the numbers divided by the number of studies. The sum of the 32 numbers equals -.84, and so the mean or \bar{X} equals -.84 divided by 32, which equals -.02625. The mode, the median, and the mean virtually agree for this example and they all indicate that therapists show little or no bias on the average.

In 1960 Stanley Milgram conducted an experiment on obedience to authority. Residents of New Haven, Connecticut, were asked to shock someone who failed to learn material. Actually, the learner was a paid employee of Milgram and was not actually shocked. The "shocks" started at a low level but gradually escalated to a very high voltage. The person who was shocked begged the subject to stop shocking him and complained of a heart condition. The largest possible shock that could be administered, 450 volts, was labeled "danger—severe shock."

Psychiatrists had predicted that subjects would not shock the learner beyond 300 volts. What Milgram actually found is contained in Table 3.1.

TABLE 3.1 Maximum Shock Level Administered by 40 Subjects in the Milgram Obedience Experiment

300	330	450	450
300	345	450	450
300	360	450	450
300	375	450	450
300	450	450	450
315	450	450	450
315	450	450	450
315	450	450	450
315	450	450	450
330	450	450	450

Quite clearly subjects were very obedient. Of 40 subjects, 26 administered the maximum voltage.

The mode is 450, the maximum value possible. Because $n = 40$, the median is the average of the 20th and 21st observations. Both the 20th and 21st observations are 450, and so the median is also 450. The sum of the 40 observations is 16,200, and so the mean is $16,200/40 = 405$. Alternatively, the computations for the mean could be simplified by subtracting 300 from each number. The sum of the numbers would then be 4200. The mean would be $(4200/40) + 300$, which also equals 405, as it should. As will be seen later in this chapter, with distributions that are negatively skewed such as this one, the mean tends to be less than the median and the mode.

Three measures of central tendency have been defined: the mode, the median, and the mean. It should be noted that they do not always point to the same typical value. More often than not they disagree, if only by a small amount. The source of the disagreement is due to the shape of the distribution. Distributions with certain shapes yield a mode, median, and mean that differ from each other. When the distribution is exactly symmetric, the median and the mean are equal to each other. The earlier discussed gender bias data set has a fairly symmetric distribution, and the mean and the median are nearly equal to each other. Symmetry in the distribution is a sufficient but not necessary condition for the mean to equal the median. That is, in a symmetric distribution the mean must equal the median, but in a distribution that is *not* symmetric, the mean and median may also be equal. For instance, the following sample is not symmetric

1, 1, 2, 6, 7, 8, 17

but the mean and median are both equal to 6.

For symmetric distributions, the mode equals the median and the mean when the distribution is unimodal. A unimodal, symmetric distribution must be peaked in the center of the distribution. For bimodal, symmetric distribu-

tions, the peaks are not in the center and the mode does not equal the mean or the median.

For skewed distributions (see Chapter 2 for a definition of skew) the following generally holds ($<$ symbolizes less than)


Positive skew: mode $<$ median $<$ mean

Negative skew: mean $<$ median $<$ mode

which is shown graphically in Figure 3.2. So, for a positive skew, the mean is usually the largest measure of central tendency, whereas for a negative skew the mean is usually the smallest.

Because the value of central tendency depends on what measure of the central tendency is used, which one should be preferred? There is no universal answer to this question. The answer depends on what the *purpose* is in determining the central tendency. If the purpose is ease of computation, the mode is probably the best measure. However, for large data sets (sample size greater than 100), the mean is probably easier to determine than either the mode or the median. If the purpose is prediction, and given a symmetric distribution with observations bunched in the center, the mean tends to be a better predictor of future values than the median or the mode. For ease of interpretation the median is useful. For multimodal distributions reporting the mode is best. So, the determination of which measure of central tendency is best depends on the shape of the distribution, ease of computation, and simplicity of interpretation.

FIGURE 3.2 Direction of skew and relationship between the mode, median, and mean.



