

Reasoning and Decision Making

12

Deductive Reasoning: Thinking Categorically

Validity and Truth in Syllogisms
How Well Can People Judge Validity?
Mental Models of Deductive Reasoning

Deductive Reasoning: Thinking Conditionally

Forms of Conditional Syllogisms
Why People Make Errors in Conditional Reasoning:
The Wason Four-Card Problem
■ **Demonstration:** Wason Four-Card Problem

Test Yourself 12.1

Inductive Reasoning: Reaching Conclusions From Evidence

The Nature of Inductive Reasoning
The Availability Heuristic
■ **Demonstration:** Which Is More Prevalent?
The Representativeness Heuristic
■ **Demonstration:** Judging Occupations
■ **Demonstration:** Description of a Person
■ **Demonstration:** Male and Female Births
The Confirmation Bias

Culture, Cognition, and Inductive Reasoning

■ **Demonstration:** Questions About Animals

Decision Making: Choosing Among Alternatives

The Utility Approach to Decisions
Decisions Can Depend on How Choices Are Presented
■ **Demonstration:** What Would You Do?
Justification in Decision Making

The Physiology of Thinking

The Prefrontal Cortex
Neuroeconomics: The Neural Basis of Decision Making

Something to Consider: Is What Is Good for You Also Good for Me?

■ **Demonstration:** A Personal Health Decision

Test Yourself 12.2

Chapter Summary

Think About It

If You Want to Know More

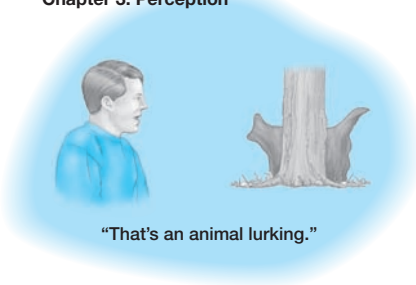
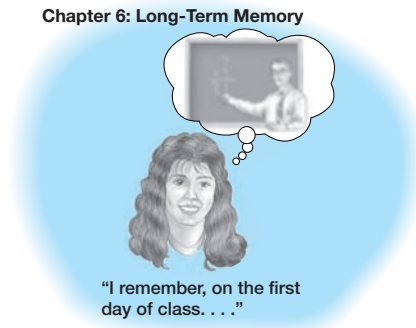
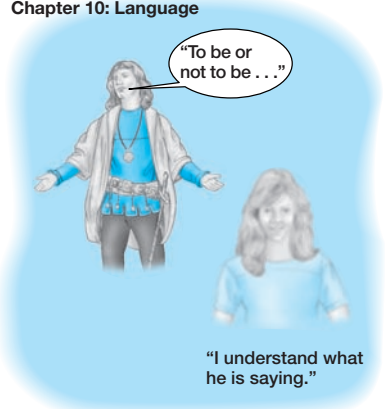
Key Terms

CogLab: Wason Selection Task; Typical Reasoning;
Risky Decisions; Decision Making: Monty Hall

Some Questions We Will Consider

- Do people reason logically, or do they make errors in reasoning? (438)
- How does reasoning operate in the discoveries made by scientists? (454)
- What kinds of reasoning “traps” do people get into when reasoning and when making decisions? (456)
- How does the fact that people sometimes feel a need to justify their decisions affect the process by which they make these decisions? (471)

What is **reasoning**? One definition is *the process of drawing conclusions* (Leighton, 2004). Another, more specific, definition is *cognitive processes by which people start with information and come to conclusions that go beyond that information* (Kurtz et al., 1999). Whatever definition we use, there is no question that reasoning is relevant to material in most of the chapters in this book. For example, in Chapter 3, on perception, we saw that perceiving an object can involve inference from incomplete information (Figure 12.1); in

COGNITION	REASONING
<p>Chapter 3: Perception</p>  <p>“That’s an animal lurking.”</p>	<p>Animal is inferred from ambiguous shapes.</p>
<p>Chapter 6: Long-Term Memory</p>  <p>“I remember, on the first day of class. . . .”</p>	<p>Memories are constructed from what we remember, plus perhaps other information.</p>
<p>Chapter 10: Language</p>  <p>“To be or not to be. . . .”</p> <p>“I understand what he is saying.”</p>	<p>Meaning is created by using knowledge we obtained earlier to help interpret Hamlet’s statement.</p>

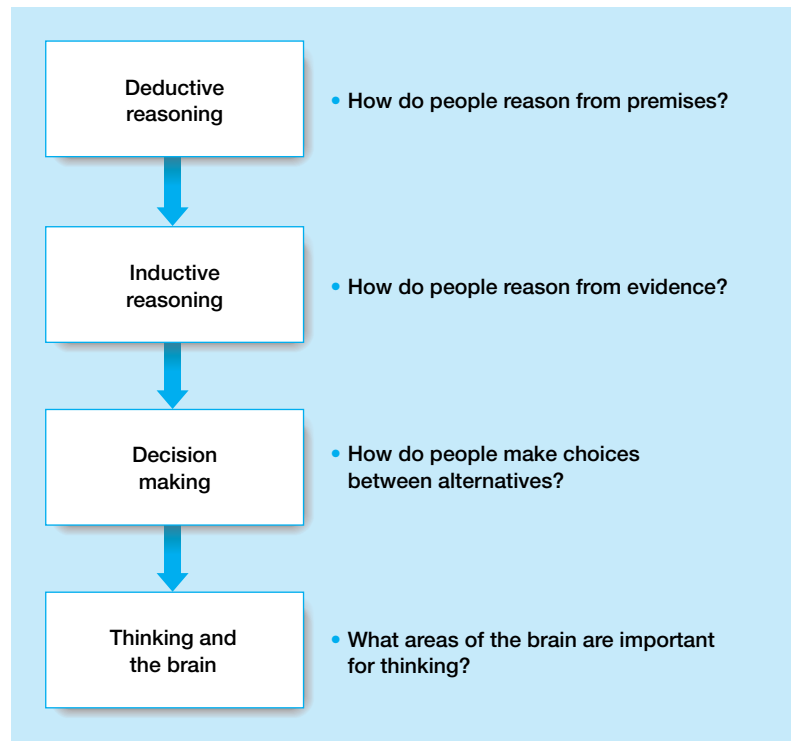
■ **Figure 12.1** Some examples of processes that fit the definition of “reasoning” that we have encountered in previous chapters.

Chapter 6, on long-term memory, we saw that our memories of events from the past are created by a process of construction, also from incomplete information; and in Chapter 10, on language, we saw how understanding one part of a story can depend on inferences based on what you know has happened before.

As these examples show, reasoning is involved in a large portion of what we study in cognitive psychology. In this chapter we are going to focus on how cognitive psychologists have studied reasoning by focusing on two specific types: deductive reasoning and inductive reasoning. We first consider **deductive reasoning**, which involves sequences of statements called **syllogisms** (Figure 12.2). For example, if we know that at least a C average is required to graduate from State U., and that Josie is graduating from State U., we can logically conclude that Josie had at least a C average.

We then consider **inductive reasoning**, in which we arrive at conclusions about what is *probably* true, based on evidence. Thus, if we know that Richard attended State U. for 4 years and that he is now the vice president of a bank, we might conclude that it is likely that he graduated. Notice, however, that in this example, we cannot say that he *definitely* graduated (maybe he never completed all the requirements, and his father, who is president of the bank, made him a vice president). Thus, we can make *definite* conclusions based on deductive reasoning and *probable* conclusions based on inductive reasoning. Studying both kinds of reasoning provides insights both about how the mind works and about everyday thinking.

■ **Figure 12.2**
Flow diagram for
this chapter.



We will also consider **decisions**, which usually involve making choices between alternatives, and which can involve both inductive and deductive reasoning. Finally, we will describe how the brain is involved in thinking.

Deductive Reasoning: Thinking Categorically

Aristotle is considered the father of deductive reasoning because he introduced the basic form of deductive reasoning called the *syllogism*. A syllogism includes two statements, called **premises**, followed by a third statement, called the **conclusion**. We will first consider **categorical syllogisms**, in which the premises and conclusion describe the relation between two categories by using statements that begin with *all*, *no*, or *some*. An example of a categorical syllogism is:

Syllogism (1) Premise 1: All birds are animals.
 Premise 2: All animals eat food.
 Conclusion: Therefore, all birds eat food.

We will refer to categorical syllogisms as simply *syllogisms* in the discussion that follows.

Validity and Truth in Syllogisms

The word “valid” is often used in everyday conversation to mean something is true or might be true. For example, saying “Susan has a valid point” could mean that what Susan is saying is true, or possibly that it should be considered further. However, when used in conjunction with categorical syllogisms, the term *validity* has a very specific meaning: A syllogism is valid when its conclusion follows *logically* from its two premises. Conversely, the syllogism is invalid when the conclusion does not follow from the premises.

Determining whether a particular syllogism is valid or invalid is not always easy because there are many different combinations of premises and conclusions. Instead of considering all of the possibilities (which would be more appropriate for a course in logic), we will consider a few examples of valid and invalid syllogisms. The syllogism below is the classic example of a valid syllogism. It is called Aristotle’s “perfect” syllogism because it was introduced by Aristotle, and it is almost immediately obvious that the conclusion follows from the two premises.

Syllogism (2) Premise 1: All A are B.
 Premise 2: All B are C.
 Conclusion: Therefore, all A are C.

The conclusion begins with “therefore” to indicate that it follows from the two premises. For brevity, we will omit the *therefore* in the rest of the syllogisms, with the understanding that in valid syllogisms the conclusion always follows from the premises.

Syllogisms can be stated abstractly, in terms of A, B, and C, as in syllogism 2, or more concretely, using meaningful terms, as in syllogism 1. Notice that syllogism 1 has

exactly the same form as syllogism 2, with A becoming “bird,” B becoming “animals,” and C becoming “eat food.”

Although syllogisms may seem like an artificial way of studying thinking, the advantage of using them is that we can construct syllogisms that are either **valid** (like syllogisms 1 and 2, in which the conclusion follows logically from the two premises), or not valid (the conclusion does not follow from the premises). These syllogisms can then be used to determine how well people evaluate validity.

A basic principle of deductive reasoning states that if the two premises of a valid syllogism are true, the syllogism’s conclusion must be true. Syllogism 1 illustrates this principle because both premises are true, and the syllogism is logically valid. But what happens when one or more of the premises are not true? Consider the following syllogism:

Syllogism (3) All birds are animals.
All animals have four legs.
All birds have four legs.

In this case, the second premise is not true, and the conclusion is not true. Nonetheless, the syllogism is still *valid* because validity depends on the *form* of the syllogism, not its *content*. This apparent conflict between “validity” and “truth” is often confusing, because “valid” is sometimes used interchangeably with “true” in everyday conversation. But remember that with regard to syllogisms, validity refers only to the *logical progression* of the premises and conclusions. As we will see, it is often difficult to separate validity and truth, and this sometimes leads to errors in reasoning. Here is an example of an invalid syllogism that contains premises and a conclusion that could be true.

Syllogism (4) All of the students are tired.
Some tired people are irritable.
Some of the students are irritable.

If you think this syllogism is valid, consider what happens when we replace “students” with “men” and “irritable” with “women.”

Syllogism (5) All of the men are tired.
Some tired people are women.
Some of the men are women.

Now it is clear that this syllogism is not valid because even though two premises could be true, it is clear that the conclusion is not. Thus, sometimes it is easy to see that a syllogism is valid, as in syllogism 1, but sometime it isn’t easy, as in syllogism 4.

How Well Can People Judge Validity?

Why are psychologists interested in whether people can judge the validity of a syllogism? One reason is to help answer the question, “Do people think logically?” One approach to this question was taken by early philosophers, who said that people’s minds work logically, and so if they do make errors in judging validity it means that they were being careless or were not paying attention.

Another approach, which has been adopted by most cognitive psychologists, is that logic is not necessarily built into the human mind, and so if people do make errors, these errors tell us something about how the mind operates. Psychologists are interested in determining what errors occur and what contributes to these errors.

People's performance in judging syllogisms has been determined using two methods: (1) *evaluation*: Present two premises and a conclusion, and ask people to indicate whether the conclusion logically follows from the premises; (2) *production*: Present two premises and ask people to indicate what conclusion logically follows from the premises, or if no conclusion logically follows. We will focus our attention on the evaluation task because researchers have studied deductive reasoning extensively using this method.

When people are tested using the evaluation method, they make errors for all syllogisms except for syllogism 2, and for some syllogisms the error rate can be as high as 70–80 percent (Gilhooly, 1988). The exact error rate depends on a number of factors, including whether the syllogism is stated abstractly (in terms of A's, B's, and C's) or in real-world terms (for example, in terms of birds and animals).

