Chapter seven

Defining the Observed Significance Level of a Test: A Simple Example Using the Binomial Distribution

Sampling distributions

What are They?

What Role Do They Play in Inferential Statistics?

Probabilities and probability distributions

How Does One Calculate the Probability of a Given Outcome?

What are Probability Distributions?

How are They Used?

The binomial distribution

What is It?

How is It Calculated?

What are Its Characteristics?

D. Weisburd and C. Britt, *Statistics in Criminal Justice*, DOI 10.1007/978-1-4614-9170-5_7,

When we make inferences to a population, we rely on a statistic in our sample to make a decision about a population parameter. At the heart of our decision is a concern with Type I error. Before we reject our null hypothesis, we want to be fairly confident that it is in fact false for the population we are studying. For this reason, we want the observed risk of a Type I error in a test of statistical significance to be as small as possible. But how do statisticians calculate that risk? How do they define the observed significance level associated with the outcome of a test?

The methods that statisticians use for calculating the observed significance level of a test of statistical significance vary depending on the statistics examined. Sometimes these methods are very complex. But the overall logic that underlies these calculations is similar, irrespective of the statistic used. Thus, we can take a relatively simple example and use it as a model for understanding how the observed significance level of a test is defined more generally in statistics. This is fortunate for us as researchers, because it means that we do not have to spend all of our time developing complex calculations to define risks of error. Once we understand how risks of error are defined for one problem, we can let statisticians calculate the risks for other more complex problems. Our concern is not with the calculations themselves, but with understanding the general logic that underlies them.

We begin this chapter by discussing a very simple decision. When should we begin to suspect that a coin used in a coin toss is unfair or biased? Ordinarily, we might come to a conclusion based on common sense or intuition. In statistics, we take a more systematic approach, relying on the logic of hypothesis testing and a type of distribution called a sampling distribution. Using this example of the coin toss and a sampling distribution called the binomial distribution, we illustrate how statisticians use probability theory to define the observed significance level, or risk of Type I error, for a test of statistical significance.

The Fair Coin Toss

Imagine that you and your friends play a volleyball game each week against a group of criminal justice students from another school. You always begin the game with a coin toss to decide who will serve the ball first. Your opponents bring an old silver dollar, which you have agreed to use for the toss. They choose heads and continue to choose heads each time you play. At first, this does not seem like a problem. However, each week you play, the coin comes up heads and they serve the ball.

Suppose that this happened for four straight weeks. Would you begin to become suspicious? What if it went on for six weeks? How many times in a row would they have to win the coin toss before you and your team accused them of cheating? Would they have to win for ten or twenty weeks? You might worry about accusing them too quickly, because you know that even if the coin is fair it sometimes happens that someone is lucky and just keeps on winning. You would want to be fairly certain that the coin was biased before concluding that something was wrong and taking some action.

In everyday life, you are likely to make this decision based on intuition or prior experience. If you ask your classmates, each one is likely to come up with a slightly different number of coin tosses before he or she would become suspicious. Some students may be willing to tolerate only four or five heads in a row before concluding that they have enough evidence to accuse their opponents of cheating. Others may be unwilling to reach this conclusion even after ten or fifteen tosses that come up heads. In part, the disagreement comes from personality differences. But more important, guesswork or common sense does not give you a common yardstick for deciding how much risk you take in coming to one conclusion or another.

Sampling Distributions and Probability Distributions

Statistical inference provides a more systematic method for making decisions about risk. The coin toss can be thought of as a simple test of statistical significance. The research hypothesis is that the coin is biased in favor of your opponents. The null hypothesis is that the coin is fair. Each toss of the coin is an event that is part of a sample. If you toss the coin ten times, you have a sample of ten tosses. Recall from Chapter 6 that Type I error is the error of falsely rejecting the null hypothesis that the coin is fair. If you follow the common norm in criminal justice, then you are willing to reject the null hypothesis if the risk of a Type I error is less than 5%.

But how can we calculate the risk of a Type I error associated with a specific outcome in a test of statistical significance, or what we generally term the observed significance level of a test? One simple way to gain an estimate of the risk of unfairly accusing your friends is to check how often a fair coin would give the same result as is observed in your series of volleyball games. For example, let's say that you played ten games in a season and in all ten games the old silver dollar came up heads (meaning that the opposing team won the toss). To check out how often this might happen just by chance when the coin is in fact a fair one, you might go to a laboratory with a fair coin and test this out in practice. One problem you face is deciding how many samples or trials you should conduct. For example, should you conduct one trial or sample by flipping your fair coin just ten times and stopping? Or should you conduct multiple trials or samples, each with ten tosses of the fair coin? Clearly, one trial or sample of ten tosses will not tell you very much. Indeed, one of the reasons you have gone to the laboratory is that you know it sometimes happens that a fair coin will come out heads ten times in a row. What you want to know is how rare an event this is. How often would you gain ten heads in a row in a very large number of samples or trials of a fair coin?

The distribution that is gained from taking a very large number of samples or trials is called a **sampling distribution.** In principle, one could create a sampling distribution by drawing thousands and thousands of samples from a population. For example, in the case of our coin toss, we might conduct thousands of trials of ten flips of a fair coin. If we recorded the outcome for each trial and placed our results in a frequency distribution, we would have a sampling distribution for a sample of ten tosses of a fair coin.

This sampling distribution would allow us to define the risk of a Type I error we would face in rejecting the null hypothesis that the old silver dollar is fair. For example, suppose that in the sampling distribution we gained a result of ten heads in only 1 in 1,000 samples. If we reject the null hypothesis in this case, our risk of making a Type I error, according to the sampling distribution, is only 0.001. This is the observed significance level of our test of statistical significance. In only 1 in 1,000 samples of ten tosses of a fair coin would we expect to gain a result of ten heads. If the old silver dollar was indeed a fair coin, it would seem very unlikely that on our one trial of ten tosses of the coin each toss would come out heads. Of course, in making our decision we cannot be certain that the silver dollar used in the volleyball toss is not a fair coin. While the occurrence of ten heads in ten tosses of a fair coin is rare, it can happen about once in every 1,000 samples.

Building a sampling distribution provides a method for defining our risk of a Type I error. However, it is very burdensome to create a sampling distribution by hand or even in the laboratory. If you try out our example of ten tosses of a fair coin, you will see that developing even 100 samples is not easy. If we had to actually construct a sampling distribution every time we wanted to make a decision about a hypothesis, it would be virtually impossible to make statistical inferences in practice.

