CHAPTER 8

PROBABILITY AND HYPOTHESIS TESTING

In earlier chapters, we discussed the steps which are important in conducting research: posing your research questions, identifying the variables, selecting the best way to measure the variables, constructing an appropriate design, and finally gathering the data and analyzing them.

We have also talked about how we may present the data as straightforward description: for nominal data, as frequency tables; for ordinal and interval data, in terms of some measure of central tendency ($\overline{X}$, median, or mode) and some measure of dispersion from the central tendency (variance, standard deviation, range). We have also talked about how we can standardize scores so that we can compare results from different test scores.

In one sense, then, we are now able to talk about our sample in a sensible way. Usually, however, we don't want to just *describe* our data. Our goal in research is to be able to say something about the population as a whole on the basis of our sample. We don't want to talk about just our 30 Thai students studying English but about Thai students in general. We don't want to just describe our 35 level 1 students but to say something about beginning language learners. We have already stressed the importance of selecting a sample which represents the population. We hope the 30 Thai *S*s are representative of all Thai students studying English. We have already suggested that one way to get a representative sample is to select the *S*s on a random basis. Still, even with random selection, there is always room for error. We can be sure that no matter how random the selection of the 30 *S*s was, the next 30 Thai *S*s would be slightly different.

Suppose we know that the mean listening comprehension score on a 35-item test for our Thai *S*s is 18. If we selected a second random sample of Thai *S*s, it is unlikely that we would come up with the same $\overline{X}$ of 18 for them. The mean scores of the two groups will not be exactly the same. Assuming that we do not have time to give the test to all Thai students studying English, and assuming that we may commit errors in sampling (even though we have selected the *S*s randomly), what can be done? How can we infer anything from our sample?

Don't despair; it's possible. To do this, we will use *inferential statistics* which build on what we have already learned in the past three chapters. The difference is that inferences can be made on the basis of the data from our sample by using statistical techniques and the principles of probability.

Probability may be a difficult concept to define, but everybody has a notion of what's probable and what is not. Weather reporters talk about the probability of rain, gamblers about the probability of winning a trip for two to beautiful Acapulco in the grocery store bingo game. We are all surprised when something improbable happens. We don't expect to win the trip to Acapulco nor do we expect to walk into class and find all our ESL students speaking perfect English. The probability of that happening is not absolutely zero, but it's not highly probable.

Obviously, if our predictions are based on the information we have, the probability of our making a correct prediction is higher than when we just wildly guess about something. When the weatherman says that the probability of rain is high, he uses previous information about rain at this time of year given the prevailing conditions, information collected for many years. When a salesman at a private language school tells you that the probability of your learning Mandarin during the course is 65%, you assume (perhaps naively) that he has information showing good results in learning Mandarin by a whole collection of people like yourself.

In using inferential statistics we are concerned with probability for we want to know how probable it is that we are making correct inferences. *In an experiment, the probability of getting the results we got is the proportion of times that* that *outcome would happen if the experiment were repeated an infinite number of times.* Remember that we've already said that if we measure any human behavior many, many times it will gradually approach a normal distribution (the bell-shaped curve). So, in our Thai example, we're asking how likely it is that the results from that sample represent the data we would get if we repeated the experiment over and over and over with all Thai students studying English. If we repeated the experiment over and over and over, we would expect the data to approach a normal distribution. So we can ask how probable it is that our sample fits into that normal distribution.

The usual example given to clarify probability is the flip of coins. Suppose you start flipping a fair coin (not bent and not weighted). You know that your chance of it coming up heads—that is, the probability—is 50%. You can test this out by recording the number of heads obtained. At the beginning, the proportion is very much dependent on the outcome of each flip, but as the number of times accumulates, the probability gradually settles down to .50:

Heads	4	53	139	237	480
Flips	10	100	300	500	1000
%(probability)	.40	.53	.46	.47	.48

If you had tails on both sides of the coin, the outcome of heads would be impossible, the probability zero. If you had a two-headed coin, the outcome would be certain and the probability 1.0. Probability, then, is always going to be somewhere betwen zero and 1, impossible and certain. If you have an unfairly weighted coin, there would still be two outcomes (heads or tails). If the heads turned out to have a probability of .82, the probability of tails must be .18.