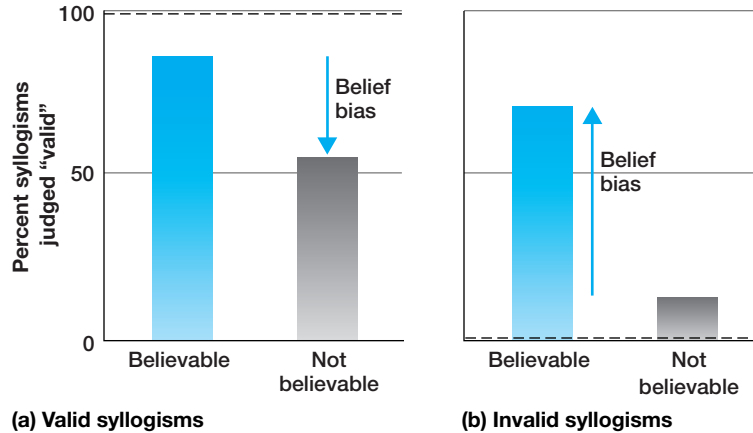
There are many reasons that people make errors in syllogisms. We will focus on two that have been widely studied, the *atmosphere effect* and the *belief bias*. The **atmosphere effect** states that the words *All*, *Some*, and *No* in the premises creates an overall “mood” or “atmosphere” that can influence the evaluation of the conclusion. According to the atmosphere effect, two *All*'s generally suggests an *All* conclusion. One or two *No*'s suggest a *No* conclusion; and one or two *Some*'s suggest a *Some* conclusion. If you look back at the syllogisms we have considered so far, you can see that application of these ideas can lead to a correct evaluation. For example, syllogism 1, which has two *All* premises and an *All* conclusion, is valid. However, for syllogism 4, which has a *Some* premise and a *Some* conclusion, the atmosphere effect would lead to the incorrect conclusion that this invalid syllogism is valid. Thus, just considering the initial terms of premises sometimes leads to correct evaluations of validity, but also leads to errors.

According to the **belief bias**, if a syllogism's conclusion is true or agrees with a person's beliefs, this increases the likelihood that the syllogism will be judged as valid. In addition, if the conclusion is false, this increases the likelihood that the syllogism will be judged as invalid. The belief bias could lead to the erroneous conclusion that syllogism 4 is valid because it seems possible that “Some of the students are irritable” could be true. Here is an example of two invalid syllogisms—both of which have the same form, but which have conclusions that differ in believability.

- Syllogism (6) No police dogs are vicious.
Some highly trained dogs are vicious.
Some police dogs are not highly trained.
- Syllogism (7) No addictive drugs are inexpensive.
Some cigarettes are inexpensive.
Some addictive drugs are not cigarettes.

When people are presented with these two syllogisms, they are more likely to indicate that syllogism 6 is invalid because they doubt the truth of the conclusion, but will

■ **Figure 12.3** Effect of the belief bias on judgments of validity. (a) Results for valid syllogisms. (b) Results for invalid syllogisms. The dashed lines indicate the percent of syllogisms that were actually valid. The arrows indicate the effect of the belief bias (Evans et al., 1983).



often indicate that syllogism 7 is valid because the conclusion is believable (Evans & Feeney, 2004).

The results of a study in which people were presented with valid and invalid syllogisms with both believable and not-believable conclusions are shown in Figure 12.3 (Evans et al., 1983). Figure 12.3a shows the results for valid syllogisms, which, if people judge validity correctly, should result in 100 percent “valid” responses, as indicated by the dashed line. It is no surprise that when the conclusion is believable, 86 percent of the judgments were “valid.” However, when the conclusion was not believable, the rates of “valid” responses dropped to 55 percent. This decrease in “valid” responses due to the unbelievability of the conclusion is an example of the belief bias.

Figure 12.3b shows the results for invalid syllogisms (Evans et al., 1983). In this case the belief bias has a huge effect. Even though the syllogism is invalid (so the percent of “valid” responses should be 0, as shown by the dashed line at the bottom of the graph), 70 percent of the responses with believable conclusions were judged to be “valid.”

The data in Figure 12.3 indicate that people are influenced by the believability of the conclusion, and therefore they do not always follow the rules of logic.

Mental Models of Deductive Reasoning

We have presented some data on deductive reasoning, but we still haven’t considered how people might go about determining whether a syllogism is valid or invalid. One way is to use diagrams, such as Venn diagrams or Euler circles (see If You Want To Know More on page 481 for more about using diagrams to determine syllogism validity). But Phillip Johnson-Laird (1999a) wonders whether people would use these diagram methods if they hadn’t been taught about them, and also points out that some of these methods don’t work for some of the more complex syllogisms.

So what could people be doing? To begin a discussion of this question, Johnson-Laird (1995) posed a problem similar to this one (try it):

On a pool table there is a black ball directly above the cue ball. The green ball is on the right side of the cue ball, and there is a red ball between them. If I move so the red ball is between me and the black ball, the cue ball is to the _____ of my line of sight.

How did you go about solving this problem? Johnson-Laird points out that the problem can be solved by applying logical rules, but that most people solve it by imagining the way the balls are arranged on the pool table (see Figure 12.21 on page 484). This idea, that people will imagine situations, is the basis of Johnson-Laird's proposal that people use mental models to solve deductive reasoning problems.

A **mental model** is a specific situation that is represented in a person's mind that can be used to help determine the validity of syllogisms in deductive reasoning. The basic principle behind mental models is that people create a model, or representation of the situation, for a reasoning problem. They generate a tentative conclusion based on this model, and then look for exceptions that might falsify the model. If they do find an exception, they modify the model. Eventually, when they can find no more exceptions, they accept their final model.


We can illustrate how this would work for a categorical syllogism by using the following example (from Johnson-Laird, 1999a):

None of the artists are beekeepers.
All of the beekeepers are chemists.
Some of the chemists are not artists.

We will imagine that we are visiting a meeting of the Artists, Beekeepers, and Chemists Society (the ABC Society, for short). We know that everyone who is eligible to be a member must be an artist, a beekeeper, or a chemist, and that they must also abide by the following rules, which correspond to the first two premises of the syllogism above:

No artists can be beekeepers.
All of the beekeepers must be chemists.

Our task is made easier because we can tell what professions people have by what they are wearing. As shown in Figure 12.4, artists are wearing berets, beekeepers are



■ **Figure 12.4** Costumes of the Artists, Beekeepers, and Chemists Society (see text).

wearing protective beekeeper's veils, and chemists are wearing molecule hats. According to the rules, no artists can be beekeepers, so people wearing berets can never wear beekeeper's veils. Also the fact that all beekeepers must be chemists means that everyone wearing a beekeeper's veil must also be wearing a molecule hat. When we meet Alice, we know she is an artist because of her beret, and we notice she is following the rule of not being a beekeeper (Figure 12.5a). Then we meet Beechem, who is wearing a combination beekeeper-molecule getup, in line with the rule that all beekeepers must be chemists (Figure 12.5b). Remember that the conclusion that has been proposed has to do with artists and chemists. Based on what we have seen so far, we can formulate our first model: *No artists are chemists*.

Remember, however, that once we have proposed our first model, we need to look for possible exceptions that would falsify this model. We do this by milling around in the crowd until we meet Cyart, who is wearing a beret and a molecule hat (Figure 12.5c). We note that he is not violating the rules, so we now know that *No artists are chemists* cannot be true, and, thinking back to Beechem, the beekeeper-chemist, we revise our model to *Some of the chemists are not artists*.

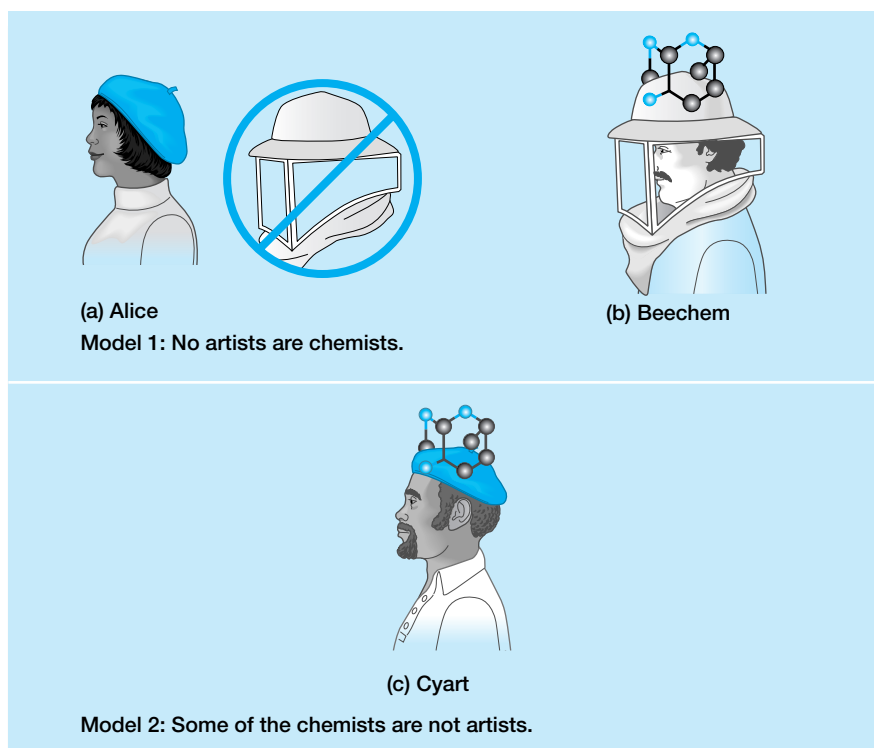


Figure 12.5 Costumes worn by (a) Alice (artist and not beekeeper); (b) Beechem (beekeeper and chemist); (c) Cyart (artist and chemist).

We keep looking for an exception to this rule, but find only Clara, who is a chemist, which is also allowed by the membership regulations. But this case does not refute our new model, and after more searching, we can't find anyone else in the room whose existence would refute this syllogism's conclusion, so we accept it. This example illustrates the basic principle behind the mental model theory: A conclusion is valid only if it cannot be refuted by any model of the premises.

The mental model theory is attractive because it can be applied without training in the rules of logic, and because it makes predictions that can be tested. For example, the theory predicts that syllogisms that require more models will be more difficult to solve, and this prediction has been confirmed in experiments (Bucciarelli & Johnson-Laird, 1999).

There are other proposals about how people might test syllogisms (see Rips, 1995, 2002), but there isn't agreement among researchers regarding the correct approach. We have presented the mental model theory because it is supported by the results of a number of experiments and because it is one of the models that is easiest to apply and explain. However, a number of challenges face researchers who are trying to determine how people evaluate syllogisms. These problems include the fact that people use a variety of different strategies in reasoning, and that some people are much better at solving syllogisms than others (Bucciarelli & Johnson-Laird, 1999). Thus, the question of how people go about solving syllogisms remains to be answered.

The categorical syllogisms we have been considering so far have premises and conclusions that begin with *All*, *Some*, or *No*. We will now consider another type of syllogism called *conditional* syllogisms, in which the first premise has the form "If . . . then. . . ."

Deductive Reasoning: Thinking Conditionally

Conditional syllogisms have two premises and a conclusion, like the ones we have been discussing, but the first premise has the form "If . . . then. . . ." This kind of deductive reasoning is common in everyday life. For example, let's say that you lent your friend Steve \$20, but he has never paid you back. Knowing Steve, you might say to yourself that you knew this would happen. Stated in the form of a syllogism, your reasoning might look like this: If I lend Steve \$20, then I won't get it back. I lent Steve \$20. Therefore, I won't get my \$20 back.

Forms of Conditional Syllogisms

There are four major types of conditional syllogisms. Each of them is listed in Table 12.1. They are presented in abstract form (p and q) on the left, and in concrete "everyday" form on the right. For conditional syllogisms, the notations p and q are typically used instead of A and B . P is the first, or "If," term, called the **antecedent**, and q is the second, or "then," term, called the **consequent**.

Table 12.1 Four Syllogisms That Begin With the Same First Premise

First premise of all syllogisms: If p , then q (abstract version), or If I study, then I'll get a good grade (concrete example)				
Syllogism	Second Premise	Conclusion	Is it Valid?	Performance
<i>Syllogism 1:</i> Affirming the antecedent	p (Abstract) I studied (Concrete)	Therefore, q Therefore, I'll get a good grade	Yes	97%
<i>Syllogism 2:</i> Denying the consequent	Not q I didn't get a good grade	Therefore, not p Therefore, I didn't study	Yes	60%
<i>Syllogism 3:</i> Affirming the consequent	q I got a good grade	Therefore, p Therefore, I studied	No	40%
<i>Syllogism 4:</i> Denying the antecedent	Not p I didn't study	Therefore, not q Therefore, I didn't get a good grade	No	40%

Source: Performance data from Evans et al. (1993).