Fortunately, there is another method we can use for creating sampling distributions. This method relies on probability theory, rather than a burdensome effort to collect samples in the real world. Because we use probabilities, the distributions that are created using this method are called **probability distributions.** Importantly, though we rely on probability theory because it is very difficult to develop sampling distributions in practice, we do not suffer for our approach. This is because probability theory allows us to calculate the outcomes one would expect in a perfect world. In the real world, we might flip the coin slightly differently as we got tired or the coin might become worn on one side or another, thus affecting the outcomes we gain. In probability theory, we remove the imperfections of the real world from our estimates.

The Multiplication Rule

In order to estimate the risk of a Type I error in the case of a series of tosses of a fair coin, we can use the **multiplication rule,** a simple rule about probabilities drawn from probability theory. The multiplication rule tells you how likely we are to gain a series of events one after another—in this case, a series of outcomes in a toss of a coin. It allows us to estimate theoretically how often we would gain a specific series of events if we drew an infinite number of samples. The multiplication rule generally used to establish probabilities in statistics is based on the assumption that each event in a sample is **independent** of every other event. In the case of the coin toss, this means that the outcome of one toss of a coin is unaffected by what happened on the prior tosses. Each time you toss the coin, it is as if you started with a clean slate. That would seem a fairly reasonable assumption for our problem. What worries us is that the coin is unfair overall, not that it is becoming less or more unfair as time goes on.

An example of a series of events that are not independent is draws from a deck of cards. Each time you draw a card, you reduce the number of cards left in the deck, thus changing the likelihood of drawing any card in the future. For example, let's say that on your first draw from a deck of 52 cards you drew an ace of spades. On your second draw, you cannot draw an ace of spades because you have already removed it from the deck. The likelihood of drawing an ace of spades on the second draw has thus gone from 1 in 52 to 0 in 51. You have also influenced the likelihood of drawing any other card because there are now 51, not 52, cards left in the deck. If you want a series of draws from a deck to be independent of one another, you have to return each card to the deck after you draw it. For example, if you returned the ace of spades to the deck, the chance of choosing it (assuming the deck was mixed again) would be the same as it was on the first draw. The chances of choosing any other card would also be the same because you once again have all 52 cards from which to draw.

The multiplication rule for four independent events is stated in Equation 7.1. It says that the likelihood of any series of events, represented as *A, B, C,* and *D,* happening one after another is equal to the probability of event *A* times the probability of event *B* times the probability of event *C* times the probability of event *D*. The rule can be extended to as many events as you like. We have chosen four here, because this was the number of volleyball games we began with. If you wanted to extend the rule, you would simply include the additional events on the left side of the equation (for example, *E* and *F*) and include the probability of each on the right side of the equation [e.g., $P(E) \cdot P(F)$].

$$P(A \,\&\, B \,\&\, C \,\&\, D) = P(A) \cdot P(B) \cdot P(C) \cdot P(D) \qquad \textbf{Equation 7.1}$$

Extending this logic to our example of the coin toss is straightforward. The probability of *A* and *B* and *C* and *D* can represent the probability that four tosses in a row come up heads. Our main problem is to establish what the probability is of a head coming up on any particular toss of the coin. In this, we are helped by our null hypothesis, which states that the coin is fair. If the coin is fair, then there should be an even chance of a head or a tail coming up on any particular toss of the coin. Put differently, under the assumption of the null hypothesis that the coin is fair, the likelihood of a head coming up is 0.50.

What, then, does the multiplication rule tell us about the chances of getting four heads in a row on four tosses of a fair coin? In part a of Table 7.1, we calculate that probability by multiplying 0.50 (the likelihood of gaining a head on any toss of a fair coin) by itself four times, to represent four tosses of an unbiased coin. The result is 0.0625. If you had decided at the outset to make a decision about the null hypothesis—that the coin is fair—after four tosses of the coin, then you have conducted a type of test of statistical significance. For this test, the observed significance level (or risk of a Type I error) of rejecting the null hypothesis on the basis of four heads is 0.0625.

If you use the norms of criminal justice research, this is not enough, however, for you to reject the null hypothesis that the coin is fair. Criminal justice researchers generally want the risk of falsely rejecting the null hypothesis to be less than 5%. A bit over 6% is still more than the 5% significance criterion that is used by convention. So if you had decided to

Table 7.1 Probabilities Associated with Tosses of a Fair Coin

a. $P(A \,\&\, B \,\&\, C \,\&\, D) = P(A) \cdot P(B) \cdot P(C) \cdot P(D) = (0.50)(0.50)(0.50)(0.50) = 0.0625$

b. $P(A \,\&\, B \,\&\, C \,\&\, D \,\&\, E) = P(A) \cdot P(B) \cdot P(C) \cdot P(D) \cdot P(E) = (0.50)(0.50)(0.50)(0.50)(0.50) = 0.0313$

c. $P(A \,\&\, B \,\&\, C \,\&\, D \,\&\, E \,\&\, F \,\&\, G \,\&\, H \,\&\, I \,\&\, J) = P(A) \cdot P(B) \cdot P(C) \cdot P(D) \cdot P(E) \cdot P(F) \cdot P(G) \cdot P(H) \cdot P(I) \cdot P(J)$
$= (0.50)(0.50)(0.50)(0.50)(0.50)(0.50)(0.50)(0.50)(0.50)(0.50) = 0.0010$

make a decision about the fairness of the coin after four coin tosses, you would probably not want to reject the null hypothesis that the coin is fair and confront your opponents. Under this criterion, the likelihood of falsely rejecting the null hypothesis, or the observed significance level of your test, would have to be below 0.05.

What if you had decided at the outset to make a decision about the null hypothesis after five tosses of a coin? Would a result of five heads in a row lead you to reject the null hypothesis? As illustrated in part b of Table 7.1, the multiplication rule tells you that the likelihood of getting five heads in a row if the coin is fair is 0.0313. This is less than our threshold of 0.05, and thus would lead you to reject the null hypothesis. Is this consistent with your earlier commonsense conclusions? Students are usually surprised at how quickly they reach the 0.05 significance threshold in this example.

If you had decided at the outset that you would need ten or fifteen heads in a row, you may want to reconsider, given what we have learned from the multiplication rule. The likelihood of getting ten heads in a row in ten tosses of a fair coin is only 1 in 1,000 (see part c of Table 7.1). The likelihood of getting fifteen heads in a row in fifteen tosses of a fair coin is even lower: about 3 in 100,000. In both of these cases, you would take a very small risk of a Type I error if you rejected the null hypothesis. Nonetheless, the multiplication rule tells us that, even if the coin is fair, it is possible to get ten or even fifteen heads in a row. It just does not happen very often.