One guideline to determining which measure of central tendency to use is the level of measurement. For variables at the nominal level of measurement, the mode is the appropriate measure. For variables at the ordinal level of measurement, the median is the appropriate measure. Finally for variables at the interval level of measurement, the mean is the most appropriate measure of central tendency. But even if the level of measurement is at the interval level of measurement, the shape of the distribution may require using the median or the mode as the measure of central tendency.

Properties of the Mean

The mean, or the arithmetic average, is the standard measure of central tendency. Given a sample of numbers, most researchers almost automatically compute the mean. Because it is an important and common statistic, special attention must be paid to its interpretation.

The mean is not necessarily equal to any observation in the sample. This point is abundantly clear for the variable of family size. Imagine that in a given area, the average size of a family is 4.2 persons. Of course, no family has 4.2 members, but 4.2 is a useful number nonetheless. Say, a housing development of 50 units is planned. Using the mean as a guide for forecasting, the expectation is for 50 times 4.2 or 210 to live in the development. The number 210 can be very useful in planning the need for social services in the community.

The mean can be quite misleading if there is an outlier in the sample. Say, the variable is income and the annual income of six college students is measured and the following numbers are obtained:

\$2,700
\$3,600
\$2,800
\$6,300
\$1,800
\$1,040,000

The millionaire student's income influences the mean so extremely as to result in a mean of \$176,200, not at all typical of the income of the other five college students. Means can be grossly distorted by the presence of outliers. The median is much less affected. The median of the six incomes is \$3,200.

When the sample contains a mixture of scores from two very different types of persons, the distribution is bimodal. In such cases the mean is not very informative.

The mean is the only measure of central tendency such that it can be subtracted from each observation and the sum of these differences is always zero. As a formula, $\sum(X - \bar{X})$ equals zero. One consequence of this fact is that

the sum of observations above the mean is as far above the mean as the sum of the observations below the mean. It is this property that makes the mean the number that is closest to the other numbers in the sample.

A related property of the mean is that if a constant is subtracted from each score, and this difference is squared and summed across all the scores in the sample, this quantity is at its smallest value when the constant chosen is the mean. The mean is said to be a *least-squares* estimate of central tendency.

Grouped Data

Sometimes a researcher has a frequency table but does not have access to original scores. The computation of the mode, median, and mean must be modified.

The mode of data in a frequency table is the midpoint of the class with the largest frequency. (As defined in Chapter 2, the midpoint is one-half the sum of the upper and lower limits for a class.) The median can be obtained by adding the relative frequencies starting with the lowest class. The median is the midpoint of the class interval whose cumulative relative frequency contains 50%.

When observations are grouped together into classes the mean can be computed as follows:

1. Multiply each score by its frequency.
2. Sum the score frequency products.
3. Divide this sum by the sample size.

For instance, in the Milgram study of obedience the shock level is first multiplied by the frequency.

<i>Score</i>	<i>Frequency</i>			
300	×	5	=	1,500
315	×	4	=	1,260
330	×	2	=	660
345	×	1	=	345
360	×	1	=	360
375	×	1	=	375
450	×	26	=	11,700

The sum of these products equals 16,200. The mean or \bar{X} then equals 16,200 divided by 40, which is 405.

When numbers are grouped by class intervals and there is no access to the original data, one can use the above technique to approximate the mean. The midpoint of the class interval is used as the number to be multiplied by the frequency. For instance, in Chapter 2 the following class intervals were set up for Smith's review of 32 gender bias studies of therapists and counselors. Below are the class intervals and frequencies used in Chapter 2.

<i>Class Interval</i>	<i>Frequency</i>
-1.10 to -.91	1
-.90 to -.71	0
-.70 to -.51	1
-.50 to -.31	4
-.30 to -.11	4
-.10 to .09	13
.10 to .29	5
.30 to .49	1
.50 to .69	3

The midpoint of the class intervals times the frequencies are

<i>Midpoint</i>	<i>Frequency</i>	
-1.005	× 1	= -1.005
-.805	× 0	= 0.000
-.605	× 1	= -.605
-.405	× 4	= -1.620
-.205	× 4	= -.820
-.005	× 13	= -.065
.195	× 5	= .975
.395	× 1	= .395
.595	× 3	= 1.785

The sum of these numbers is $-.960$. The mean is then estimated by $-.960$ divided by 32, which equals $-.03$. This value closely approximates the actual mean of $-.02625$.