Syllogisms 1 and 2 are valid. Syllogism 1 is called **affirming the antecedent** because the antecedent, p , or *studying*, in the second premise is affirmed. Syllogism 2 is called **denying the consequent** because the consequent, q , or *getting a good grade*, in the second premise is negated.

Syllogism 3 is called **affirming the consequent**, because q is affirmed in the second premise. This conclusion is invalid, because even though you didn't study, it is still possible that you could have received a good grade. Perhaps the exam was easy, or maybe you knew the material because it was about your job experience. If that explanation is not convincing, consider the following syllogism, which has the same form as syllogism 3, with "studying" replaced by "robin" and "good grade" replaced by "bird."

If it's a robin, then it's a bird.
It's a bird.
Therefore, it's a robin.

When stated in this way, it becomes more obvious that the affirming the consequent form of the syllogism is invalid.

Syllogism 4 is called **denying the antecedent**, because p is negated (not p) in the second premise. As in syllogism 3, you can probably think of situations that would contradict the conclusion, in which a good grade was received even though the person didn't study. Again, this fact that this syllogism is invalid becomes more obvious when restated in terms of birds and robins:

If it is a robin, then it's a bird.
It's not a robin.
Therefore, it's not a bird.

How well can people judge the validity of these syllogisms? The results of many experiments, shown in the far right column of Table 12.1, indicate that most people (close to 100 percent in most experiments) correctly judge that syllogism 1 is valid, but that performance is lower on syllogism 2, which is also valid, and 3 and 4, which are not valid. These percentages are the average results from many studies in which the letters p and q were used for the antecedent and the consequent. We have already seen that the ease of determining the validity of conditional syllogisms can be greatly affected by whether the task is stated abstractly (in terms of p 's and q 's) or concretely (studying and grades; robins and birds). In the next section we will describe a reasoning problem that has been studied using an abstract task and many different forms of real-world tasks.

Why People Make Errors in Conditional Reasoning: The Wason Four-Card Problem

CogLab

Wason
Selection
Task

If reasoning from conditional syllogisms depended only on applying rules of formal logic, then it wouldn't matter whether the syllogism was stated in terms of abstract symbols, such as p 's and q 's, or in terms of real-world items, such as studying or robins. However, we know that people are often better at judging the validity of syllogisms when real-world items are substituted for abstract symbols, and we also know that real-world items can sometimes lead to errors, as when people are influenced by the belief bias. Evidence for the effect of using real-world items in a conditional-reasoning problem is provided by a series of experiments involving the **Wason four-card problem**. Try this task in the following demonstration.



Demonstration

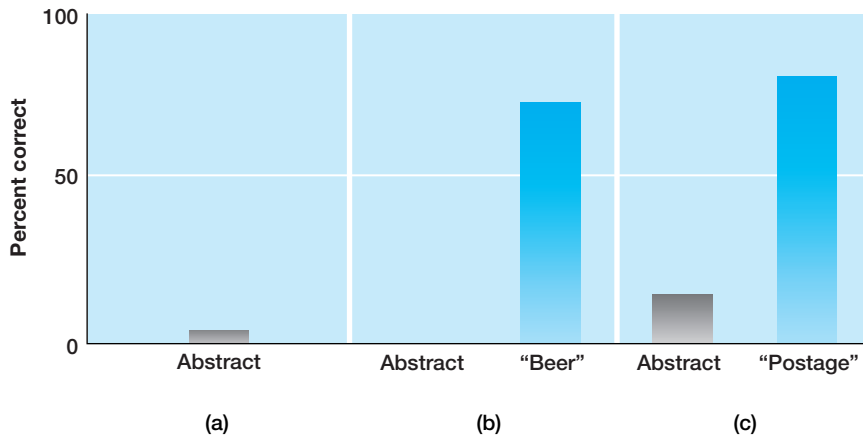
Wason Four-Card Problem

Four cards are shown in Figure 12.6. There is a letter on one side of each card and a number on the other side. Your task is to indicate the *minimum* number of cards you would need to turn over to test the following rule: *If there is a vowel on one side, then there is an even number on the other side.*



■ **Figure 12.6** The Wason four-card problem (Wason, 1966). Follow the directions in the demonstration and try this problem.

If vowel, then even number.



■ **Figure 12.7** Performance on different versions of the four-card problem. (a) Abstract version in Figure 12.6. (b) Abstract version and the beer/drinking-age version in Figure 12.8. (c) Abstract version and the postage version in Figure 12.9. Performance is better on concrete versions of the problem.

When Wason (1966) posed this task (which we will call the abstract task from now on), 53 percent of his participants indicated that the E must be turned over. This is correct because turning over the E directly tests the rule. (If there is an E, then there must be an even number, so if there is an odd number on the other side, this would prove the rule to be false.) However, another card needs to be turned over to fully test the rule. Forty-six percent of Wason’s participants indicated that in addition to the E, the 4 would need to be turned over. The problem with this answer is that if a vowel is on the other side of the card, this is consistent with the rule, but if a consonant is on the other side, turning over the 4 tells us nothing about the rule, because having a consonant on one side and a vowel on the other does not violate the rule. As shown in Figure 12.7a, only 4 percent of Wason’s participants came up with the correct answer that the card with the 7 also needs to be turned over. Turning over the 7 is important because revealing a vowel would disconfirm the rule.

The key to solving the card problem is to be aware of the **falsification principle**: *To test a rule, it is necessary to look for situations that falsify the rule.* As you can see from Table 12.2, the only two cards that have the potential to achieve this are the E and the 7. Thus, these are the only two cards that need to be turned over to test the rule.

The Role of “Regulations” in the Wason Task The Wason task has generated a great deal of research. One reason for the degree of interest in this problem is because it is a conditional-reasoning task. (Note that the problem is stated as an “If . . . then . . .” statement.) But the main reason researchers are interested in this problem is that they want to know why participants make so many errors.

One way researchers have gone about answering this question is to determine how participants perform when the problem is restated in real-world terms. In one of these

Table 12.2 Outcomes of Turning Over Each Card in the Wason Task

The Rule: If there is a vowel on one side, then there is an even number on the other side.


If turn over . . .	And the result is . . .	Then this _____ the rule.
E	Even	confirms
E	Odd	falsifies
K	Even	is irrelevant to*
K	Odd	is irrelevant to
4	Vowel	confirms
4	Consonant	is irrelevant to
7	Vowel	falsifies
7	Consonant	is irrelevant to

*This outcome of turning over the card is irrelevant because the rule does not say anything about what should be on the card if a consonant is on one side. Similar reasoning holds for all of the other irrelevant cases.

experiments Richard Griggs and James Cox (1982) used the following version of the problem:


Four cards are shown in Figure 12.8. Each card has an age on one side and the name of a beverage on the other side. Imagine you are a police officer who is applying the rule “If a person is drinking beer, then he or she must be over 19 years old.” (The participants in this experiment were from Florida, where the drinking age was 19 at the time.) Which of the cards in Figure 12.8 must be turned over to determine whether the rule is being followed?

This beer/drinking-age version of Wason’s problem is identical to the abstract version except that concrete everyday terms (beer, soda, and ages) are substituted for the letters and numbers. Griggs and Cox found that for this version of the problem, 73 percent of their participants provided the correct response: It is necessary to turn over the “beer” and the “16 years” cards. In contrast, none of their participants answered the ab-



■ **Figure 12.8** The beer/drinking-age version of the four-card problem (Griggs & Cox, 1982).

If drinking beer, then over 19 years old.



If sealed, then 5d stamp.

■ **Figure 12.9** Postage version of the four-card problem (Johnson-Laird et al., 1972).

stract task correctly (Figure 12.7b). Why is the concrete task easier than the abstract task? Apparently, being able to relate the beer task to regulations about drinking makes it easier to realize that the “16 years” card must be turned over.


The idea that knowing about regulations helps solve the Wason task has also been demonstrated using the cards in Figure 12.9 and the following instructions:

Pretend you are a postal worker sorting letters. According to postal regulations, if a letter is sealed, it must have a 5d stamp on it (*d* is pence in Great Britain). Which of the four envelopes in Figure 12.9 would you have to turn over to determine whether the rule is being obeyed?

The answer to the English postal version of the problem is that the sealed envelope and the one with a 4d stamp on it must be turned over to check the rule. When Philip Johnson-Laird and coworkers (1972) presented this problem to English participants who were familiar with an actual postal regulation similar to the one stated in the problem, performance was 81 percent correct, compared to 15 percent for the abstract task (Figure 12.7c). When Griggs and Cox (1982) tested American participants on the postal regulation task, they did not observe the large improvement in performance observed for the English participants. The reason for this result appears to be that the American participants were not familiar with postal regulations that specified different postage for opened and sealed envelopes.

The Role of “Permissions” in the Wason Task Patricia Cheng and Keith Holyoak (1985) took the Wason task a step further by proposing the concept of *pragmatic reasoning schemas*. A **pragmatic reasoning schema** is a way of thinking about cause and effect in the world that is learned as part of experiencing everyday life. One schema that people learn is the **permission schema**, which states that if a person satisfies condition A (such as being the legal age for drinking), then they get to carry out action B (being served alcohol). The permission schema for the drinking problem “If you are 19, then you got to drink beer” is something that most of the participants in this experiment had learned, so they were able to apply that schema to the card task.

This idea that people apply a real-life schema like the permissions schema to the card task makes it easier to understand the difference between the abstract version of the card task and the “drinking beer” or “postal regulation” versions. The abstract task



■ **Figure 12.10** Cholera version of the four-card problem (Cheng & Holyoak, 1985).

If entering, then cholera is listed.

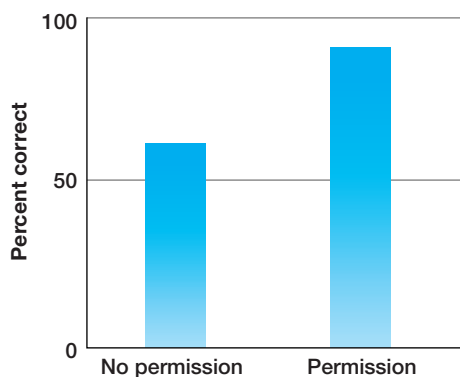
is set up so that participants approach it as a problem in which their goal is to indicate whether an abstract statement about letters and numbers is true. But the drinking-beer task and postal-regulation tasks are set up so participants approach it as problems in which their goal is to be sure that a person has permission to drink alcohol or mail a letter. Apparently, activating the permission schema helps people to focus attention on the far right card, which participants often ignore for the abstract task.

To test this idea that a permissions schema may be involved in reasoning about the card task, Cheng and Holyoak ran an experiment with two groups of participants who both saw the cards in Figure 12.10. One of the groups was read the following directions:

You are an immigration officer at the International Airport in Manila, capital of the Philippines. Among the documents you have to check is a sheet called Form H. One side of this form indicates whether the passenger is entering the country or in transit, and the other side of the form lists names of tropical diseases. You have to make sure that if the form says “Entering” on one side, that the other side includes cholera among the list of diseases.* Which of the following forms would you have to turn over to check? Indicate only those that you need to check to be sure. [*The asterisk is explained in the text that follows.]

Sixty-two percent of the participants in this group chose the correct cards (“Entering” and “Typhoid, Hepatitis”). Participants in the other group saw the same cards and heard the same instructions as the first group, but with the following changes: Instead of saying that the form listed tropical diseases, it said that the form listed *inoculations the travelers had received in the past 6 months*. In addition, the following sentence was added where indicated by the asterisk (*): *This is to ensure that entering passengers are protected against the disease.*

The changes in the instructions were calculated to achieve a very important effect: Instead of checking just to see whether the correct diseases are listed on the form, the immigration officer is checking to see whether the travelers have the inoculations necessary to *give them permission* to enter the country. These instructions were intended to activate the participants’ permissions schema, and apparently this happened, because 91 percent of the participants in this condition picked the correct cards (Figure 12.11).



■ **Figure 12.11** Results of Cheng and Holyoak’s (1985) experiment that used two versions of the cholera problem. When “permissions” are implied by the instructions, performance is better.

An Evolutionary Approach to the Four-Card Problem One of the things we have learned from our descriptions of cognitive psychology research is that one set of data can be interpreted in different ways by different investigators. We saw this in the case of the misinformation effect in Chapter 7, in which memory errors were caused by presenting misleading postevent information (MPI) after a person witnessed an event. We saw that one group of researchers explained these errors by stating that the MPI distorted existing memories (Loftus, 1993), but that other researchers offered explanations based on the effect of retroactive interference and source monitoring errors (Lindsay, 1990).