The multiplication rule allows us to estimate how often we would expect to get a series of specific outcomes in a very large number of trials or samples, without actually going out and doing the hard work of constructing a sampling distribution in the real world. However, the problem as examined so far assumes that the coin will come up heads every time. What if the coin comes up heads generally, but not all the time? For example, what if you play ten games and the coin comes up heads nine times? The situation here is not as one-sided. Nonetheless, it still seems unlikely that your opponents would win most of the time if the coin were fair. The multiplication rule alone, however, does not allow us to define how likely we are to get such a result.

Different Ways of Getting Similar Results

The multiplication rule allows us to calculate the probability of getting a specific ordering of events. This is fine so far in our coin toss because in each example we have chosen there is only one way to get our outcome. For example, there is only one way to get five heads in five coin

Table 7.2

Arrangements for Nine Successes in Ten Tosses of a Coin

Arrangement 1	●	○	○	○	○	○	○	○	○	○
Arrangement 2	○	●	○	○	○	○	○	○	○	○
Arrangement 3	○	○	●	○	○	○	○	○	○	○
Arrangement 4	○	○	○	●	○	○	○	○	○	○
Arrangement 5	○	○	○	○	●	○	○	○	○	○
Arrangement 6	○	○	○	○	○	●	○	○	○	○
Arrangement 7	○	○	○	○	○	○	●	○	○	○
Arrangement 8	○	○	○	○	○	○	○	●	○	○
Arrangement 9	○	○	○	○	○	○	○	○	●	○
Arrangement 10	○	○	○	○	○	○	○	○	○	●

○ = Head; ● = Tail

tosses or ten heads in ten coin tosses. In each case, your opponents must toss a head before each game. This would be the situation as well if your opponents tossed tails ten times in ten coin tosses. However, for any outcome in between, there is going to be more than one potential way to achieve the same result.

For example, if your opponents tossed nine heads in ten coin tosses, they could win the coin toss nine times (with a head) and then lose the toss in the tenth game (with a tail). Or they could lose the first toss (with a tail) and then win the remaining nine. Similarly, they could lose the second, third, fourth, fifth, sixth, seventh, eighth, or ninth coin toss and win all the others. Each of these possible ordering of events is called an **arrangement.** As is illustrated in Table 7.2, there are ten possible arrangements, or different ways that you could get nine heads in ten coin tosses. In the case of ten heads in ten coin tosses, there is only one possible arrangement.

It is relatively simple to list all of the arrangements for our example of nine heads in ten coin tosses, but listing becomes very cumbersome in practice as the split of events becomes more even. For example, if we were interested in how many ways there are of getting eight heads in ten coin tosses, we would have to take into account a much larger number of arrangements. As Table 7.3 illustrates, it takes a good deal of effort to list every possible arrangement even for eight heads. In the case of a more even split of events—for example, five heads in ten tosses—it becomes extremely cumbersome to list each arrangement one by one. Because of this, we generally use the formula in Equation 7.2 to define the number of arrangements in any series of events.

$$\binom{N}{r} = \frac{N!}{r!(N-r)!} \qquad \textbf{Equation 7.2}$$

On the left side of this equation we have N "choose" r, where N is the number of events in the sample and r is the number of successes in the

Table 7.3 Arrangements for Eight Successes in Ten Tosses of a Coin

1: ●●○○○○○○○○	16: ○●○○○○○○●○	31: ○○○○●●○○○○
2: ●○●○○○○○○○	17: ○●○○○○○○○●	32: ○○○○●○●○○○
3: ●○○●○○○○○○	18: ○○●●○○○○○○	33: ○○○○●○○●○○
4: ●○○○●○○○○○	19: ○○●○●○○○○○	34: ○○○○●○○○●○
5: ●○○○○●○○○○	20: ○○●○○●○○○○	35: ○○○○●○○○○●
6: ●○○○○○●○○○	21: ○○●○○○●○○○	36: ○○○○○●●○○○
7: ●○○○○○○●○○	22: ○○●○○○○●○○	37: ○○○○○●○●○○
8: ●○○○○○○○●○	23: ○○●○○○○○●○	38: ○○○○○●○○●○
9: ●○○○○○○○○●	24: ○○●○○○○○○●	39: ○○○○○●○○○●
10: ○●●○○○○○○○	25: ○○○●●○○○○○	40: ○○○○○○●●○○
11: ○●○●○○○○○○	26: ○○○●○●○○○○	41: ○○○○○○●○●○
12: ○●○○●○○○○○	27: ○○○●○○●○○○	42: ○○○○○○●○○●
13: ○●○○○●○○○○	28: ○○○●○○○●○○	43: ○○○○○○○●●○
14: ○●○○○○●○○○	29: ○○○●○○○○●○	44: ○○○○○○○●○●
15: ○●○○○○○●○○	30: ○○○●○○○○○●	45: ○○○○○○○○●●

○ = Head; ● = Tail

total number of events. In our case, N is the number of coin tosses and r is the number of times that the coin comes up heads. Put together, this statement establishes our question: How many ways are there of gaining r heads in N tosses of a coin? To answer our question, we need to solve the right side of the equation. Each of the terms in the equation is defined as a **factorial,** indicated by the symbol !. When we take a factorial of a number, we merely multiply it by all of the whole numbers smaller than it. For example, 3! is equal to (3)(2)(1), or 6. Because factorials get very large very quickly, a table of factorials is presented in Appendix 1. Note that 0! = 1. Applied to our problem of nine heads in ten coin tosses, Equation 7.2 is worked out below:

Working It Out

$$\binom{N}{r} = \frac{N!}{r!(N-r)!}$$

$$\binom{10}{9} = \frac{10!}{9!(10-9)!}$$

$$= \frac{10!}{9!\,1!}$$

$$= \frac{3{,}628{,}800}{362{,}880(1)}$$

$$= 10$$

Using this method, we get the same result as before. There are ten possible arrangements to get nine heads in ten tosses of a coin. When we apply Equation 7.2 to the problem of five heads in ten coin tosses, its usefulness becomes even more apparent. There are 252 different ways of getting five heads in ten tosses. Listing each would have taken us considerably longer than the calculation below.