The mode for the gender bias data is $-.005$ because the interval $-.10$ to $.09$ has the largest frequency. The median is also $-.005$. Both of these values are close to the mode and median of the original data, each of which is $.00$.

Occasionally, there is a choice of what factor to use as the frequency when computing a measure of central tendency. What to use as the frequency is determined by what variable the researcher seeks to measure. Imagine a car manufacturer which produces five different kinds of cars. These cars have the following mean miles-per-gallon ratings.

18, 24, 27, 30, 45

What is the average miles per gallon for the car company's fleet? It all depends on what exactly the question is. If the mileage of the five cars is desired, the average the five numbers would suffice. However, if the mileage of cars sold is desired, the number of cars sold must be used as the frequency. Alternatively, if the interest is in the number of gallons of gas, the frequency to be used is the number of miles driven.

Summary

A set of scores can be summarized by a measure of central tendency. The *central tendency* estimates a sample's typical value. Measures of central

tendency are used to simplify a mass of data and to facilitate prediction. There are three common measures of central tendency: the mode, the median, and the mean. The *mode* is the most frequent number in the sample, the *median* is the middle of the distribution, and the *mean* is the arithmetic average.

The mode is directly determined by the most frequent score in the sample. The median is determined by first rank ordering the scores from smallest to largest. Given n scores, for odd-sized samples, the median is the $(n + 1)/2$ th largest score. For even-sized samples, the median is the average of the $n/2$ th and the $(n/2 + 1)$ th largest scores. The mean, which is symbolized by \bar{X} , is the sum of all the scores divided by the sample size.

For symmetric distributions the mean and the median are equal. For positively skewed distributions the mean tends to be greater than the median, whereas for negatively skewed distributions the mean tends to be less than the median. The mean is very sensitive to outliers, but the mode and median are less affected.

For data grouped into classes the mode, median, and mean of the original data can be closely approximated. The mode is the midpoint of the class interval that has the largest frequency. The median is the midpoint of the class interval that contains the cumulative relative frequency of 50. The mean is found by multiplying each class midpoint by its class frequency, summing the products, and then dividing by the sample size.

The measures of central tendency provide one way of summarizing a distribution. The next chapter presents various ways of determining how meaningful the central tendency is as a summary of the distribution. This can be done by computing a measure of variability.

Problems

1. Compute the mode, median, and mean for the following samples.
 - a. 6, 8, 3, 5, 6, 2, 7
 - b. 4, 3, 2, 6, 2, 2, 5, 4
 - c. 2, 8, 4, 2, 5, 8, 10, 1, 8, 2
 - d. 2, 4, 3, 7, 5, 8, 3, 96
2. In which of the four samples in problem 1 is the mean less informative than the other measures of central tendency? Why?
3. a. Compute the mean, mode, and median of the following sample of numbers.

2, 3, 3, 3, 3, 4, 4, 5, 6, 6
- b. Compute the mean, mode, and median for the same sample of numbers, but including the number 100.

2, 3, 3, 3, 3, 4, 4, 5, 6, 6, 100

Compare the results with the answers for part (a). Which measure of central tendency changed the most?

- c. Again take the sample of numbers in part (a), include the number 10, and compute the mean, median, and mode.

2, 3, 3, 3, 3, 4, 4, 5, 6, 6, 10

Compare the results with the answers for part (a) and part (b). Which measure of central tendency was most affected by the size of the additional number?

4. A class, consisting of an even number of students, takes an exam, and 14 students score above the median. How many students are in the class?
5. Below is a table of the area of the New England states in square miles.