Our consideration of the Wason four-card problem now leads us to another controversy, in which different explanations have been offered to explain the results of various experiments. For example, one proposed alternative to a permissions schema is that performance on the Wason task is governed by a built-in cognitive program for detecting cheating. Let's consider the rationale behind this idea.

Leda Cosmides and John Tooby (1992) are among psychologists who have an **evolutionary perspective on cognition**. They argue that we can trace many properties of our minds to the evolutionary principles of natural selection. According to natural selection, adaptive characteristics—characteristics that help a person or animal survive to pass their genes to the next generation—will, over time, become basic characteristics of humans. Charles Darwin originally proposed this theory based on observations of physical characteristics. For example, Darwin observed that birds in a specific area had beaks with shapes adapted to enable them to obtain the food that was available.

Applying this idea to cognition, it follows that a highly adaptive feature of the mind would, through a similar evolutionary process, become a basic characteristic of the mind. One such characteristic, according to the evolutionary approach, is related to **social-exchange theory**, which states that an important aspect of human behavior is the ability for two people to cooperate in a way that is beneficial to both people. Thus, when caveman Morg lends caveman Eng his carving tool in exchange for some food that Eng has brought back from the hunt, both people benefit from the exchange.

Everything works well in social exchange as long as each person is receiving a benefit for whatever he or she is giving up. However, problems arise when someone cheats. Thus, if Morg gives up his carving tool, but Eng fails to give him the food, this does not bode well for Morg. It is essential, therefore, that people be able to detect cheating behavior so they can avoid it. According to the evolutionary approach, people who can do this will have a better chance of surviving, so “detecting cheating” has become a part of the brain's cognitive makeup.

The evolutionary approach proposes that the Wason problem can be understood in terms of cheating. Thus, people do well in the postal version of the four-card problem (Figure 12.9) because they can detect cheaters—someone who mails a letter with incorrect postage. Similarly, people do well in the cholera task (Figure 12.10) because they can detect someone who cheats by entering the country without a cholera shot.

To test the idea that cheating (and not permissions) is the important variable in the four-card problem, Cosmides and Tooby (1992) devised a number of four-card sce-

narios involving unfamiliar situations. Remember that one idea behind the permissions schema is that people perform well because they are familiar with various rules.

To create unfamiliar situations, Cosmides and Tooby created a number of experiments that took place in a hypothetical culture called the Kulwane. Participants in these experiments read a story about this culture, which led to the conditional statement, “If a man eats cassava root, then he must have a tattoo on his face.” Participants saw the following four cards: (1) eats cassava roots; (2) eats molo nuts; (3) tattoo; and (4) no tattoo. Their task was to determine which cards they needed to turn over to determine whether the conditional statement above was being adhered to. This is a situation unfamiliar to the participants, and one in which cheating could occur, because a man who eats the cassava root without a tattoo would be cheating.

Cosmides and Tooby found that participants’ performance was high on this task, even though it was unfamiliar. They also ran other experiments in which participants did better for statements that involved cheating than for other statements that could not be interpreted in this way (Cosmides, 1989; also see Gigerenzer & Hug, 1992).

However, in response to this proposal, other researchers have created scenarios that involve unfamiliar permission rules. For example, Ken Manktelow and David Over (1990) tested people using a rule that said, “If you clean up spilt blood, you must wear gloves.” Note that this is a “permission” statement that most people have not heard before, and which does not involve cheating. However, stating the problem in this way caused an increase in performance, just like many of the other examples of the Wason task that we have described.

The controversy continues among those who feel permissions are important, those who focus on cheating, and researchers who have proposed other explanations for the results of the Wason task. This is mainly because evidence has been presented for and against every proposed mechanism (Johnson-Laird, 1999b; Manktelow, 1999).

We are left with the important finding that the context within which conditional reasoning occurs makes a big difference. Stating the four-card problem in terms of familiar situations can often generate better reasoning than abstract statements or statements that people cannot relate to. However, familiarity is not always necessary for conditional reasoning (as in the tattoo problem), and situations have also been devised in which people’s performance is not improved, even in familiar situations (Evans & Feeney, 2004; Griggs, 1983; Manktelow & Evans, 1979).

Sometimes controversies such as this one are frustrating to read about because, after all, aren’t we looking for “answers”? But another way to look at controversies is that they illustrate the complexity of the human mind and the challenge facing cognitive psychologists. Remember that at the beginning of this book, we described an experiment by Donders that involved simply indicating when a light was presented or whether the light was presented on the right or on the left (see Chapter 1, page 6). We described Donders’ experiment to illustrate the basic principle that cognitive psychologists must infer the workings of the mind from behavioral observations. It is fitting, therefore, that in this, the last chapter of the book, we are now describing a task that involves

mental processes far more complex than judging whether a light has flashed, but which illustrates exactly the same principle—the workings of the mind must be inferred from behavioral observations.

We see in this controversy, about how people deal with the Wason task, how a number of different hypotheses about what is happening in the mind can be plausibly inferred from the same behavioral evidence. Perhaps, in the end, the actual mechanism will be something that has yet to be proposed, or perhaps the mind, in its complexity, has a number of different ways of approaching the Wason task, depending on the situation.



Test Yourself 12.1

1. What is deductive reasoning? What does it mean to say that the conclusion to a syllogism is “valid”? How can a conclusion be valid but not true? True but not valid?
2. How well can people judge the validity of syllogisms? Why are psychologists interested in the errors that people make in judging validity?
3. What is a categorical syllogism, and how do the atmosphere effect and belief bias influence evaluation of categorical syllogisms?
4. What is a mental model for reasoning? Explain how this model can be applied to the Artists, Beekeepers, and Chemists syllogism on page 441.
5. What is a conditional syllogism? Which of the four types of syllogisms described in the chapter are valid, which are not valid, and how well can people judge the validity of each type?
6. What is the Wason four-card problem, and what do the results of experiments that have used abstract and concrete versions of the problem indicate about the roles of (a) concreteness; (b) knowledge of regulations; and (c) permissions schemas in solving this problem?
7. How has the evolutionary approach to cognition been applied to the Wason four-card problem? What can we conclude from all of the experiments on the Wason problem?



Inductive Reasoning: Reaching Conclusions From Evidence

In deductive reasoning, premises are stated as a fact, such as “All robins are birds.” However, in inductive reasoning, premises are based on observation of one or more specific cases, and we generalize from these cases to a more general conclusion.

The Nature of Inductive Reasoning

For inductive reasoning, conclusions do not *definitely* follow from premises, and conclusions are *suggested*, with varying degrees of certainty. This is illustrated by these two inductive arguments:

Observation: All the crows I've seen in Pittsburgh are black. When I visited my brother in Washington, DC, the crows I saw there were black too.
Conclusion: I think it is a pretty good bet that all crows are black.

Observation: Here in Nashville, the sun has risen every morning.
Conclusion: The sun is going to rise in Nashville tomorrow.

Notice there is a certain logic to each argument, but the second argument is more convincing than the first. Let's consider some of the things that influence the strength of an inductive argument.

Determining the Strength of an Inductive Argument In evaluating inductive arguments, we do not consider validity, as we did for deductive arguments, but instead we decide how strong the argument is. Strong arguments are more likely to result in conclusions that are true, and weak arguments to result in conclusions that are not as likely to be true. Remember that for inductive arguments we are dealing with what is *probably* true, not what is *definitely* true.

There are a number of factors that can contribute to the strength of an inductive argument. Among them are the following:

- *Representativeness of observations*: How well do the observations about a particular category represent all of the members of that category? Clearly, the crows example suffers from a lack of representativeness because it does not consider crows from other parts of the country. If there are rare blue crows in California, then the conclusion is not true.
- *Number of observations*: The argument about the crows is made stronger by adding the Washington, DC, observations to the Pittsburgh observations. Adding more observations would strengthen it further. The conclusion about the sun rising in Nashville is extremely strong because it is supported by a very large number of observations.
- *Quality of the evidence*: Stronger evidence results in stronger conclusions. For example, although the conclusion that "The sun will rise in Nashville" is extremely strong because of the number of observations, it becomes even stronger when we consider scientific descriptions of how the earth rotates on its axis and revolves around the sun. Thus, adding the observation that "scientific measurements of the rotation of the earth indicate that every time the earth rotates the sun will appear to rise" strengthens the conclusion even further.

Bringing in scientific evidence to support an inductive argument illustrates the connection between inductive and deductive reasoning because the scientific observation about the rotation of the earth can be stated as the following *deductive* syllogism:

If the earth rotates around its axis, then the place where Nashville is located will experience sunrise for each rotation.

The earth rotates on its axis.

Therefore, Nashville will experience sunrise for each rotation.

The possibility of using scientific evidence in both inductive and deductive arguments illustrates that although it is important to distinguish between the two types of reasoning, the borderline between them can sometimes become fuzzy. We will describe the link between inductive and deductive reasoning further as we consider how inductive reasoning is used in science.

Inductive Reasoning in Science Inductive reasoning is the basic procedure used to make scientific discoveries. The goal in science is to discover something new. To achieve this, scientists often make systematic observations. These observations can include taking a poll about political attitudes, observing social behavior in a shopping mall, or doing a laboratory experiment on the Wason four-card problem. If these observations yield interesting data, these observations can be generalized to a larger population and perhaps be used to create a theory that goes beyond the specific observations.

In most scientific research, and especially in psychology, we base our conclusions on more than one observation. We test a large number of participants and may run an experiment in a number of different ways. Added participants and obtaining similar results in variations of an experiment all strengthen our conclusions. The strength of scientific conclusions also depends on the representativeness of our observations.

An obvious example of the importance of representative observations is determining attitudes by polling. Predicting the presidential election by taking a poll of students in your cognitive psychology class might tell you something about the cognitive psychology class, but would not necessarily provide an accurate prediction of the election results. To make a statement about attitudes in the United States it is necessary to elicit opinions from a representative cross-section of people in the United States. Similarly, conclusions from laboratory experiments in cognitive psychology can be safely generalized only to the population represented in the sample of people who participate in the experiment. Thus, it would be a mistake to say that the results for the beer/drinking-age version of the Wason four-card experiment will necessarily generalize to people in a society in which there are no laws regulating drinking.

Inductive reasoning is used not only to make the jump from specific observations to more general conclusions, but to create hypotheses for further experiments. An example of how inductive reasoning has been used to devise scientific experiments is provided by the way Cheng and Holyoak devised the cholera experiment we described on page 449 (Holyoak, 2003). The first step was to observe that people's performance on the Wason task improves when the task is stated in terms of receiving "permission" to do something, as is the case in the beer/drinking-age and stamp/sealed-envelope versions of the task.

Based on this observation, Cheng and Holyoak reasoned that making people more aware that permissions are involved when they are trying to solve the Wason task should improve their performance. They devised a way to test this idea by creating the scenario for the cholera problem, which included one condition in which permissions were

not emphasized and another condition in which they were. We can state the reasoning behind their thinking about this experiment (which, as we have described above, was created using inductive reasoning), by the following *deductive* syllogism:

If a permissions schema is activated, then performance on the Wason task should improve.

In this experiment, one of the groups will read a sentence that will activate a permissions schema.

Therefore, the performance of this group will improve.

As we saw when we described the “cholera” experiment, adding a sentence that was designed to activate a permissions schema did, in fact, increase performance, thereby confirming the “If . . . then . . .” premise of the conditional argument. Thus, inductive reasoning can be used to generate a hypothesis about what might be going on, and deductive reasoning sets forth the rationale of the experiment that is designed to test this hypothesis.

Inductive Reasoning in Everyday Life Inductive reasoning is used not just for creating scientific experiments, but for determining many of the choices we make in everyday life. For example, Sarah has observed, from a course she took with Professor X, that he asked a lot of questions about experimental procedures on his exams. Based on this observation, Sarah concludes that the exam she is about to take in another of Professor X’s courses will probably be similar. In another example, Sam has bought merchandise from mail-order company Y before and got good service, so he places another order based on the assumption that he will continue to get good service. Thus, anytime we make a prediction about *what will happen* based on our observations about *what has happened in the past*, we are using inductive reasoning.