We said that in an experiment probability is the proportion of times that we would get the same outcome if the experiment were repeated indefinitely. Suppose, then, that there are more than two possible outcomes. You know that when you roll a die, you can't get a probability of .50 because there are six possible outcomes. You may get a 1, 2, 3, 4, 5, or 6. If you want to get a 6, the probability of your getting it in one throw is 1 ÷ 6, the desired outcome divided by all the possible outcomes. The same thing happens in a multiple-choice test. If there are four answers, and the correct answer is (b), your chance of *guessing* accurately is 1 ÷ 4. That's what this direction says:

$$P\text{(probability)} = \frac{\text{number of desired events}}{\text{number of possible outcomes}}$$

Estimating probability may seem complicated when we get beyond coins or dice, but it's really not. Consider the probability of answering *two* true-false questions correctly just by guessing. You know there are four possible outcomes (the event you want—right, right; right, wrong; wrong, right; and wrong, wrong); so the probability is 1 ÷ 4, or .25. The probability of getting *one of the two* true-false questions right is 2 ÷ 4, or .50. Now we need to move to the probability of obtaining our research outcome if the research were repeated an infinite number of times. To do that, we must turn again to the normal distribution.

The probability of a certain score in the normal distribution

Let's go back to our beautiful bell-shaped curve, the normal distribution. In a normal distribution, half the scores (observations or *S*s) are above the mean and half below. Since the curve is symmetric, we can show the breakdown of the distribution under any part of the curve. We are easily able to discover the probability of any observation occurring at any point on the curve.

To do this we will begin by talking about the probability of *individual* scores fitting into the normal distribution. First we need to know what the mean and standard deviation are for the normal distribution of the data. As an example, let's pretend that all ESL teachers have had to take an applied linguistics examination in order to qualify for their state teaching credentials. According to the test publishers, the mean on the test is 65 and the standard deviation is 15. Now we give the exam to the teachers in our program. The first paper is corrected. What is the probability that the score will be higher than 65? Our knowledge of the normal distribution tells us that half the scores will be above the mean and half will be below the mean (see Figure 8-1). Since the mean is 65, we know that the probability of getting a score above 65 is .50. What are the chances the score will be between 50 and 80? This is where all the work on *z* scores comes in handy. We said that the standard deviation on our applied linguistics exam for ESL teachers was 15; so we can chart the curve in 15-point intervals and label the *z* scores as we have (see Figure 8-2). We can see that