Working It Out

$$\binom{N}{r} = \frac{N!}{r!(N - r)!}$$

$$\binom{10}{5} = \frac{10!}{5!(10 - 5)!}$$

$$= \frac{10!}{5!\ 5!}$$

$$= \frac{3{,}628{,}800}{120(120)}$$

$$= 252$$

Solving More Complex Problems

Now that we have a method for calculating arrangements, we can return to our original problem, which was to define the probability of your opponents tossing the coin in ten games and getting heads nine times. Because there are ten different ways of getting nine heads in ten coin tosses, you need to add up the probabilities associated with these ten sequences. This is what is done in Table 7.4. The multiplication rule is used to calculate the probability for each sequence, or arrangement, under the assumption of the null hypothesis that the coin is fair. Because our null hypothesis states that the coin is fair, we can assume that the chances of gaining a head and a tail are even. The probability of any event, whether a head or a tail, is 0.50, and the probability of a sequence of ten events is always the same. This makes our task easier. But it is important to note that if the null hypothesis specified an uneven split (for example, 0.75 for a head and 0.25 for a tail), then each of the sequences would have a different probability associated with it. In any case, the likelihood of getting any one of these sequences is about 0.001, rounded to the nearest thousandth. When we add together the ten sequences, we get a probability of 0.010.

Table 7.4 The Sum of Probabilities for All Arrangements of Nine Heads in Ten Tosses of a Fair Coin

											PROBABILITY
Arrangement 1	●	○	○	○	○	○	○	○	○	○	0.001
Arrangement 2	○	●	○	○	○	○	○	○	○	○	0.001
Arrangement 3	○	○	●	○	○	○	○	○	○	○	0.001
Arrangement 4	○	○	○	●	○	○	○	○	○	○	0.001
Arrangement 5	○	○	○	○	●	○	○	○	○	○	0.001
Arrangement 6	○	○	○	○	○	●	○	○	○	○	0.001
Arrangement 7	○	○	○	○	○	○	●	○	○	○	0.001
Arrangement 8	○	○	○	○	○	○	○	●	○	○	0.001
Arrangement 9	○	○	○	○	○	○	○	○	●	○	0.001
Arrangement 10	○	○	○	○	○	○	○	○	○	●	0.001
									Total Probability:		0.01

Probability of throwing each arrangement of 10 throws
$= P(A) \cdot P(B) \cdot P(C) \cdot P(D) \cdot P(E) \cdot P(F) \cdot P(G) \cdot P(H) \cdot P(I) \cdot P(J)$
$= (0.50)(0.50)(0.50)(0.50)(0.50)(0.50)(0.50)(0.50)(0.50)(0.50)$
$= 0.001$

This means that we would expect to get nine heads in ten coin tosses of a fair coin in only about 1 in 100 samples in a very large number of trials of a fair coin. But is this the observed significance level of a test of statistical significance in which we gain nine heads in ten tosses of a coin? Or put in terms of Type I error, is this the total amount of risk we face of falsely rejecting the null hypothesis when we gain nine heads? The answer to this question is no, although it may be difficult at first to understand why. If we are willing to reject the null hypothesis based on an outcome of nine heads in ten trials, then we are, by implication, also willing to reject the null hypothesis if our outcome is ten heads in ten trials. In calculating our total risk of a Type I error, we must add together the risk of all potential outcomes that would lead us to reject the null hypothesis. This is why, when testing hypotheses, we generally do not begin with an estimate of the specific probability associated with a single outcome, but rather with the sampling distribution of probabilities of all possible outcomes.

The Binomial Distribution

To construct a probability or sampling distribution for all of the possible outcomes of ten coin tosses, we could continue to compute the number of permutations and the likelihood of any particular arrangement. However, Equation 7.3 provides us with a more direct method for calculating the probability associated with each of the potential outcomes in our sample. Equation 7.3 is generally defined as the **binomial formula,** and the distribution created from it is called the **binomial distribution.** As the name suggests, the binomial distribution is concerned with events in which there are only two possible outcomes—in our example, heads and tails.

$$P\binom{N}{r} = \frac{N!}{r!(N-r)!}p^r(1-p)^{N-r}$$ **Equation 7.3**

The binomial formula may look confusing, but most of it is familiar from material already covered in this chapter. The left-hand side of the equation represents the quantity in which we are interested—the probability of getting r successes (in our case, r heads) in a sample of N events (for us, ten tosses of a coin). The first part of the equation provides us with the number of arrangements for that number of heads. This quantity is then multiplied by $p^r(1-p)^{N-r}$, where p is the probability of a successful outcome (a head) under the null hypothesis and r is the number of successes. This formula gives us the probability associated with each arrangement. Although this part of the equation looks somewhat different from the multiplication rule we used earlier, it provides a shortcut for getting the same result, as the example below illustrates.

We have already calculated the likelihood of getting nine or ten heads in ten coin tosses if the coin is fair. To complete our sampling distribution, we need to compute probabilities associated with zero through eight heads as well. Let's begin with eight heads in ten coin tosses of an unbiased coin:

$$P\binom{10}{8} = \frac{10!}{8!(10-8)!}(0.50)^8(1-0.50)^{10-8}$$

Step 1: Calculating the number of arrangements

Working It Out

$$\binom{10}{8} = \frac{10!}{8!(10-8)!}$$

$$= \frac{10!}{8!\ 2!}$$

$$= \frac{3{,}628{,}800}{40{,}320(2)}$$

$$= 45$$

In step 1 we simply follow the same method as we did earlier in establishing the number of ways of getting eight heads in ten tosses of a coin. Our conclusion is that there are 45 different arrangements.

Step 2: Calculating the probability of any specific arrangement

Working It Out

$$\begin{aligned} p^r(1 - p)^{N-r} &= (0.50)^8(1 - 0.50)^{10-8} \\ &= (0.50)^8(0.50)^2 \\ &= (0.50)^{10} \\ &= 0.00098 \end{aligned}$$

Step 2 provides us with the likelihood of getting any particular arrangement under the assumption of the null hypothesis that the coin is fair. By the null hypothesis, p is defined as 0.50, and r is the number of successes (heads) in our example, or 8. So p^r is $(0.50)^8$, and $(1 - p)^{N-r}$ is $(1 - 0.50)^{10-8}$, or $(0.50)^2$. The outcome of this part of the equation can be simplified to $(0.50)^{10}$. This in turn is the same outcome that we would obtain using the multiplication rule, because the expression $(0.50)^{10}$ means that we multiply the quantity 0.50 by itself 10 times. Using the multiplication rule, we would have done just that.