It makes sense that we make predictions and choices based on past experience, especially when predictions are based on familiar situations such as studying for an exam or buying merchandise by mail. However, we make so many assumptions about the world, based on past experience, that we are using inductive reasoning constantly, often without even realizing it. For example, did you run a stress test on the chair you are sitting in to be sure it wouldn’t collapse when you sat down? Probably not. You assumed, based on your past experience with chairs, that it would not collapse. This kind of inductive reasoning is so automatic that you are not aware that any kind of “reasoning” is happening at all. Think about how time consuming it would be if you had to approach every experience as if you were having it for the first time. Inductive reasoning provides the mechanism for using past experience to guide present behavior.

When people use past experience to guide present behavior, they often use shortcuts to help them reach conclusions rapidly. After all, we don’t have the time or energy to stop and gather every bit of information that we need to be 100 percent certain that every conclusion we reach is correct. These shortcuts take the form of heuristics, which are “rules of thumb” that are likely to provide the correct answer to a problem, but which are not foolproof.

Using heuristics may sound familiar because we saw that people use heuristics to help them understand what they are seeing (Chapter 3, page 79) and what a sentence means (Chapter 10, page 373). There are a number of heuristics that people use in reasoning that often lead to the correct conclusion, but that sometimes do not. We will now describe two of these heuristics, the *availability heuristic* and the *representative heuristic*.

The Availability Heuristic

The following demonstration introduces the availability heuristic.



Demonstration

Which Is More Prevalent?


Answer these questions.

- Which is more prevalent, words that begin with the letter *r*, or words in which *r* is the third letter?
- Each item in the following list consists of two different possible causes of death. Your task is to judge, for each pair of cases, which cause of death you consider more likely for people in the United States. Think about this question in this way: Imagine you randomly picked someone in the United States. Will that person be more likely to die next year from cause A or cause B?

Cause A	Cause B	Most likely?
Homicide	Appendicitis	
Auto-train collision	Drowning	
Measles	Smallpox	
Botulism	Asthma	
Asthma	Tornado	
Appendicitis	Pregnancy	

When faced with a choice, we are often guided by what we remember from the past. The **availability heuristic** states that events that are more easily remembered are judged as being more probable than events that are less easily remembered (Tversky & Kahneman, 1973). Consider, for example, the problems we posed in the demonstration. When participants were asked to judge whether there are more words with *r* in the first position or the third, 70 percent stated that there are more words that begin with *r*, even though in reality there are three times more words that have *r* in the third position (Tversky & Kahneman, 1973; but see also Gigerenzer & Todd, 1999).

Figure 12.12 shows the results of experiments in which participants were asked to judge the relative prevalence of various causes of death (Lichtenstein et al., 1978). The height of the bars indicates the percentage of participants who picked the least likely al-



	Least	vs.	Most
1	Appendicitis	vs.	Homicide (20)
2	Auto/train collision	vs.	Drowning (5)
3	Smallpox	vs.	Measles (∞)
4	Botulism	vs.	Asthma (920)
5	Tornado	vs.	Asthma (20)
6	Pregnancy	vs.	Appendicitis (2)

■ **Figure 12.12** Likely-causes-of-death experiment results. Pairs of “causes of death” are listed below the graph, with the least likely cause on the left. The number in parentheses on the right indicates how many more times more people were actually killed by the cause on the right. The bars in the graph indicate the number of people who judged the *least likely* alternative in each pair as causing the most deaths. (Adapted from Lichtenstein et al., 1978.)

ternative. The key below the graph indicates the pairs for each bar, with the least likely cause indicated first, followed by the more likely cause. The number in parentheses indicates the relative frequency of the more likely cause compared to the less likely cause. For example, the far left bar indicates that 9 percent of the participants thought it was more likely that a person would die from appendicitis compared to homicide. However, actual mortality data indicate that 20 times more people die from homicide than from appendicitis. Thus, in this case, 9 percent of the participants made an incorrect judgment, but most of the participants made an accurate judgment regarding the most likely cause of death.

For the other causes of death, a substantial proportion of the participants misjudged which cause is more likely. In these cases, large numbers of errors were associated with causes that were publicized by the media. For example, 58 percent thought that more deaths were caused by tornados than by asthma, when in reality, 20 times more people

die from asthma than from tornados. Particularly striking is that fact that 41 percent of the participants thought that botulism caused more deaths than asthma, even though 920 times more people die of asthma.

The explanation for these misjudgments appears linked to availability. When you try to think of words that begin with *r* or that have *r* in the third position, it is much easier to think of words that begin with *r* (run, rain, real . . .) than words that have *r* in their third position (word, car, arranged . . .). When someone dies of botulism or tornados it is front-page news, whereas deaths from asthma go virtually unnoticed by the general public (Lichtenstein et al., 1978).

An experiment by Stuart McKelvie (1997) demonstrates the availability heuristic in another way. McKelvie presented lists of 26 names to participants. In the “famous men” condition, 12 of the names were famous men (Ronald Reagan, Mick Jagger) and 14 were nonfamous women. In the “famous women” condition, 12 of the names were famous women (Tina Turner, Beatrix Potter) and 14 were nonfamous men. When participants were asked to estimate whether there were more males or more females in the list they had heard, their answer was influenced by whether they had heard the male-famous list or the female-famous list. Seventy-seven percent of the participants who had heard the male-famous list stated that there were more males in their list (notice that there were actually fewer), and 81 percent of the participants who had heard the female-famous list stated that there were more females in their list. This result is consistent with the availability heuristic, because the famous names would be more easily remembered and would stand out when participants were asked to decide whether there were more male or female names.

An example of operation of the availability heuristic in everyday life is the drop in the number of people flying on commercial airlines that occurred during the year following the 9/11 terrorist attacks. The persistent images of airplanes smashing into the World Trade Center have led many people to avoid air travel in favor of driving even though, according to the National Transportation Safety Board, the fatality rate is about 500 times greater for driving than for flying in a commercial airplane. Although factors other than thinking about safety are undoubtedly involved in the drop in air travel, it is likely that the ready availability of images and descriptions of this disaster is one of the factors involved.

Another example of the availability heuristic is provided by how people’s judgments are based on correlations they observe between events. For example, you might know from past observations that when it is cloudy and there is a certain smell in the air, it is likely to rain later in the day. Being aware of correlations can be extremely useful. If you have observed that your boss is more likely to grant your requests when he or she is in a good mood, you can use this knowledge to determine the best time to ask for a raise.

Although knowledge of correlations between events can be useful, sometimes people fall into the trap of creating illusory correlations. **Illusory correlations** occur when a correlation between two events appears to exist, but in reality doesn’t exist or is much weaker than it is assumed to be.

Illusory correlations can occur when we expect two things to be related, and so we fool ourselves into thinking they are related even when they are not. These expectations

often take the form of **stereotypes**—an oversimplified generalization about a group or class of people that often focuses on the negative. Often, people’s stereotype about the characteristics of a particular group leads them to pay particular attention to behaviors associated with that stereotype, and this attention creates an illusory correlation, which reinforces the stereotype. This is related to the availability heuristic because selective attention to the stereotypical behaviors makes these behaviors more “available” (Chapman & Chapman, 1969; Hamilton, 1981).

We can appreciate how illusory correlations reinforce stereotypes by considering the stereotype that gay males are effeminate. A person who believes this stereotype might pay particular attention to effeminate gay characters on TV programs or in movies, and to situations in which they see a person who they know is gay acting effeminate. Although these observations support a correlation between being gay and being effeminate, the person has ignored the large number of cases in which gay males are not effeminate. This may be because these cases do not stand out or because the person chooses not to pay attention to them. Whatever the reason, selectively taking into account only the situations that support the person’s preconceptions can create the illusion that a correlation exists, when there may be only a weak correlation or none at all.

The Representativeness Heuristic

The representativeness heuristic is based on the idea that people often make judgments based on how much one event resembles another event.

Making Judgments Based on Resemblances The **representativeness heuristic** states that the probability that an event A comes from class B can be determined by how well A resembles the properties of class B. To put this in more concrete terms, consider the following demonstration.



Demonstration

Judging Occupations

We randomly pick one male from the population of the United States. That male, Robert, wears glasses, speaks quietly, and reads a lot. Is it more likely that Robert is a librarian or a farmer? 🐦

When Amos Tversky and Daniel Kahneman (1974) presented this question in an experiment, more people guessed that Robert was a librarian. Apparently the description of Robert as wearing glasses, speaking quietly, and reading a lot matched these people’s image of a typical librarian (see “illusory conjunctions” above). Thus, they were influenced by the representativeness heuristic into basing their judgment on how closely they think Robert’s characteristics (which correspond to “A” in our definition of the representativeness heuristic) match those of a “typical” librarian (“B”). However, in doing this they were ignoring another important source of information—the base rates

of farmers and librarians in the population. The **base rate** is the relative proportion of different classes in the population. In 1972, when this experiment was carried out, there were many more male farmers than male librarians in the United States, and this base rate leads to the conclusion that it is much more likely that Robert is a farmer (remember that he was randomly chosen from the population).

One reaction to the farmer–librarian problem might be that the participants might not have been aware of the base rates for farmers and librarians, and so didn't have the information they needed to make a correct judgment. The effect of knowing the base rate has been demonstrated by presenting participants with the following problem:

In a group of 100 people, there are 70 lawyers and 30 engineers. What is the chance that if we pick one person from the group at random that the person will be an engineer?

Participants given this problem correctly guessed that there would be a 30 percent chance of picking an engineer. However, for some participants the following description of the person who was picked was added:

Jack is a 45-year-old man. He is married and has four children. He is generally conservative, careful, and ambitious. He shows no interest in political and social issues and spends most of his free time on his many hobbies, which include home carpentry, sailing, and mathematical puzzles.

Adding this description caused participants to greatly increase their estimate of the chances that the randomly picked person (Jack, in this case) was an engineer. Apparently, when only base-rate information is available, people use that information to make their estimates. However, when any descriptive information is available, people disregard the base-rate information, and this often causes errors in reasoning.


Making Judgments Without Considering the Conjunction Rule The following demonstration illustrates another characteristic of the representativeness heuristic.



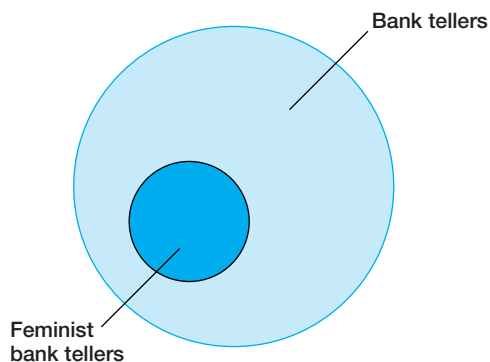
Demonstration

Description of a Person

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in antinuclear demonstrations. Which of the following alternatives is more probable?

1. Linda is a bank teller.
2. Linda is a bank teller and is active in the feminist movement. 

The correct answer to this problem is that it is more likely that Linda is a bank teller, but when Tversky and Kahneman (1983) posed this problem to their participants, 85 percent picked statement 2. It is easy to see why they did this. They were influenced by the representativeness heuristic, because the description of Linda fits people's idea of



■ **Figure 12.13** Because feminist bank tellers are a subset of bank tellers, it is always more likely that someone is a bank teller than a feminist bank teller.

a typical feminist. However, in doing this they violated the **conjunction rule**, which states that the probability of a conjunction of two events (A and B) cannot be higher than the probability of the single constituents (A alone or B alone). For example, the probability that Anne has a red Corvette cannot be greater than the probability that she has a Corvette, because the two constituents together (Corvette *and* red) define a smaller number of cars than one constituent (Corvette) alone. Similarly, there are fewer feminist bank tellers than bank tellers, so stating that Linda is a bank teller *includes* the possibility that she is a *feminist* bank teller (Figure 12.13).

People tend to violate the conjunction rule even when it is clear that they understand it. The culprit is the representativeness heuristic; in the example just cited, the participants saw Linda's characteristics as more representative of *feminist bank teller* than *bank teller*.


Incorrectly Assuming That Small Samples Are Representative People also make errors in reasoning by ignoring the importance of the size of the sample on which observations are based. The following demonstration illustrates the effect of sample size.

Demonstration

Male and Female Births

A certain town is served by two hospitals. In the larger hospital about 45 babies are born each day, and in the smaller hospital about 15 babies are born each day. As you know, about 50 percent of all babies are boys. However, the exact percentage varies from day to day. Sometimes it may be higher than 50 percent, sometimes lower.

For a period of 1 year, each hospital recorded the days on which more than 60 percent of the babies born were boys. Which hospital do you think recorded more such days?

- The larger hospital?
- The smaller hospital?
- About the same (that is, within 5 percent of each other)? 

When participants were asked this question in an experiment (Tversky & Kahneman, 1974), 22 percent picked the larger hospital, 22 percent picked the smaller hospital, and 56 percent stated that there would be no difference. The group that thought there would be no difference was presumably assuming that the birthrate for males and females in both hospitals would be representative of the overall birthrate for males and females. However, the correct answer is that there would be more days with over 60 percent male births in the small hospital.