Figure 8-1

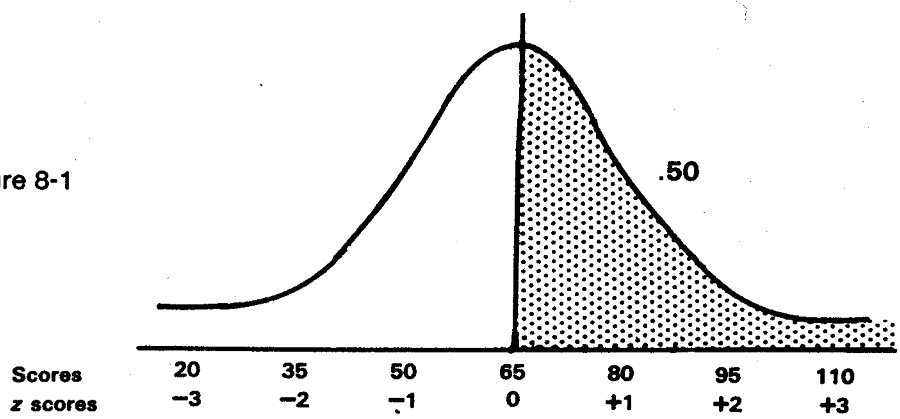

Figure 8-2

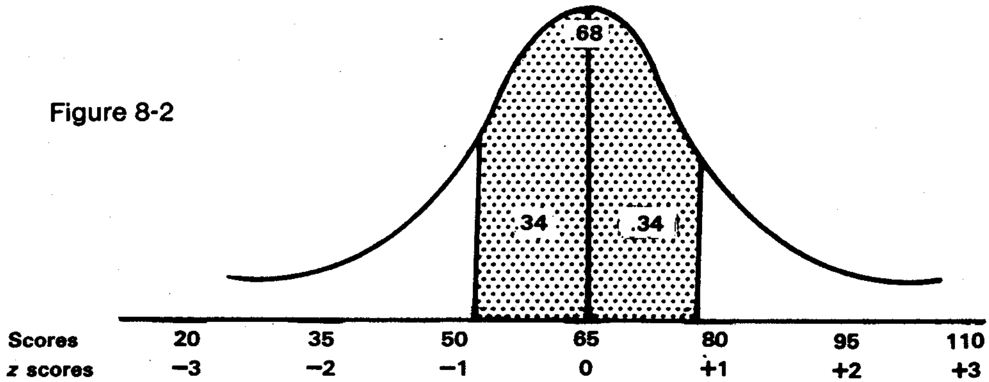

there is 1 s below the mean and 1 s above the mean for the score range of 50 to 80. We can simply look this up in the Appendix z score tables. It shows that between 0 (the mean) and ± 1.0, we will find .3413 of the observations. So we add these together and get .6826 or .68 as the probability of getting a z score between -1 and $+1$ (a raw score between 50 and 80). If we wonder how likely it is that the score on the paper will be between 35 and 95, a z score of $+2$ to -2, the table tells us that the chances are very good: .9544 or .95.

Now let's assume that we think that the teachers in our training program aren't ordinary. What is the probability that the S we have selected will score *over* 70? We might make the prediction thinking that our program is, of course, better than most programs! First we need to know how many standard deviations (the z score) 70 is above 65. Just think of standard deviation as a ruler (in this case 15 units long) and the z score as the number of ruler lengths. 70 is 5 points above the mean. Our ruler for standard deviation for the test is 15. Thus, $5 \div 15$ gives a z score of .33. We look this up in the Appendix. The table says that .1293 of the scores will fall between a z score of 0 and .33 (between 65 and 70). Since half the scores fall above the mean, we can subtract $(.50 - .13)$ to get the probability of .37. This means that there are 37 chances in 100 that the subjects' scores will be 70 or over (see Figure 8-3).

Figure 8-3

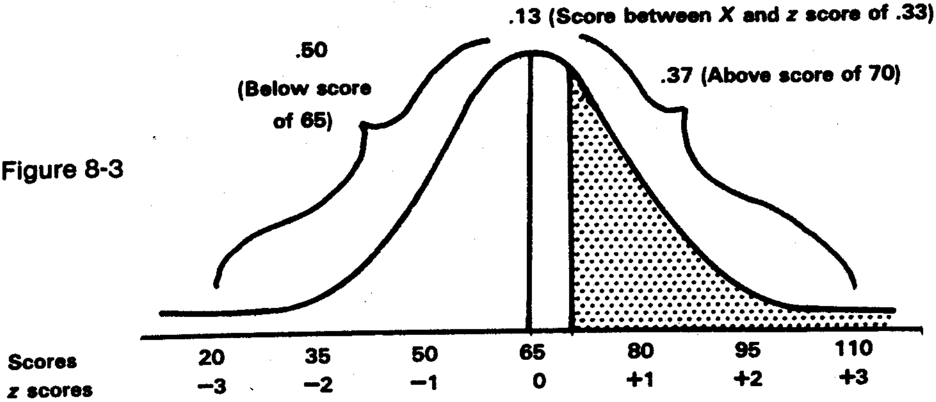

What is the probability that the score for our teacher-trainee will be over 100? (Meaning that we think our program doesn't just prepare people better but is definitely the best program in the country!) Again, we first measure off the number of standard deviations such an outcome would be above the mean of 65:

$$z \text{ score} = \frac{X - \overline{X}}{s} = \frac{100 - 65}{15} = \frac{35}{15} = 2.33$$

The table shows us that between a z score of 0 and 2.3 fall .4893 of the observations. We subtract this from .50 and find that the probability of obtaining this score is .01 (see Figure 8-4). It's highly unlikely that our teacher-trainee would be able to score that high. If she did, since it's improbable, then you would want to know why. It's not likely that it would happen by *chance.* We *could* say it was chance because there *is* one chance in one hundred that that might happen. However, we might also say it is so *un*likely that this would happen by