Step 3: Combining the two outcomes

Working It Out

$$\begin{aligned} P\binom{N}{r} &= \frac{N!}{r!(N - r)!}\, p^r(1 - p)^{N-r} \\ P\binom{10}{8} &= 45(0.00098) \\ &= 0.0441 \end{aligned}$$

Combining the two parts of the equation, we find that the likelihood of tossing eight heads in ten tosses of a fair coin is about 0.044. In Table 7.5, we calculate the probabilities associated with all the potential

Table 7.5 Computation of Probability Distribution for Ten Tosses of a Fair Coin

	$\binom{N}{r} = \frac{N!}{r!\,(N-r)!}$	$\binom{N}{r} p^r (1-p)^{N-r}$
0 heads	$\frac{3,628,800}{1(10-0)!} = \frac{3,628,800}{3,628,800} = 1$	$1(0.00098) = 0.0010$
1 head	$\frac{3,628,800}{1(10-1)!} = \frac{3,628,800}{362,880} = 10$	$10(0.00098) = 0.0098$
2 heads	$\frac{3,628,800}{2(10-2)!} = \frac{3,628,800}{80,640} = 45$	$45(0.00098) = 0.0441$
3 heads	$\frac{3,628,800}{6(10-3)!} = \frac{3,628,800}{30,240} = 120$	$120(0.00098) = 0.1176$
4 heads	$\frac{3,628,800}{24(10-4)!} = \frac{3,628,800}{17,280} = 210$	$210(0.00098) = 0.2058$
5 heads	$\frac{3,628,800}{120(10-5)!} = \frac{3,628,800}{14,400} = 252$	$252(0.00098) = 0.2470$
6 heads	$\frac{3,628,800}{720(10-6)!} = \frac{3,628,800}{17,280} = 210$	$210(0.00098) = 0.2058$
7 heads	$\frac{3,628,800}{5,040(10-7)!} = \frac{3,628,800}{30,240} = 120$	$120(0.00098) = 0.1176$
8 heads	$\frac{3,628,800}{40,320(10-8)!} = \frac{3,628,800}{80,640} = 45$	$45(0.00098) = 0.0441$
9 heads	$\frac{3,628,800}{362,880(10-9)!} = \frac{3,628,800}{362,880} = 10$	$10(0.00098) = 0.0098$
10 heads	$\frac{3,628,800}{3,628,800(10-10)!} = \frac{3,628,800}{3,628,800} = 1$	$1(0.00098) = 0.0010$
		$\Sigma = 1.0$*

*The total in the last column is in fact slightly greater than 100%. This is due to rounding the numbers to the nearest decimal place in order to make the calculation more manageable.

outcomes in this binomial distribution. The resulting sampling distribution is displayed in Table 7.6.

The probability or sampling distribution for ten tosses of a fair coin illustrates how likely you are to get any particular outcome. All of the outcomes together add up to a probability of 1.[1] Put differently, there is a 100% chance that in ten tosses of a coin you will get one of these 11 potential outcomes. This is obvious, but the sampling distribution allows you to illustrate this fact. Following what our common sense tells us, it also shows that an outcome somewhere in the middle of the distribution

[1]Because of rounding error, the total for our example is actually slightly larger than 1 (see Table 7.5).

Table 7.6 Probability or Sampling Distribution for Ten Tosses of a Fair Coin

0 heads	0.001
1 head	0.010
2 heads	0.044
3 heads	0.118
4 heads	0.206
5 heads	0.247
6 heads	0.206
7 heads	0.118
8 heads	0.044
9 heads	0.010
10 heads	0.001

is most likely. If the coin is fair, then we should more often than not get about an even split of heads and tails.

The largest proportion (0.247) in the sampling distribution is found at five heads in ten tosses of a coin. As you move farther away from the center of the distribution, the likelihood of any particular result declines. The smallest probabilities are associated with gaining either all heads or no heads. Like many of the distributions that we use in statistics, this distribution is symmetrical. This means that the same probabilities are associated with outcomes on both sides.

In Chapter 6, we talked about the fact that samples vary from one another. This is what makes it so difficult to make inferences from a sample to a population. Based on a sample statistic, we can never be sure about the actual value of the population parameter. However, as illustrated in this sampling distribution, samples drawn from the same population vary in a systematic way in the long run. It is very unlikely to draw a sample with ten heads in ten tosses of a fair coin. On the other hand, it is very likely to draw a sample with four, five, or six heads in ten tosses.

Using the Binomial Distribution to Estimate the Observed Significance Level of a Test

Using the sampling distribution, we can now return to the problem of identifying the risks of error associated with rejecting the null hypothesis that the coin brought by the other volleyball team is fair. Earlier we suggested that you might want to use a 5% significance level for this test, in part because it is the standard or conventional significance level used by most criminal justice researchers. This means that in order to reject the null hypothesis you would require that the observed significance level (p) of your test (or the risk of making a Type1 error by incorrectly rejecting the null hypothesis) be less than 5% (or $p < .05$). Using this level, when would you be willing to reject the null hypothesis that the coin is fair and confront your opponents?

Applying the Binomial Distribution to Situations Where p ≠ 0.5

The examples in the text focus on applying the binomial distribution to situations where the probability of a success is equal to 0.5. There are other situations where we are interested in the probability of multiple successes (or failures), but success and failure are not equally likely. For example, many of the games of chance that a person might play at a casino are constructed in such a way that winning and losing are not equally likely—the chances of losing are greater than the chances of winning—but use of the binomial distribution would allow for calculation of the chances of winning over several plays of the game.

Consider the following more detailed example. Suppose that we have a quiz with five questions and we are interested in the probability of a student correctly guessing all of the answers on the quiz. If the only possible answers are true or false, then the probability of guessing the correct response on any single question is $p = 1/2 = 0.5$. We can then apply the binomial in the same way as we have in the previous examples to determine the probability of some number of correct answers. The following table presents the numbers of correct answers and the corresponding probabilities.