We can understand why this result would occur by considering a statistical rule called the **law of large numbers**, which states that the larger the number of individuals that are randomly drawn from a population, the more representative the resulting group will be of the entire population. Conversely, samples of small numbers of individuals will be less representative of the population. Thus, in the hospital problem it is more likely that the percentage of boys born on any given day will be near 50 percent in the large hospital and farther from 50 percent in the small hospital. To make this conclusion clear, imagine that there is a very small hospital that records only one birth each day. Over a period of a year there will be 365 births, with about 50 percent being boys and 50 percent being girls. However, on any given day, there will be either 100 percent boys or 100 percent girls—clearly percentages that are not representative of the overall population. The problem for reasoning, however, is that people often assume representativeness holds for small samples, and this results in errors in reasoning. (See Gigerenzer & Hoffrage, 1995; Gigerenzer & Todd, 1999, for additional perspectives on how statistical thinking and heuristics operate in reasoning.)

The Confirmation Bias

One of the major roadblocks to accurate reasoning is the **confirmation bias**, our tendency to selectively look for information that conforms to our hypothesis and to overlook information that argues against it. This effect was demonstrated by Wason (1960), who presented participants with the following instructions:

You will be given three numbers which conform to a simple rule that I have in mind. . . . Your aim is to discover this rule by writing down sets of three numbers together with your reasons for your choice of them. . . . After you have written down each set, I shall tell you whether your numbers conform to the rule or not. . . . When you feel highly confident that you have discovered the rule, you are to write it down and tell me what it is. (p. 131)

After Wason presented the first set of numbers, 2, 4, and 6, the participants began creating their own sets of three numbers and receiving feedback from Wason. Note that Wason told participants only whether their numbers fit *his* rule. The participants did not find out whether *their* rationale was correct until they felt confident enough to actually announce their rule. The most common initial hypothesis was “increasing intervals of two.” Because the actual rule was “three numbers in increasing order of magnitude,” the rule “increasing intervals of two” is incorrect even though it creates sequences that satisfy Wason’s rule.

The secret to determining the correct rule is to try to create sequences that *don’t* satisfy the person’s current hypothesis, but which *do* satisfy Wason’s rule. Thus, determining that the sequence 2, 4, 5 is correct, allows us to reject our “increasing intervals of two” hypothesis and formulate a new one. The few participants whose rule was correct on their first guess followed the strategy of testing a number of hypotheses by creating sequences that were designed to *disconfirm* their current hypothesis. In contrast,

participants who didn't guess the rule correctly on their first try tended to keep creating sequences that confirmed their current hypothesis.

The confirmation bias acts like a pair of blinders—we see the world according to rules we think are correct and are never dissuaded from this view because we seek out only evidence that confirms our rule. The confirmation bias is so strong that it can affect people's reasoning by causing them to ignore relevant information. Charles Lord and coworkers (1979) demonstrated this in an experiment that tested how people's attitudes are affected by exposure to evidence that contradicts the attitudes.

By means of a questionnaire, Lord identified one group of participants in favor of capital punishment and another group against it. Each participant was then presented with descriptions of research studies on capital punishment. Some of the studies provided evidence that capital punishment had a deterrent effect on murder; others provided evidence that capital punishment had no deterrent effect. When the participants reacted to the studies, their responses reflected the attitudes they had at the beginning of the experiment. For example, an article presenting evidence that supported the deterrence effect of capital punishment was rated as “convincing” by proponents of capital punishment and “unconvincing” by those against capital punishment. This is the confirmation bias at work—people's prior beliefs caused them to focus only on information that agreed with their beliefs and to disregard information that didn't.

Culture, Cognition, and Inductive Reasoning

When we discussed categories in Chapter 8, we described research comparing the way Itza and American participants think about categories (see Chapter 8, page 315, for a description of the Itza). We saw that Itza and American participants have different ideas about which level of categories is “basic,” with Itza focusing on the level “sparrow, oak” and Americans focusing on the level “bird, tree.” We also saw that there are similarities, with both American and Itza using categories similarly for an inductive-reasoning task.

We now return to these two groups to consider some further comparisons of how they use categories in reasoning. Let's start with a demonstration.

Demonstration

Questions About Animals

The following two questions are about animals that live on an island.

Question 1:

Porcupines have a disease.

Squirrels have another disease.

Do you think all other mammals on the island have the disease of *porcupines* or the disease of *squirrels*?

Question 2:

Wolves and deer have a disease.

Wolves and coyotes have another disease.

Do you think all other mammals on the island have the disease of *wolves and deer* or of *wolves and coyotes*? 🐾

The preceding questions both involve induction, because they require reasoning from specific observations to more general conclusions, and because the answer is a “probably” answer rather than a “definitely” answer. These items are based on a model of reasoning about categories called the **similarity-coverage model** (Osherson et al., 1990). The goal of this model is to explain how people’s conceptions of different categories influence the strength of inductive arguments. We will discuss this model by posing a few more problems, and will then return to the questions from the demonstration. One principle of the model, called the **typicality principle**, is illustrated by asking which of the following is most likely to be true.

Premise: Robins have a higher potassium concentration in their blood than humans.

Conclusion 1: Therefore, all birds have a higher potassium concentration in their blood than humans.

Premise: Penguins have a higher potassium concentration in their blood than humans.

Conclusion 2: Therefore, all birds have a higher potassium concentration in their blood than humans.

According to the typicality principle, the argument with the most typical example of a category in the premise is the strongest argument. Thus, if people think robins are more typical of birds than penguins, they will pick conclusion 1 as more likely to be true.

Here is an example that illustrates another principle. Which of the following arguments are more likely to be true?

Premise: Hippopotamuses have an ulnar artery.
Hamsters have an ulnar artery.

Conclusion 1: All mammals have an ulnar artery.

Premise: Hippopotamuses have an ulnar artery.
Rhinoceroses have an ulnar artery.

Conclusion 2: All mammals have an ulnar artery.

According to the **diversity principle**, the argument with the greatest coverage of a category is stronger. Because hippopotamus and hamster taken together are more diverse than hippopotamus and rhinoceros, they have higher coverage of the category *mammal*. Therefore, according to the principle of diversity, conclusion 1 is stronger.

Now return to the demonstration “Questions About Animals,” and decide which argument illustrates typicality and which illustrates diversity. (Note that the questions in the demonstration are stated differently from the examples we have just considered, but

it is still possible to determine which principles apply. Stop and decide, before reading further.)

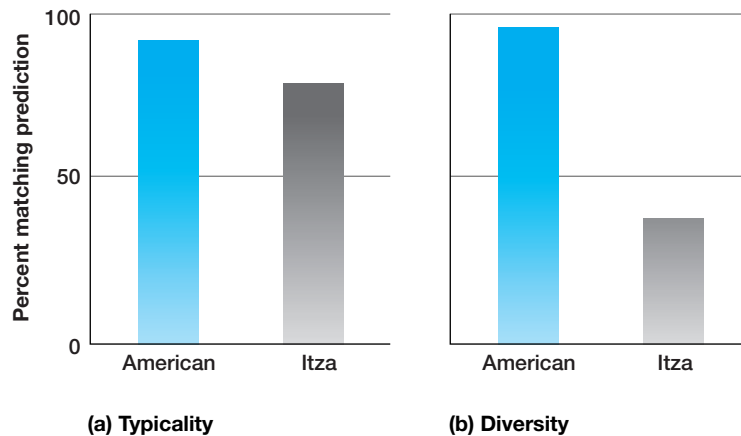
When Alejandro Lopez and coworkers (1997) presented problems like question 1 in the demonstration to Itza and U.S. participants, they found that the overall responses of both groups agreed with the principle of typicality (Figure 12.14a; picking *squirrel* rather than *porcupine*). Thus, both groups' answers reflected the animals in the premises that they thought were more typical.

Although the results for typicality were similar for both groups, the results for diversity were different. Thus, when the two groups were presented with problems like question 2 in the demonstration, almost all of the U.S. participants (96 percent) gave an answer that corresponded to the principle of diversity (picking *wolves and deer*), whereas only 38 percent of the Itza participants gave an answer that corresponded to this principle (Figure 12.14b).

What does this result mean? Why do Itza ignore diversity in making their choices? One reason is that the Itza, but not the U.S. participants, are influenced by ecological considerations. For example, one Itza participant was presented with the following problem.

Rats and pocket mice have a disease. Tapirs and squirrels have a disease. Do you think all other mammals on the island would have the disease of rats and pocket mice or the disease of tapirs and squirrels?

The person chose “rats and pocket mice,” even though the principle of diversity predicts that tapirs and squirrels, which are more different, would be chosen. The person explained that tapirs and squirrels are less likely to pass on the disease because they probably got it from another agent, such as a bat biting them, whereas rats and pocket mice are similar enough so they don't need an agent like a bat biting them to both get the disease. Thus, the Itza participant was not being illogical. Rather, she was using logic that was based on her knowledge of the animals in their environment.



■ **Figure 12.14**
(a) Percentage of judgments of inductive arguments that follow the principle of typicality for American and for Itza participants. (b) Percentage of judgments that follow the principle of diversity for American and for Itza participants (Lopez et al., 1997).

Interestingly, similar effects of knowledge on reasoning about categories have been observed for different groups of U.S. participants. For example, when asked to sort trees into different categories, taxonomists (who are schooled in scientific biology) sorted the trees by their scientific biological category, whereas park maintenance workers (who are responsible for maintaining the trees) sorted the trees based on their experience in taking care of trees. When tested, the park maintenance workers rejected the diversity principle in favor of ecologically based reasoning—just like the Itza (Medin et al., 1997)!

These cross-cultural and cross-occupational studies emphasize that in studying cognition, we can't assume that people always think in exactly the same way or use the same information to support their thinking.

Decision Making: Choosing Among Alternatives

We make decisions every day, from relatively unimportant ones (what clothes to wear, what movie to see) to those that can have great impact on our lives (what college to attend, whom to marry, what job to choose). The process of decision making can involve both inductive and deductive reasoning, so we have already considered some of the principles that apply to the study of how people make decisions.

When we discussed the availability and representativeness heuristics we used examples in which people were asked to make judgments about things like causes of death or people's occupations. As we discuss decision making, our emphasis will be on how people make judgments that involve choices between different courses of action. These choices can be concerned with personal decisions, such as deciding what school to attend or whether to fly or drive to a destination, or they can be concerned with decisions that a person might make in conjunction with their profession, such as "Which advertising campaign should my company run?" or "What is the best economic policy for the United States?" We begin by considering one of the basic properties of decision making: Decisions involve both benefits and costs.

The Utility Approach to Decisions

Much of the early theorizing on decision making was influenced by **economic utility theory**, which is based on the assumption that people are basically rational, so if they have all of the relevant information, they will make a decision which results in the maximum expected utility. **Utility** refers to outcomes that are desirable because they are in the person's best interest (Manktelow, 1999; Reber, 1995). The economists who studied decision making thought about utility in terms of monetary value, with the goal of good decision making being to make choices that resulted in the maximum monetary payoff.

One of the advantages of the utility approach is that it specifies procedures that make it possible to determine which choice would result in the highest monetary value. For example, if we know the odds of winning when playing a slot machine in a casino,

and also know the cost of playing and the size of the payoff, it is possible to determine that, in the long run, playing slot machines is a losing proposition. Thus, in terms of monetary payoff, it would be unwise to play the slots.

A Problem for the Utility Approach: People Do Not Necessarily Act to Maximize Monetary Value Even though most people realize that in the long run the casino wins, the huge popularity of casinos indicates that many people have decided to patronize casinos anyway. Observations such as this, as well as the results of many experiments, have led psychologists to conclude that people do not always make decisions that maximize their monetary outcome. This does not necessarily mean that people are irrational, but that they find value in things other than money. Thus, for some people the fun of gambling might outweigh the probable loss of some money, and, of course, there's the thrill of thinking about the possibility of "beating the odds" and being the one who hits the jackpot.

Another problem with the utility approach is that many decisions do not involve payoffs that can be calculated. As the popular television advertisement for a credit card says, *Tickets to the ball game, \$60; Hot dogs, \$10; Your team's baseball cap, \$20; Seeing the game with your son or daughter, "priceless."* Thus, utility is not always reducible to dollars and cents, but is often in the mind of the person.