Figure 8-4

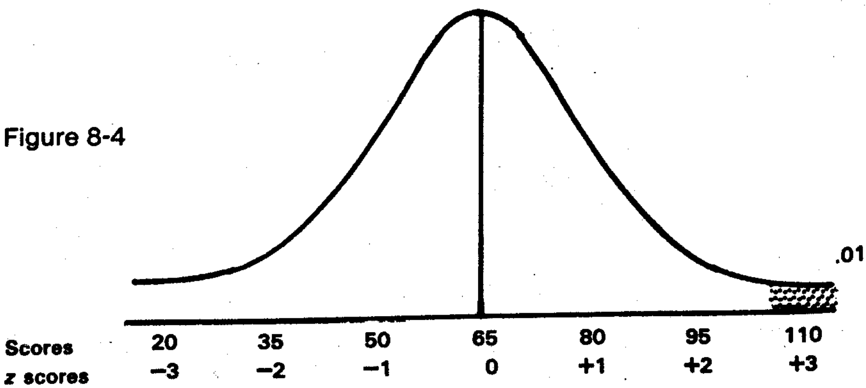

chance that there may be something about our program that made it possible for our teacher-trainee to obtain such a high score. Outsiders might check to see if we sold her a copy of the test for $100 on the day before the exam, or to see if the testing procedure was changed so that she heard the answers over a headset. She definitely doesn't "belong" in this distribution, even though she is part of the distribution.

To summarize this section, to find the probability of any individual score in the normal distribution:

1. Change the outcome to a z score by measuring the number of standard deviations from the mean.
2. The z score for the mean in a normal distribution is 0; the tables give you the probability that the score falls between 0 and your obtained z score.
3. To find the probability of getting *higher* than that z score, subtract from .50.
4. If the score falls below the mean, you treat it just the same way. For example, a -2 z score shows that .477 of the scores fall between the mean and -2, and the probability of getting a z score *below* -2 is .02 (very unlikely).

The process of hypothesis testing

We usually think about two things when we are making decisions (that is, when we are taking risks). First, what is the probability of getting the results we expect—how likely is it that we are right? Inferential statistics will help us answer this question. The second thing to think about in decision making is the cost—what do we lose if we're wrong? Here, inferential statistics cannot help us.

In our earlier example, we asked how probable it would be that a student from our teacher-training program would receive a particular score on a state test for teachers. We found that there was only 1 chance in 100 that the student would score over 100. It is not a typical score for the distribution (though of course we can't really claim that the only reason for such a high score is that our program is the best in the world!). We are often interested in the likelihood of a single score occurring at some point in a distribution of scores. However, it is more often the case that we want to discover the probability that our predictions, the hypotheses we make about research outcomes, are correct. To do this, we first must state our hypotheses in a way that allows us to compare the sample data with that of the population from which the sample is drawn.

For example, we might want to compare the data on a group of students who have received some special instruction with a population of foreign students who have not received that treatment; we expect that the experimental group will give us results which are better than the rest of the population. We expect a difference. We might, on the other hand, want to compare the data on the 30 Thai students with the population of Thai students studying English and hope that they are the same—that there is no difference. There are a number of possible hypotheses to reflect these expectations. However, the most common hypothesis is the null hypothesis, which states that there is no difference between the sample and the population. If you have strong evidence that leads you to expect not only a difference but the direction of that difference as well, you may use a directional hypothesis. The possible hypotheses, then, are:

H_0 = null hypothesis: there is no difference between the sample drawn from the population and the population.

H_1 = positive, directional hypothesis: there is a difference between the sample and the population; the sample *S*s will do better than the population.

H_1 = negative, directional hypothesis: there is a difference; the sample will not do as well as the population from which it was drawn.

H_1 = alternative hypothesis, no direction: there is a difference but the direction of the difference is not specified.

When we hope that there will be a difference between our sample and the population (that is, we hope that some special teaching techniques have helped our sample so that they will perform better than the population from which they were selected), we use the null hypothesis. This may seem strange since we hope there *is* a difference. The reason for this is that evidence that agrees with our predictions can't be conclusive grounds for accepting the hypothesis. Evidence that is *in*consistent with the hypothesis is good enough grounds for discarding it. This may seem a strange state of affairs, but remember that data that support your hypothesis might also be consistent with lots of other hypotheses or explanations as well as the one you wish to suggest. It doesn't tell you which of all the possible explanations is the most correct. Therefore, we use the null hypothesis and try to reject it. If we are able to reject the null hypothesis, we have support for the alternative hypothesis, the hypothesis of difference.

The next problem is to decide just how improbable a finding must be before we are allowed to reject the null hypothesis. Usually we want the probability of the findings falling in the distribution where they do to be very low indeed. The practice in most fields is not to reject the null hypothesis unless there are fewer than 5 chances in 100 (.05 probability level) that it could happen by chance. Others require an .01 level—1 chance in 100. The probability level chosen to reject the null hypothesis is called the *level of significance*.