Computation of Binomial Probabilities for Five True-False Questions

NUMBER OF CORRECT ANSWERS	$\binom{N}{r} p^r (1 - p)^{N-r}$
0 correct	$\frac{5!}{0!(5-0)!} 0.5^0 (1 - 0.5)^{5-0} = 0.03125$
1 correct	$\frac{5!}{1!(5-1)!} 0.5^1 (1 - 0.5)^{5-1} = 0.15625$
2 correct	$\frac{5!}{2!(5-2)!} 0.5^2 (1 - 0.5)^{5-2} = 0.3125$
3 correct	$\frac{5!}{3!(5-3)!} 0.5^3 (1 - 0.5)^{5-3} = 0.3125$
4 correct	$\frac{5!}{4!(5-4)!} 0.5^4 (1 - 0.5)^{5-4} = 0.15625$
5 correct	$\frac{5!}{5!(5-5)!} 0.5^5 (1 - 0.5)^{5-5} = 0.03125$

Now suppose that the questions are worded as multiple-choice items and the student has to choose one answer from four possibilities. For any single question, the probability of guessing the correct answer is $p = 1/4 = 0.25$. Given that we have multiple questions, we can again calculate the probability for the number of correct responses using the binomial distribution, but we need to replace $p = 0.5$ with $p = 0.25$ in the equations to reflect the different probability of a correct answer. The following table presents the numbers of correct responses and the corresponding probabilities for the multiple-choice response set.

Computation of Binomial Probabilities for Five Multiple-Choice Questions

NUMBER OF CORRECT ANSWERS	$\binom{N}{r} p^r (1 - p)^{N-r}$
0 correct	$\frac{5!}{0!(5-0)!} 0.25^0 (1 - 0.25)^{5-0} = 0.2373$
1 correct	$\frac{5!}{1!(5-1)!} 0.25^1 (1 - 0.25)^{5-1} = 0.3955$
2 correct	$\frac{5!}{2!(5-2)!} 0.25^2 (1 - 0.25)^{5-2} = 0.2637$
3 correct	$\frac{5!}{3!(5-3)!} 0.25^3 (1 - 0.25)^{5-3} = 0.0879$
4 correct	$\frac{5!}{4!(5-4)!} 0.25^4 (1 - 0.25)^{5-4} = 0.0146$
5 correct	$\frac{5!}{5!(5-5)!} 0.25^5 (1 - 0.25)^{5-5} = 0.0010$

It is important to note that the distribution presented in the second table is no longer symmetrical, reflecting the fact that the probability of a correct response is no longer equal to the probability of an incorrect response. For the true-false questions, where the probabilities of correct and incorrect answers are the same, we see that the probabilities of zero and five correct responses are equal, the probabilities of one and four correct responses are equal, and the probabilities of two and three correct responses are equal. In contrast, when we look at the probabilities for multiple-choice questions with four possible answers, there is no such symmetry. The most likely outcome is one correct response, with a probability of 0.3955. The probability of guessing four or five correct multiple-choice answers is much lower than the probability of guessing four or five correct true-false answers. In general, the probabilities in the table show that increasing the number of possible answers makes it much more difficult for the student to correctly guess all the answers and increases the chances of getting no correct responses or only one correct response.

At first glance, you might decide to reject the null hypothesis for outcomes of zero, one, two, eight, nine, and ten heads. Each of these is below the threshold of 0.05 that we have suggested. However, at the outset we stated in our research hypothesis that we were concerned not that the coin was biased per se, but that it was biased against your team. This means that we set up our research hypotheses in such a way that we would reject the null hypothesis only if the outcomes were mostly heads. Although tossing zero, one, or two heads is just as unlikely as tossing eight, nine, or ten heads, our research hypothesis states our intention not to consider the former outcomes.

What about the risk of falsely rejecting the null hypothesis in the case of eight, nine, or ten heads? As we noted earlier, in calculating the risk of a Type I error, we must add up the probabilities associated with all the outcomes for which we would reject the null hypothesis. So, for example, if we want to know the risk of falsely rejecting the null hypothesis on the basis of eight heads in ten coin tosses, we have to add together the risks associated with eight, nine, and ten heads in ten tosses. The question we ask is, What is the risk of falsely rejecting the null hypothesis if we gain eight or more heads in ten coin tosses? The total risk, or observed significance level, would be about 0.055 (that is, 0.044 + 0.010 + 0.001), which is greater than our threshold of 0.05 for rejecting the null hypothesis. It is too large an outcome for you to confront your opponents and accuse them of cheating.

In the case of nine heads, the outcome is well below the threshold of a Type I error we have chosen. By adding together the probabilities associated with gaining nine or ten heads in ten coin tosses, we arrive at a risk of 0.011 of falsely rejecting the null hypothesis. If we decided to reject the null hypothesis that the coin is fair on the basis of an outcome of nine heads, then the observed significance value for our test would be 0.011. For ten heads, as we noted earlier, the risk of a Type I error is even lower ($p = 0.001$). Because there are no outcomes more extreme than ten heads in our distribution, we do not have to add any probabilities to it to arrive at an estimate of the risk of a Type I error.

You would take a very large risk of a Type I error if you decided in advance to reject the null hypothesis that the coin is fair based on six heads in ten tosses of a coin. Here, you would have to add the probabilities associated with six (0.206), seven (0.118), eight (0.044), nine (0.010), and ten heads (0.001).

As the coin toss example illustrates, sampling distributions play a very important role in inferential statistics. They allow us to define the observed significance level, or risk of a Type I error, we take in rejecting the null hypothesis based on a specific outcome of a test of statistical significance. Although most sampling distributions we use in statistics are considerably more difficult to develop and involve much more complex mathematical reasoning than the binomial distribution, they follow a logic similar to what we have used here. For each distribution, statisticians use probabilities to define the likelihood of gaining particular outcomes. These sampling distributions provide us with a precise method for defining risks of error in tests of statistical significance.

What you have learned here provides a basic understanding of how sampling distributions are developed from probability theory. In later chapters, we will rely on already calculated distributions. However, you should keep in mind that steps similar to those we have taken here have been used to construct these distributions.

Chapter Summary

Whereas a sample distribution is the distribution of the results of one sample, a **sampling distribution** is the distribution of outcomes of a very large number of samples, each of the same size. A sampling distribution that is derived from the laws of probability (without the need to take countless samples) may also be called a **probability distribution.** A sampling distribution allows us to define the observed significance level of a test of statistical significance, or the estimated risk of a Type I error we take in rejecting the null hypothesis based on sample statistics. To guide our decision as to whether to reject or fail to reject the null hypothesis, we compare the observed significance level with the criterion significance level set at the outset of the test of statistical significance.

By using the **multiplication rule,** we can calculate the probability of obtaining a series of results in a specific order. The number of **arrangements** is the number of different ways of obtaining the same result. The total probability of obtaining any result is the individual probability multiplied by the number of different possible arrangements.