The idea that utility can be in the person's mind brings up another potential problem with the utility approach. When people have to make decisions that will affect their lives, they often create mental simulations, which can sometimes be misleading (Kahneman & Tversky, 1982; Dunning & Parpal, 1989). **Mental simulations** are models that people create about what will happen following different decisions. For example, if Roberta is trying to decide whether to go to "University A" or "University B," she may imagine what it would be like attending each school. In doing this she may imagine life at University A to be intellectual and the atmosphere at University B to be more socially oriented.

Although the procedure of creating mental simulations can be useful, there is a danger that it may not lead to accurate predictions. After all, Roberta hasn't actually attended either school and has no experience with college, so she is just guessing what each school would be like. In fact, people often make inaccurate predictions about what will happen in a particular situation. For example, when people win the lottery, they may initially see nothing but positive outcomes, such as being able to quit their job, buy a new house, and finance their children's college educations. Later, however, they become aware of negative aspects, such as being hounded by other people who want a piece of the action, lack of privacy, losing friends, and worries about what's happening to their investments in the stock market. Events that people imagine will occur are often different from the events that actually do occur (T. D. Wilson et al., 2000).

People Are Often Not Good at Predicting Their Emotional Reactions T. D. Wilson and coworkers (2000) point out that even if people were able to accurately predict what would happen, they are often poor at predicting how happy or unhappy the event will cause them to feel. One of the things responsible for this lack of accuracy in predicting their emotions is called the **focusing illusion**, which occurs when people focus their attention on just one

aspect of a situation and ignore other aspects of a situation that may be important. For example, when college students were asked the questions “How happy are you?” and “How many dates did you have last month?” their answers depended on the order in which the questions were asked. When the happiness question was asked first, the correlation between the answers to the two questions was 0.12, but when the dating question was asked first, the correlation rose to 0.66. Apparently, asking the dating question first caused participants to focus on dating as being an important determinant of happiness, and so they rated themselves as happier if they had a large number of dates (Strack et al., 1988).

The focusing illusion has also been demonstrated in a study that considered people’s perceptions of how satisfied a target person would be if they lived in different locations. The participants for this study were students at two Midwestern universities (the University of Michigan and the University of Ohio) and at two California universities (the University of California at Irvine and UCLA).

There were two groups, the *self* group and the *other* group. The self group rated *themselves* on overall life satisfaction. The *other* group predicted how a hypothetical *target person* who was similar to themselves would rate their life satisfaction if they lived in California and if they lived in the Midwest. (Both groups also rated other things as well, but for our purposes we will focus on overall life satisfaction.)

The results of this study showed that there was no difference in how the California and Midwest students in the *self* group rated their own overall life satisfaction, but both California and Midwest students in the *other* group predicted that their hypothetical target person would be happier in California (Table 12.3; Schkade & Kahneman, 1998).

Why did both California and Midwest students in the *other* group predict that their target person would be happier in California? The experimenters suggest that the higher ratings for California were probably caused by the participants’ tendency to focus on the most easily observed and distinctive differences between the two locations, such as good weather and natural beauty (which people generally associate with California), and to ignore other factors, such as job prospects, academic opportunities, and financial situation. The message here is that before you decide that moving to California will make you happier, it is important to consider a wide range of outcomes from this decision—not just that the weather will be better.

All of the evidence above shows that people often make decisions that do not result in maximizing monetary value. In addition, they may not be able to accurately predict

Table 12.3 Focusing Illusion Experiment

Group	Question Asked	Answer
Self	How would you rate your life satisfaction?	No difference in self-ratings for California and Midwest students.
Other	Would a person be happier living in California or the Midwest?	Both California and Midwest students answered that the hypothetical person would have higher life satisfaction in California.

what outcome a particular decision will bring or how they will feel about the outcome when it happens. We will now consider research on decision making that has shown that people's evaluation of different choices can depend on the way these choices are presented.

Decisions Can Depend on How Choices Are Presented

Our discussion of deductive and inductive reasoning has shown that reasoning is affected by more than just the facts of the situation. This also happens in decision making, when a person's judgments are affected by the way choices are stated. For example, take the decision about whether to become a potential organ donor. Although a poll has found that 85 percent of Americans approve of organ donation, only 28 percent have granted permission by signing a donor card. This signing of the card is called the **opt-in procedure**, because it involves the person taking an active step (Johnson & Goldstein, 2003).

The low American consent rate for organ donation also occurs in other countries, such as Denmark (4 percent), the United Kingdom (27 percent), and Germany (12 percent). One thing that these countries have in common is that they all use the opt-in procedure. However, in France and Belgium the consent rate is over 99 percent. These countries use the **opt-out procedure**, in which the person is a potential organ donor unless he or she requests not to be.

Besides having important ramifications for public health (in 1995 more than 45,000 people in the United States died waiting for a suitable donor organ), the difference between opt-in and opt-out procedures has important implications for the theory of decision making, because according to the utility approach, people make their decisions based on expected utility value. According to this approach, people's decisions shouldn't depend on how the potential choices are stated. However, the opt-in versus opt-out results indicates that the procedure used to identify people's willingness to be organ donors does have an effect.

But what about when people are confronted with hypothetical situations in which they are forced to choose between two alternatives? The following demonstration provides an example of such a situation.

Demonstration

What Would You Do?

Imagine that the United States is preparing for the outbreak of an unusual Asian disease that is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:

- If Program A is adopted, 200 people will be saved.
- If Program B is adopted, there is a 1/3 probability that 600 people will be saved, and a 2/3 probability that no people will be saved.

CogLab

Risky
Decisions

CogLab

Decision
Making

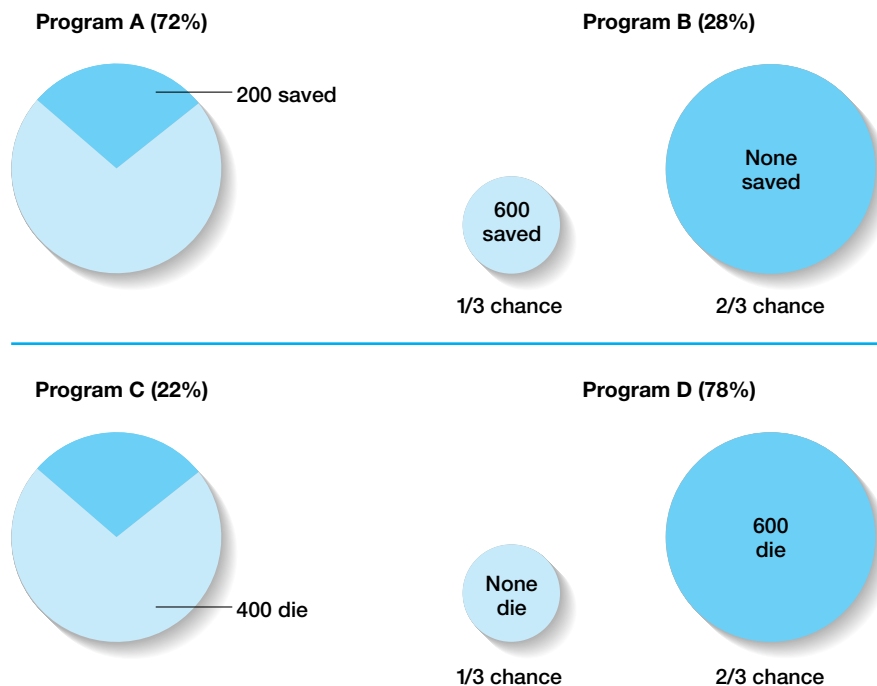
Which of the two programs would you favor?

Now consider the following additional proposals for combating the same disease:

- If Program C is adopted, 400 people will die.
- If Program D is adopted, there is a 1/3 probability that nobody will die, and a 2/3 probability that 600 people will die.

Which of these two programs would you pick? 🐦

For the first pair of proposals, Program A was chosen by 72 percent of the students in an experiment by Tversky and Kahneman (1981) and the rest picked Program B (Figure 12.15). The choice of Program A represents a **risk-aversion strategy**. The idea of saving 200 lives with certainty is more attractive than the risk that no one will be saved. However, when Tversky and Kahneman presented the descriptions of Programs C and D to another group of students, 22 percent picked Program C and 78 percent picked Program D. This represents a **risk-taking strategy**. The certain death of 400 people is less acceptable than a 2 in 3 chance that 600 people will die.



■ **Figure 12.15** How framing affects decision making. These pie charts diagram the conditions set forth for Programs A, B, C, and D in the text. Note that the number of deaths and probabilities for programs A and B are exactly the same as for programs C and D. The percentages indicate the percentage of participants who picked each program when given choices between A and B or between C and D (Tversky & Kahneman, 1981).

Tversky and Kahneman concluded that, in general, when a choice is framed in terms of gains (as in the first problem, which is stated in terms of saving lives) people use a risk-aversion strategy, and when a choice is framed in terms of losses (as in the second problem, which is stated in terms of losing lives), people use a risk-taking strategy.

But if we look at the four programs closely, we can see that they are identical pairs (Figure 12.15). Programs A and C both result in 200 people living and 400 people dying. Yet 72 percent of the participants picked program A and only 22 percent picked program C. A similar situation occurs if we compare programs B and D. Both lead to the same number of deaths, yet one was picked by 28 percent of the participants and the other by 78 percent. These results illustrate the **framing effect**—decisions are influenced by how a decision is stated, or *framed*.

One reason people's decisions are affected by framing is that the way a problem is stated can highlight some feature of the situation (for example, that people will die) while deemphasizing other features (Kahneman, 2003). It should not be a surprise that the way a choice is stated can influence cognitive processes, because this is similar to what happens when the way a syllogism is stated influences a person's ability to determine whether the syllogism is valid (page 438). We also saw, in the chapter on problem solving, that the way a problem is stated can influence our ability to solve the problem (page 408).

Justification in Decision Making

To end our consideration of decision making, we will consider yet another factor that influences how people make decisions. This factor is the need to justify the decision. We can illustrate this by considering an experiment by Tversky and Eldar Shafir (1992), in which they presented the following problem to two groups of students. The “pass” group read the words that are underlined, and the “fail” group had these words replaced by the words in italics (which the pass group did not see).

Imagine that you have just taken a tough qualifying examination. It is the end of the semester, you feel tired and run-down, and you find out that you passed the exam/ [*failed the exam. You will have to take it again in a couple of months—after the Christmas holidays*]. You now have the opportunity to buy a very attractive 5-day Christmas vacation package to Hawaii at an exceptionally low price. The special offer expires tomorrow. Would you

- buy the vacation package?
- not buy the vacation package?
- pay a \$5 nonrefundable fee in order to retain the rights to buy the vacation package at the same exceptional price the day after tomorrow?

The results for the two groups are shown in the columns headed “passed” and “failed” in Table 12.4. Notice that there is no difference between the two groups. Fifty-four percent of the participants in the “pass” group opt to buy the vacation package, and 57 percent of the participants in the “fail” group opt for the package. The interesting

Table 12.4 Choice Behavior and Knowledge of Exam Outcome

	Passed Test	Failed Test	Test Results Not Available for 2 Days
Buy vacation package	54%	57%	32%
Don't buy	16	12	7
\$5 to keep open option to buy later	30	31	61

result happened when a third group received the same problem as above, except these participants were told that the outcome of the exam wouldn't be available for 2 more days. Notice that only 32 percent of these participants opted for the package and that 61 percent decided they would pay the \$5 so they could put off making the decision until they knew whether they had passed or failed the exam.

Apparently what happened in this experiment is that 61 percent of the participants in the “no result” group did not want to make a decision about the trip until they found out whether they passed or failed, even though the results for the other two groups indicates that passing or failing actually made no difference in the decision about the vacation packages.

To explain this result, Tversky and Shafir suggest that once the students know the outcome they can then assign a reason for deciding to buy the vacation. Participants who passed could see the vacation as a reward; participants who failed could see the vacation as a consolation that would provide time to recuperate before taking the exam again.

Although there are other possible interpretations for these results, there is a great deal of other evidence that the decision-making process often includes looking for justification so the person can state a rationale for his or her decision. This is why doctors may carry out medical tests that might not lead to different treatments but that provide additional evidence for the treatment they have recommended, thereby making it easier to justify the treatment to themselves, their patients, and, if necessary, to the courts (Tversky & Shafir, 1992).

The Physiology of Thinking

In this section we will consider the types of thinking we have discussed in this chapter and the previous one. We begin by asking the question, “How is the brain involved in problem solving, reasoning, and making decisions?” Because all of these forms of thinking involve a number of different cognitive capacities—including perception, memory, and the ability to focus and maintain attention—it isn't surprising that a number of different areas of the brain are involved. However, we will focus on one area in particular, the prefrontal cortex, because it plays such a major role in thinking.