Perhaps a graphic representation, as in Figure 8-5, will make this clearer. Let's assume that we have stated a null hypothesis, and that we have selected an .05 level of significance for rejecting the null hypothesis. If the results fall within the shaded area, the null hypothesis cannot be rejected, for the scores are typical of those that would normally be found in such a distribution. On the other hand, if the sample data fall in the area shaded in Figure 8-6 you can reject the null hypothesis. The scores are not those typically found in the distribution. If the data fall in the lower left tail, the sample is worse than the population. If the data fall in the far right tail, the sample is much better than the population. In either case, you can reject the hypothesis that the sample is no different from the population from which it was drawn (see Figure 8-6).

Figure 8-5

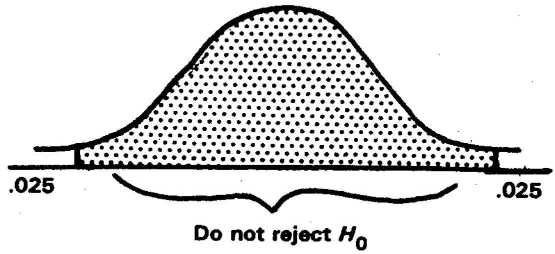

Figure 8-6

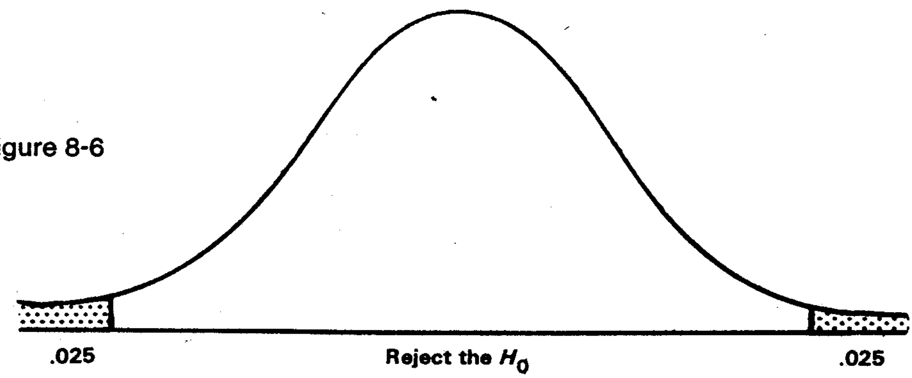

You will notice in Figures 8-5 and 8-6 that, although we chose an .05 level of significance, we have had to divide the .05 in two, giving an .025 area to each tail of the distribution. This is because we have formulated a null hypothesis which states that there is no difference between the sample and the population without saying where the difference might be. Since there is no direction specified for possible differences in the hypothesis, we must consider both tails of the distribution. This is, therefore, called a *two-tailed hypothesis.* When we reject the null hypothesis using a two-tailed test, we have evidence in support of the alternative hypothesis of difference (though the direction of the difference was not specified ahead of time).

The other hypotheses are directional. They predict a difference and the direction of the difference is specified. A pilot project or previous research points to an expected difference in one direction. This means that if we select an .05 level of significance, it will be for one tail, not two. These are *one-tailed hypotheses*. In order to find support for these hypotheses we must, again, reject the null hypothesis. However, this time we are concerned only with one tail of the distribution. In Figure 8-7, you can reject the null hypothesis and claim support for the positive directional hypothesis if the sample falls in the shaded area. For the negative direction hypothesis, we could reject the hypothesis if the data fell within the shaded area in the following diagram. Rejecting the null hypothesis gives you support for the negative directional hypothesis (Figure 8-8). As you can see, since we do not need to divide the .05 or .01 between two tails, it is easier to reject a directional, one-tailed hypothesis than a non-directional two-tailed hypothesis.