The **binomial distribution** is the sampling distribution for events with only two possible outcomes—success or failure, heads or tails, etc. It is calculated using the **binomial formula.** When deciding whether the result achieved, or observed significance level, passes the desired threshold for rejecting the null hypothesis, it is important to remember to take a cumulative total of risk.

Key Terms

arrangements The different ways events can be ordered and yet result in a single outcome. For example, there is only one arrangement for gaining the outcome of ten heads in ten tosses of a coin. There are, however, ten different arrangements for gaining the outcome of nine heads in ten tosses of a coin.

binomial distribution The probability or sampling distribution for an event that has only two possible outcomes.

binomial formula The means of determining the probability that a given set of binomial events will occur in all its possible arrangements.

factorial The product of a number and all the positive whole numbers lower than it.

independent Describing two events when the occurrence of one does not affect the occurrence of the other.

multiplication rule The means for determining the probability that a series of events will jointly occur.

probability distribution A theoretical distribution consisting of the probabilities expected in the long run for all possible outcomes of an event.

sampling distribution A distribution of all the results of a very large number of samples, each one of the same size and drawn from the same population under the same conditions. Ordinarily, sampling distributions are derived using probability theory and are based on probability distributions.

Symbols and Formulas

$!$ Factorial

r Number of successes

N Number of trials

p The probability of a success in the binomial formula. It is also used as a symbol of the observed significance level of a test of statistical significance.

To determine the probability of events A, B, C, and D occurring jointly under the assumption of independence (the multiplication rule):

$$P(A \& B \& C \& D) = P(A) \cdot P(B) \cdot P(C) \cdot P(D)$$

To determine the number of arrangements of any combination of events:

$$\binom{N}{r} = \frac{N!}{r!(N-r)!}$$

To determine the probability of any binomial outcome occurring in all its possible arrangements (the binomial formula):

$$P\binom{N}{r} = \frac{N!}{r!(N-r)!}\, p^r(1-p)^{N-r}$$

Exercises

7.1 Calculate the probability for each of the following:

a. Two tails in two tosses of a fair coin.

b. Three heads in three tosses of a fair coin.

c. Four heads in four tosses of an unfair coin where the probability of a head is 0.75.

d. Three sixes in three rolls of a fair die.

e. Five fours in five rolls of an unfair die where the probability of a four is 0.25.

7.2 All of Kate's children are boys.

a. Intuitively, how many boys do you think Kate would have to have in succession before you would be willing to say with some certainty that, for some biological reason, she is more likely to give birth to boys than girls?

b. Now calculate the number of successive births required before you could make such a decision statistically with a 5% risk of error.

c. How many successive boys would have to be born before you would be prepared to come to this conclusion with only a 1% risk of error?

7.3 The Federal Bureau of Investigation trains sniffer dogs to find explosive material. At the end of the training, Lucy, the FBI's prize dog, is let loose in a field with four unmarked parcels, one of which contains Semtex explosives. The exercise is repeated three times, and on each occasion, Lucy successfully identifies the suspicious parcel.

a. What is the chance of an untrained dog performing such a feat? (Assume that the untrained dog would always approach one of the parcels at random.)

b. If there had been five parcels instead of four and the exercise had been carried out only twice instead of three times, would the chances of the untrained dog finding the single suspicious parcel have been greater or less?

7.4 Alex, an attorney, wishes to call eight witnesses to court for an important case. In his mind, he has categorized them into three "strong" witnesses and five "weaker" witnesses. He now wishes to make a tactical decision on the order in which to call the strong and the weaker witnesses.

For example, one of his options is

Strong Weak Weak Strong Weak Weak Weak Strong

a. In how many different sequences can he call his strong and weaker witnesses?

b. If Alex decides that one of his three strong witnesses is in fact more suited to the weaker category, how many options does he now have?

7.5 In a soccer match held at a low-security prison, the inmates beat the guards 4 to 2.

a. How many different arrangements are there for the order in which the goals were scored?

b. What would your answer be if the final score were 5 to 1?

7.6 At the end of each year, the police force chooses its "Police Officer of the Year." In spite of the fact that there are equal numbers of men and women on the force, in the last 15 years, 11 of the winners have been men and 4 have been women. Paul has been investigating whether women and men are treated differently in the police force.

a. Do these figures provide Paul with a reasonable basis to suspect that the sex of the officer is an active factor? Explain your answer.

b. Looking back further into the records, Paul discovers that for the three years before the 15-year span initially examined, a woman was chosen each time. Does this affect his conclusion? Explain your answer.

7.7 Use the binomial distribution to calculate each of the following probabilities:

a. Three heads in eight tosses of a fair coin.

b. Six tails in thirteen tosses of a fair coin.

c. Four fives in five rolls of a fair die.

d. Two ones in nine rolls of a fair die.

e. Five sixes in seven rolls of a fair die.

7.8 Tracy, a teacher, gives her class a ten-question test based on the homework she assigned the night before. She strongly suspects that Mandy, a lazy student, did not do the homework. Tracy is surprised to see that of the ten questions, Mandy answers seven correctly. What is the probability that Mandy successfully guessed seven of the ten answers to the questions if

a. The questions all required an answer of true or false?

b. The questions were all in the multiple-choice format, with students having to circle one correct answer from a list of five choices?

7.9 After a supermarket robbery, four eyewitnesses each report seeing a man with glasses fleeing from the scene. The police suspect Eddy and make up an identity parade of five men with glasses. Eddy takes his place in the parade alongside four randomly chosen stooges. Of the four eyewitnesses who are brought in, three identify Eddy and the fourth points to one of the stooges. The detective in charge decides that there is enough evidence to bring Eddy to trial.

a. The detective's superior wishes to know the probability that Eddy would have been chosen by three out of the four eyewitnesses if each witness had chosen a member of the identity parade entirely at random. What is the probability?

b. What is the probability of Eddy being chosen at random by only two of the four witnesses?

7.10 A gang of five child thieves draws straws each time before they go shoplifting. Whoever draws the short straw is the one who does the stealing. By tradition, Anton, the leader, always draws first. On the four occasions that the gang has performed this ritual, Anton has drawn the short straw three times.

a. Construct a table to illustrate the binomial distribution of Anton's possible successes and failures for each of the four draws.

b. Should he accuse his fellow gang members of rigging the draw if

i. He is willing to take a 5% risk of falsely accusing his friends?

ii. He is willing to take only a 1% risk of falsely accusing his friends?