<i>State</i>	<i>Area</i>
Connecticut	5,009
Maine	33,215
Massachusetts	8,257
New Hampshire	9,304
Rhode Island	1,214
Vermont	9,609

- a. Compute the mean and the median.
- b. Which measure gives a better description of the area of a typical New England state?
6. Given below are test scores for a group of 20 subjects.

0, 3, 11, 24, 36, 47, 42, 53, 56, 59,
52, 58, 50, 64, 63, 61, 65, 78, 89, 91

- a. Compute the mean and median. Which estimate, mean or median, best describes the central tendency of the data? Why?
- b. Unfortunately, the values 0, 3, and 11 were incorrectly recorded. These three values should be 70, 73, and 81, respectively. Does this additional information change your answer to part (a)?
7. A survey of the ages of residents of nursing homes yielded the following measures of central tendency.

mean = 70 median = 78 mode = 83

In which direction is the distribution likely to be skewed?

8. Suppose that students in an introductory psychology class were tested on their knowledge of foreign policy. Suppose further that the following measures on central tendency were obtained.

mean = 80 median = 71 mode = 65

In which direction is the distribution likely to be skewed?

9. Draw histograms for samples with the following characteristics.
- The mean is greater than the mode.
 - The mean is less than the median.
 - The median is less than the mode.
 - There are two modes.

10. Schwartz and Leonard (1984) studied the learning of words that refer to objects and words that refer to actions. Their subjects were children under 1.5 years old. The children repeatedly heard 16 nonsense syllables applied to unfamiliar objects or actions. The following shows how many object and action words were acquired by each child.

Child.	Object Words	Action Words
JR	6	6
KP	8	3
LT	7	4
BG	7	7
KS	7	5
RC	3	2
CB	7	2
MK	8	2
LF	7	3
KB	7	2
BP	7	4
TH	8	2

- Compute the mean number of object words the children acquired.
 - Compute the mean number of action words.
11. From the data in problem 10, for each child subtract the action word *mean* from each child's action word score.
- How many children acquired more than the mean number of action words?
 - How many acquired less than the mean number?
 - Compute $\Sigma(X - \bar{X})$.
 - If you did the same computations for the object words, what would you find for $\Sigma(X - \bar{X})$? Why?
12. a. From the data in problem 10, find the median number of action words acquired. Subtract the median from each child's action word score, square the result, and sum across children.
- For problem 11, you computed the deviation of each score from the mean. Square these deviations, and sum across children.
 - Which is smaller, the sum of the squared deviations from the mean, or the sum of the squared deviations from the median?
13. Ballard and Crooks (1984) tried to increase the social involvement of preschool children who had low levels of interaction with others. They

observed the children after showing them a videotape in which a child modeled the activity of joining others in play. The subjects were scored on the degree of social involvement in their play and number of social interactions. Below are the mean scores per session each subject obtained after seeing the videotape.

<i>Subject</i>	<i>Interactions</i>	<i>Play Score</i>
S1	8.17	3.29
S2	8.50	2.69
S3	8.38	3.25
S4	7.92	2.73
S5	14.43	2.45
S6	14.36	3.33

- a. Compute the mean number of interactions observed after the videotape, across all subjects.
 - b. Compute the mean play score obtained after the videotape, across all subjects.
14. A university bookstore stocks a range of calculators, from simple models for general use to scientific and business calculators. During one semester they sold the following number of calculators at the prices given.

<i>Model</i>	<i>Price (in dollars)</i>	<i>Number Sold</i>
A	8	12
B	15	7
C	22	5
D	45	12
E	70	2
F	120	1

- a. Compute the measures of central tendency for the price of the calculators sold.
 - b. Which measure best describes how much was paid for a calculator?
15. The following frequency table of rent costs was compiled from a list of apartments for rent in a university town.

<i>Rent (in dollars)</i>	<i>Frequency</i>
226–250	2
251–275	7
276–300	12
301–325	8
326–350	2
351–375	2
376–400	4
401–425	1

Compute the mode, median, and mean for the data.