Figure 8-7

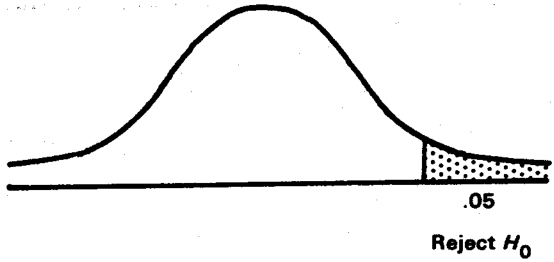

Figure 8-8

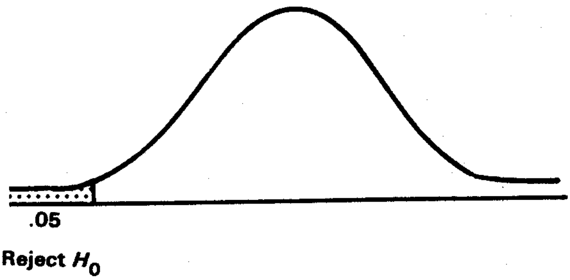

Once we have set a level of significance (whether .05 or .01) for rejecting the null hypothesis, we still need to know when a *z* score has reached that level. While you can always look the *z* value up in the Appendix to check the level of significance, you won't really need to. The value of *z* at the .05 level is 1.65 and for .025 it is 1.96. This value is referred to as the *critical value* of *z*. The Greek letter alpha α is used as an abbreviation for critical value. If the observed value of *z* is equal to or greater than the critical value, we can reject the null hypothesis. If you choose the .01 level as the level you demand before rejecting the hypothesis, then the value of *z* must be equal to or greater than 2.33 for a one-tailed hypothesis and 2.58 for a two-tailed hypothesis.

Critical values for *z* scores

1-tailed test	2-tailed test
.05 critical value = 1.65	.05 critical value = 1.96 (.025)
.01 critical value = 2.33	.01 critical value = 2.58 (.005)

Notice that the further the finding is from the mean (the larger the *z* score), the smaller the probability that the finding really "fits" into the curve with the mean that has been established for the distribution.

To summarize, we can use inferential statistics to determine how likely we are to be right. To do this, we first formulate our hypothesis, usually as a null hypothesis. By rejecting the null hypothesis we will be able to support the correctness of the alternative hypothesis. We decide on a level of significance (usually .05 or .01) at which we will reject the null hypothesis. Then we run the experiment, do the computations, and determine how small the probability is for that particular finding. If it's smaller than .05 or .01, then we can reject the null hypothesis; we can feel confident that it is an unusual finding, one that doesn't happen often by chance. If it's greater than .05 or .01, on the other hand, we must consider the null hypothesis correct. We must conclude that our finding is just part of a normal pattern; there's nothing special about it.

This procedure allows us to answer the first question—how likely we are to be right. The second question is what we lose if we're wrong. Both questions must be answered when we make decisions. In medical experiments, for example, lives depend on being right. Medical researchers might, then, decide to set their probability levels at .00000001 since they cannot afford to be wrong. In our field, we might be willing to set an even more lenient probability level than .05 or .01. For example, if you had conditional approval of state funding for a bilingual playschool for your community contingent on your raising $500 by midnight, you would most likely take any suggestion no matter what the probability was that it would succeed. That's because the answer to "what do we lose if we're wrong" is "everything." If you are extremely concerned about low reading scores of bilingual children in your school and find that bilingual children instructed in some new program make gains which are significant at the .25 level, you would undoubtedly accept that probability level, since it seems to be the best you can find. In other words, probability levels allow us to see the likelihood of our findings being not due to chance but related to our treatment. But decision making is based on many factors which make us opt for a higher or lower probability level in many instances.

In this chapter we have introduced two basic concepts: probability and the notion of hypothesis testing. In presenting the concept of probability, we began the discussion by finding the probability of some *individual* score or observation falling within the normal distribution on some test which has an established mean and standard deviation. We next discussed probability in terms of group data. While the general concept of probability in terms of a distribution does not change, we are more frequently expected to look at the probability of *group* data as being representative of or very different from the distribution curve of the population. An understanding of probability is basic to inferential statistics, the statistics that allow us to generalize from our sample data. These basic principles of probability and their relationship to inferential statistics will be discussed further in the following chapters.

ACTIVITIES

1. Try flipping a coin 100 times. Record the number of heads as they cumulate over every 10 tosses.
2. There are four schools in your district which have registered the same need for Title 1 funds. The funds will not be split; only 1 will get them all. What is the probability your school will be the winner-take-all? What are some things that might change the probability of your school being chosen?
3. The average IQ is 100; s is 15. What is the probability that your score is over 130? Assuming that it's 130, draw and label the curve showing the location of your score.
4. Look at the normal curve table. What is the probability of getting a z score greater than 2.1? Between −.6 and 1.7? Less than 1.9? Between 2.9 and 3.1? Over 2.5?
5. The distribution of language aptitude test scores is normal with a mean of 70 and a standard deviation of 15. You were one of the *S*s tested this year. What is the probability that your score was *above* 70? Above 95? Below 40?
6. Suppose we wanted to raise money to send an outstanding student to next year's TESOL Conference. We sold 548 raffle tickets (the prize was free coffee all year in the department office for the raffle winner). What is the probability that ticket 502 will win? Ticket 12? On the basis of your answer to this question, do you think you qualify to go to TESOL?
7. You are teaching reading in an ESP (English Special Purposes—Science) class. While reading technical material, *S*s' words per minute average a mean of 185, s of 12.5. What is the probability that an *S* absent that day would have a score under 164? Between 164 and 185? Over 200? What would such information tell you?
8. Decision making requires an understanding not only of probability but of costs. Assume you are head of a school. A new set of very expensive and very innovative multimedia and computer-assisted teaching materials is now available. They have an incredible record for student improvement to support the sales pitch (at .01 level of significance). Your students are also very impressed by technological toys. What do you stand to lose if you do or do not purchase them?
9. Assume you are charged with making recommendations to an appropriations committee on grant requests for additional funds for continuing programs to help economically disadvantaged children in schools.