7.11 Baron, a gambler, plays 11 rounds at a casino roulette wheel, each time placing a $100 note on either black or red.

a. Construct a table to illustrate the binomial distribution of Baron's possible successes and failures for each of the 11 rounds.

b. The casino croupiers have been told to inform the management if a client's winning streak arouses suspicion that he might be cheating. The threshold of suspicion is set at 0.005. How many successes does Baron need on 11 trials to arouse the management's suspicion?

7.12 Nicola is playing roulette on an adjacent table. On 12 successive spins of the wheel, she places a $100 note on either the first third (numbers 1–12), the second third (numbers 13–24), or the final third (numbers 25–36).

a. Construct a table to illustrate the binomial distribution of Nicola's possible successes and failures for each of the 12 spins.

b. How many times out of the 12 would Nicola need to win in order to arouse the suspicion of the casino manager that she was cheating, if the management policy is to limit the risk of falsely accusing a customer to 0.001?

7.13 A security consultant hired by store management thinks that the probability of store security detecting an incident of shoplifting is 0.1. Suppose the consultant decides to test the effectiveness of security by trying to steal an item ten different times.

a. Construct a table to illustrate the binomial distribution of possible detections for each of the ten attempted thefts.

b. Store management claims that the chances of detection are greater than 0.1. If the consultant set the threshold for detection at 0.05, how many times would she have to be detected to increase the probability of detection?

7.14 In a crime spree, Joe commits six robberies.

a. If the probability of arrest for a single robbery is 0.7, what is the probability that Joe will be arrested for three of the robberies?

b. If the probability of detection for a single robbery is 0.4, what is the probability that Joe will *not* be arrested for any of his crimes?

7.15 The arrest histories for a sample of convicted felons revealed that, with ten previous arrests, the probability of a drug arrest was 0.25. If an offender has been arrested ten times, what is the probability of

a. two drug arrests?

b. five drug arrests?

c. seven drug arrests?

Computer Exercises

The computation of binomial probabilities by hand can be quite tedious and time consuming. Spreadsheet packages typically include a function that can be inserted into a cell that will allow for the computation of a binomial probability. To compute the probability correctly, you will need to enter three pieces of information: the number of "successes" (i.e., events of interest), the total number of trials, and the probability of success on a single trial. Another item that you will need to pay attention to is whether the binomial function you are using computes a cumulative probability—the default in many spreadsheet packages. Throughout this chapter, we have not computed cumulative probabilities, but rather, what are labeled, "probability densities" in many spreadsheet packages.

Although not quite as flexible as a spreadsheet package, the computation of binomial probabilities in both SPSS and Stata is not complicated. We illustrate how the various commands work in each program.

SPSS

The computation of binomial probabilities in SPSS requires the use of the COMPUTE command that was noted in the Chapter 5 Computer Exercises. As a means of illustrating the computation of binomial probabilities in SPSS, we will reproduce the binomial probabilities listed in Table 7.5. To begin, we will create a new data set in SPSS that contains one variable: the number of success. This new variable will have values ranging from 0 to 10—enter them in order for ease of interpretation later. (After you enter these data, you should have 11 values for your new variable. For ease of illustration below, rename this variable to "successes"—without the quotation marks.)

To compute the binomial probabilities, the general form of the COMPUTE command will be:

```
COMPUTE new_var_name = PDF.BINOM(q,n,p).
EXECUTE.
```

The PDF.BINOM function will compute binomial probabilities for any given combination of successes (q), number of trials (n), and probability of success on a single trial (p). Prior to executing this command, we need to insert values for each item in the PDF.BINOM function.

- The value for "q" (the first value referenced in the parentheses) is the number of successes—this is the variable that you created with values ranging from 0 to 10. Enter the variable name here.
- The value for "n" (the second value referenced in the parentheses) is the total number of trials. For our example, enter the number 10.
- The value for "p" (the third value referenced in the parentheses) is the probability of success for a single trial. For our example, enter the value 0.5.

Assuming that you named the new variable successes, the COMPUTE command would look like:

```
COMPUTE binom_prob = PDF.BINOM(successes,10,0.5).
EXECUTE.
```

We have named the new variable "binom_prob" and once this command has been run, you should see a new variable in the second column of the SPSS data window that contains binomial probabilities. With the exception of rounding differences, these values are identical to those presented in Table 7.5.

Stata

The computation of binomial probabilities in Stata is nearly identical to that in SPSS and requires the use of the **gen** command that was noted in the Chapter 5 Computer Exercises. We walk through the same illustration of the computation of binomial probabilities as we did in SPSS, and will reproduce the binomial probabilities listed in Table 7.5 with the commands in Stata.

To create a new data set in Stata, click on the "Data Editor" button at the top center of the Stata window. You should see a spreadsheet layout, much like in SPSS. Again, we will begin by creating a new variable representing the number of success. Enter values for this new variable that range from 0 to 10. (After you enter these data, you should have 11 values for your new variable. For ease of illustration below, rename this variable to "successes"—without the quotation marks.)

To compute the binomial probabilities, the general form of the **gen** command will be:

gen new_var_name = **binomialp(n,k,p)**

Where the **binomialp** function will compute binomial probabilities for any given combination of number of trials (**n**), of number of successes (**k**), and probability of success on a single trial (**p**). Prior to executing this command, we need to insert values for each item in the **binomialp** function.

- The value for **n** (the first value referenced in the parentheses) is the total number of trials. For our example, enter the number 10.
- The value for **k** (the second value referenced in the parentheses) is the number of successes—this is the variable that you created with values ranging from 0 to 10. Enter the variable name here.
- The value for **p** (the third value referenced in the parentheses) is the probability of success for a single trial. For our example, enter the value 0.5.

Assuming that you named the new variable successes, the **gen** command would look like:

gen binom_prob = **binomialp(10,successes,0.5**)

The new variable is named "binom_prob" and once this command has been run, you should see a new variable in the second column of the data window that contains binomial probabilities. Similar to the analysis with SPSS, these values are identical to those presented in Table 7.5, with the exception of rounding differences.

Problems

1. Reproduce the tables of binomial probabilities on pages 160 and 161 in the box applying binomial probabilities when $p \neq 0.5$.
2. Verify the probabilities you calculated for any of the Exercises you worked through at the end of Chapter 7.
3. Construct a table of binomial probabilities for each of the following combinations:
   a. Number of trials = 10, probability of success = 0.2.
   b. Number of trials = 10, probability of success = 0.7.
   c. Number of trials = 15, probability of success = 0.3.
   d. Number of trials = 15, probability of success = 0.5.
   e. Number of trials = 15, probability of success = 0.8.