   #1 Presents ordinal data on the children's pronunciation of English and Spanish, letters from parents and from nationally known experts who served as consultants to the program, newspaper clippings about the program, and illustrated stories written by the children.

   #2 Presents results on the Peabody Picture Vocabulary Test in Spanish and English, ITPA scores, and gains in WISC scores which are positive but not statistically significant.

   #3 Includes pictures of children, teachers, and parents at a school fiesta, claims that three migrant families have decided to reside permanently in the community, and includes scores showing significant (.05 level) gains for *S*s on Ravem Progressive Matrices Test.

   #4 Describes its program in detail, gives raw data on SES, parental aspirations for the child, pretest data on *S*s' Draw-A-Man, Berko test of English morphology, and a self-concept test. The sponsors from a leading university regret that no further data were available at the time of the filing deadline.

   #5 Presents a 30-minute videotape sample of classroom behaviors, data showing a gain (.05) on a locally constructed proficiency test for each language, evaluation material on their paraprofessional aides, and a list of student-preferred materials.

   #6 Presents data on the effect of an intensive oral language program on reading scores. Oral language scores show gains on certain grammar structures (.01 level for 4 structures, .05 for 2, and .10 for 4). Reading scores improved but not significantly. They plan to revise the oral component to make it more effective.

   If one request is sure to be approved, what is the probability of approving each one? List the factors that you feel change each program's chances of being chosen. Which (if any) would you recommend for continuance of funding?

WORKING WITH THE COMPUTER

The purpose of the computer assignment this time is to continue your acquaintances with the computer and to show you that it can give you further information about your data. The concepts are not new; they were presented in Chapter 5.

Frequency Analysis

To run a frequency analysis, the procedures will be similar to our last computer task. The job control cards will be exactly the same. The data control cards will also be the same. However, this time we want to create another variable—the total score. To create new variables which require the computer to perform mathematical operations on the existing variables, we simply use a "compute" card and ask the computer to do the operation. For example, if we want to create the total score variable, we know that all the scores for one person must be added together. So we can punch a compute card that says:

Column 1 *Column 16*

COMPUTE TOTAL=DICTAT+CLOZE+LIS+GRAM+READ

This card will make the computer calculate total scores for all the *S*s and include it as a variable even though it is not mentioned in the variable list.

The compute card can tell the computer to add (+), subtract (−), divide (/), or multiple (*) any of your variables. The general format of the compute card for these operations will be:

Column 1 *Column 16*

COMPUTE NEW_VARIABLE=VARIABLE_X $\left\{\begin{matrix} + \\ - \\ / \\ * \end{matrix}\right\}$ VARIABLE_Y

(Of course, you name the new variable and tell the computer which variables to add or subtract or multiply or divide.)

Insert the compute card immediately after the input format card. It is one additional card, then, in the data control cards. Now you must create a new command card to run the frequency analysis:

Column 1 *Column 16*

FREQUENCIES INTEGER=DICTAT___TO___TOTAL(0,300)

This card means that all the data are in numbers (integers) and not alphabetic. It also means that you want the computer to give you a frequency analysis of all the variables, dictation to total and everything in between those two. The numbers in the parentheses indicate the minimum and maximum scores on the data. Usually, the minimum is zero; the maximum will be the highest possible score for the data (300 for our data).

Now you are ready to run the data again using the same order as in the last computer assignment: job control cards, data control cards (including the new compute card), and the command cards with the new frequencies card. Next add the data cards and then the finish and / / cards. Run the program. (Be sure to remove the Condescriptive card!)

This time your printout will have the following information for all variables:

1. Descriptive Statistics. This table includes the 12 categories discussed in the last computer assignment plus two new statistics.

MODE	The most frequently obtained score.
MEDIAN	The score at the 50th percentile.

2. Frequencies. This part includes five columns:

CODE	The score whose frequency is computed.
ABSOLUTE FREQUENCY	The number of times that score was obtained.
RELATIVE FREQUENCY (percent)	The absolute frequency divided by the total number of cases and multiplied by 100 to get the percent. $(f/N)(100)$
ADJUSTED FREQUENCY	This indicates any adjustments the computer may make because of number of occurrences or other mathematical operations (the value should not be very different from relative frequency and we need not be too concerned with this entry).
CUMULATIVE FREQUENCY	Frequencies are successively added together to get the F. This is divided by total number of cases and multiplied by 100 to get the percentile figure.

If you are confused about any of the terms on the printout, please refer back to Chapter 5 for a fuller definition of each.

If you would like to have information about the distribution of the data in graphic form, you can add a command card which asks the computer to do this. This card is called the option card, and the number related to histograms is 8. So, to get a histogram, you should use the following card:

Column 1	*Column 16*
OPTIONS	8

Insert this option card after the frequencies card and before the statistics card in your data deck. With this card, you will get a histogram (like the one below) for each variable.

(I)

VARIABLE READ

MEAN	14.000	STD ERR	1.205	MEDIAN	17.500
MODE	23.000	STD DEV	7.623	VARIANCE	58.103
KURTOSIS	-1.347	SKEWNESS	-0.439	RANGE	23.000
MINIMUM	0.0	MAXIMUM	23.000		

VALID CASES 40 MISSING CASES 0

(II) FREQUENCIES

VARIABLE READ

CATEGORY LABEL	(A) CODE	(B) ABSOLUTE FREQUENCY	(C) RELATIVE FREQUENCY (PERCENT)	(D) ADJUSTED FREQUENCY (PERCENT)	(E) CUMULATIVE ADJ FREQ (PERCENT)
	0	1	2.5	2.5	2.5
	1	2	5.0	5.0	7.5
	3	1	2.5	2.5	10.0
	4	1	2.5	2.5	12.5
	5	4	10.0	10.0	22.5
	6	2	5.0	5.0	27.5
	7	1	2.5	2.5	30.0
	8	1	2.5	2.5	32.5
	9	1	2.5	2.5	35.0
	11	2	5.0	5.0	40.0
	14	1	2.5	2.5	42.5
	16	1	2.5	2.5	45.0
	17	2	5.0	5.0	50.0
	18	4	10.0	10.0	60.0
	19	5	12.5	12.5	72.5
	20	2	5.0	5.0	77.5
	21	2	5.0	5.0	82.5
	22	1	2.5	2.5	85.0
	23	6	15.0	15.0	100.0
	TOTAL	40	100.0	100.0	

(III) HISTOGRAM

VARIABLE READ

```
CODE
       I
   0   ****** (       1)
       I
       I
   1   *********** (       2)
       I
       I
   3   ****** (       1)
       I
       I
   4   ****** (       1)
       I
       I
   5   ********************* (       4)
       I
       I
   6   *********** (       2)
       I
       I
   7   ****** (       1)
       I
       I
```

(continued, p. 94)

(continued from p. 93)

```
 8   ****** (        1)
     I
     I
 9   ****** (        1)
     I
     I
11   *********** (        2)
     I
     I
14   ****** (        1)
     I
     I
16   ****** (        1)
     I
     I
17   *********** (        2)
     I
     I
18   ********************* (        4)
     I
     I
19   ************************** (        5)
     I
     I
20   *********** (        2)
     I
     I
21   *********** (        2)
     I
     I
22   ****** (        1)
     I
     I
23   ******************************* (        6)
     I
     I.........I.........I.........-.........I.........I
     0         2         4         6         8        10
     FREQUENCY
```

Having run the sample problem data twice, and having ironed out any problems you may have encountered in learning to order the computer about, you should feel more comfortable working with the computer. The operations you have asked the computer to do have been very simple so far. You could have done the calculation as fast, or faster, by hand. In your own research, however, you are likely to have more *S*s and more variables and you may want to run many different analyses. By starting with a simple example, you will learn the mechanics of using the computer. This should give you confidence so that you will be ready to give instructions and read the answers when dealing with much more complicated data.

Suggested further reading for this chapter: Guilford and Fruchter; Slakter; and the SPSS Manual.