The Prefrontal Cortex

The prefrontal cortex (PFC; Figure 12.16) is activated by stimuli from all of the senses, by the retrieval of memories and the anticipation of future events, and can be affected by a person's emotional state (Wallis et al., 2001). It is not surprising, therefore, that the PFC plays a central role in determining complex behaviors that are involved in thinking.

One of the earliest reports of the effect of frontal lobe damage on functioning was the case of a young homemaker who had a tumor in her frontal lobe that made it impossible for her to plan a family meal, even though she was capable of cooking the individual dishes (Penfield & Evans, 1935). Results such as this led to the conclusion that the PFC plays an important role planning future activities (Owen et al., 1990).

The prefrontal cortex has been linked to problem solving in a number of ways. Damage to the PFC interferes with people's ability to act with flexibility, a key requirement for solving problems. One symptom of PFC damage is a behavior called **perseveration**, in which patients have difficulty in switching from one pattern of behavior to another (Hauser, 1999; Munakata et al., 2003). For example, patients with damage to the PFC have difficulty when the rules change in a card-sorting task. Thus, if they begin by successfully separating out the blue cards from a pack, they continue picking the blue cards even after the experimenter tells them to shift to separating out the brown cards. Clearly, perseveration would play havoc with attempts to solve complex problems for which it is necessary to consider one possible solution and then shift to another possibility if the first one doesn't work.

Because damage to the PFC results in perseveration and poor planning ability, it is not surprising that PFC damage decreases performance on tasks such as the Tower of London problem (a task similar to the Tower of Hanoi problem that involves moving colored beads between two vertical rods; Carlin et al., 2000; Owen et al., 1990),



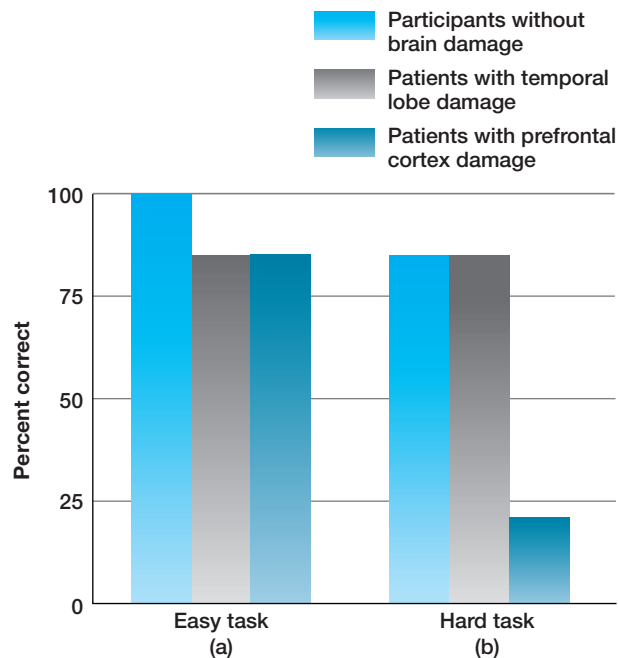
■ **Figure 12.16**
Brain showing location of the prefrontal cortex (PFC).

the Tower of Hanoi problem (Morris et al., 1997), and the Luchins water-jug problem (Colvin et al., 2001). Brain imaging has also shown that problem solving activates the PFC in normal participants (Rowe et al., 2001).

Other research has shown that the PFC is important for a number of cognitive tasks involving planning, reasoning, and making connections among different parts of a problem or a story. For example, when Tiziana Zalla and coworkers (2002) tested patients with PFC damage, she found that these patients were able to understand individual words and could identify events described in stories. However, they were unable to follow the order of events in the story or to make inferences that connected different parts of the story.

There is also a large amount of evidence that the PFC is important for reasoning. This has been demonstrated by presenting a deductive-reasoning task to people with PFC damage. Participants were presented with relationships such as the following: *Sam is taller than Nate; Nate is taller than Roger*, and their task was to arrange the names in order of the people's heights. When James Waltz and coworkers (1999) presented these tasks to patients with PFC damage and control groups of participants without brain damage and patients with temporal lobe damage, they found that all of these groups did well when the task was easy, like the previous one about Sam, Nate, and Roger (Figure 12.17a). However, when the task was made more difficult by scrambling the order of presentation (example: *Beth is taller than Tina; Amy is taller than Beth*), then the people without brain damage and the patients with temporal lobe damage still did well, but the PFC patients performed poorly (Figure 12.17b). This result confirms the conclusion of

■ **Figure 12.17** Effect of damage to the PFC on performance on a reasoning task. Participants without brain damage, participants with temporal lobe damage, and participants with PFC damage can all solve the easy task (left bars), but the PFC group's performance drops to a low level when the task is made more difficult (Waltz et al., 1999).



brain imaging studies, which show that as reasoning problems become more complex, reasoning activates larger areas of the PFC (Kroger et al., 2002). (Also see “If You Want to Know More: Neurons That Respond to Abstract Rules,” on page 483.)

Neuroeconomics: The Neural Basis of Decision Making

A new approach to studying decision making, called **neuroeconomics**, combines approaches from the fields of psychology, neuroscience, and economics (Lee, 2006; Sanfey et al., 2006). One outcome of this approach has been research that has identified areas of the brain that are activated as people make decisions while playing economic games. This research shows that decisions are often influenced by emotions, and that these emotions are associated with activity in specific areas of the brain.

To illustrate the neuroeconomic approach, we will describe an experiment by Alan Sanfey and coworkers (2003) in which people’s brain activity was measured as they played the ultimatum game. The **ultimatum game** is very simple. Two people play: One is designated as the *proposer* and one as the *responder*. The proposer is given a sum of money, say \$10, and makes an offer to the responder as to how this money should be split between them. If the responder accepts the offer, then the money is split according to the proposal. If the responder rejects the offer, neither player receives anything. Either way, the game is over after the responder makes his or her decision.

According to utility theory, the responder should accept the proposer’s offer, no matter what it is. This is the rational response, because if you accept the offer you get something, but if you refuse, you get nothing (remember that the game is only one-trial long, so there is no second chance).

In Sanfey’s experiment, participants played 20 separate games as responder: 10 with 10 different human partners and 10 with a computer partner. The offers made by both the human and computer partners were determined by the experimenters, with some being “fair” (evenly split, so the responder received \$5) and some “unfair” (the responder received \$1, \$2, or \$3). The results of responders’ interactions with their human partners (gray bars in Figure 12.18) match the results of other research on the ultimatum game—all responders accept an offer of \$5, most accept the \$3 offer, and half or more reject the \$1 or \$2 offers.

Why do people reject low offers? When Sanfey and coworkers asked participants, many explained that they were angry because they felt the offers were unfair. Consistent with this explanation, when participants received exactly the same offers from their computer partner, more accepted “unfair” proposals (blue bars in Figure 12.18). Apparently, people are less likely to get angry with an unfair computer than with an unfair person.

In addition to testing people’s behavior, Sanfey and coworkers measured brain activity in the responders as they were making their decisions. The results showed that the right anterior insula, an area located deep within the brain between the parietal and temporal lobes, was activated about three times more strongly when responders rejected an offer than when they accepted it (Figure 12.19a). Also, participants with higher activation to unfair offers rejected a higher proportion of the offers. The fact that the insula responded during rejection is not surprising when we consider that this area of

Figure 12.18 Behavioral results of Sanfey and coworkers' (2003) experiment, showing responders' acceptance rates in response to different offers made by human partners and computer partners.

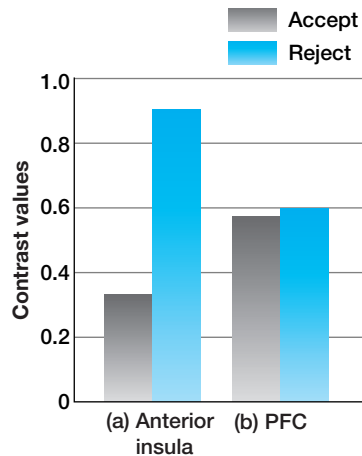
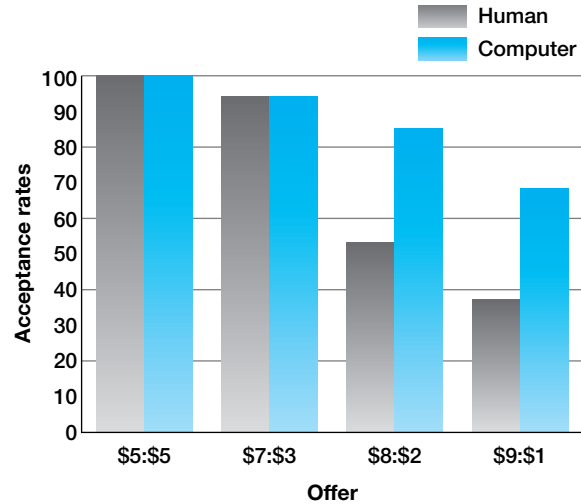


Figure 12.19 Response of the insula and PFC to “fair” and “unfair” offers (Sanfey et al., 2006).

the brain is connected with negative emotional states, including pain, distress, hunger, anger, and disgust.

What about the prefrontal cortex, which plays such a large role in complex cognitive behaviors? The PFC is also activated by the decision task, but this activation is the same for offers that are rejected and offers that are accepted (Figure 12.19b). Sanfey hypothesizes that the function of the PFC may be to deal with the cognitive demands of the task, which involves the goal of accumulating as much money as possible. Looked at in this way, each of these brain areas represents a different goal of the ultimatum game—the emotional goal of resenting unfairness is handled by the anterior insula, and the cognitive goal of accumulating money is handled by the PFC.

The results of this experiment support the idea that it is important to take emotional factors into account when considering decision making. It also illustrates the value of combining both physiological and behavioral approaches to the study of decision making.

Something to Consider


Is What Is Good for You Also Good for Me?

When we discussed how framing affects decision making, we saw that people’s decisions regarding programs to deal with the outbreak of a hypothetical Asian disease depended on how the problem was stated (pages 469–470). We now pose a similar type of medical problem, but in a more personal way, because the hypothetical decision you are asked to make could affect you personally.

Demonstration

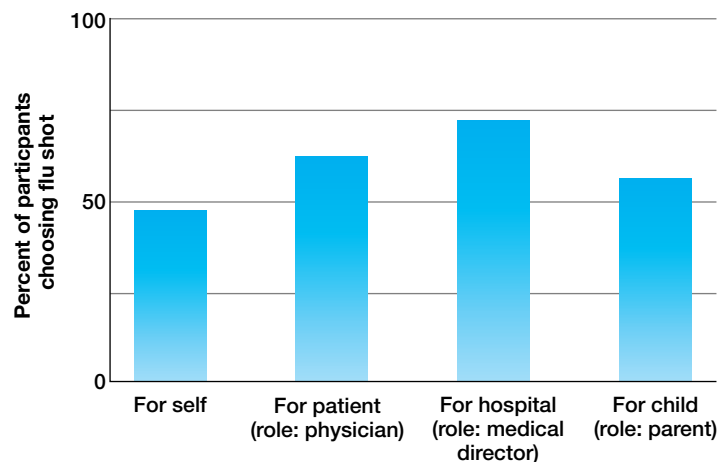
A Personal Health Decision

Imagine that there will be a deadly flu going around your area next winter. Your doctor says that you have a 10 percent chance (10 out of 100) of dying from this flu. A new flu vaccine has been developed and tested. If administered, the vaccine will prevent you from catching the deadly flu. However, there is one serious risk involved: The vaccine is made from a somewhat weaker type of flu virus, so there is a 5 percent risk (5 out of 100) that the vaccine could kill you. Considering this information, decide between the following two alternatives:

- I will not take the vaccine, and I accept the 10 percent chance of dying from this flu.
- I will take the vaccine, and I accept the 5 percent chance of dying from the weaker flu in the vaccine. (Adapted from Zikmund-Fisher et al., 2006) 

When Brian Zikmund-Fisher and coworkers (2006) gave this choice to their participants, 48 percent said they would take the vaccine. This is an interesting result, because it means that 52 percent of the participants decided to do nothing, even though statistically this doubled their chances of dying.

This result is an example of the **omission bias**—the tendency to do nothing to avoid having to make a decision that could be interpreted as causing harm. However, Zikmund-Fisher’s experiment asked participants not only to imagine that they were making a decision for themselves, as in the demonstration, but to make the decision while imagining themselves in the following three roles: (1) as a physician recommending a treatment for a patient; (2) as a hospital medical director setting treatment guidelines for all patients in the hospital; and (3) as a parent of a child who might receive the treatment. The results of this experiment, shown in Figure 12.20, indicate that people



■ **Figure 12.20** Effect of imagined decision-making role on the willingness to choose the flu vaccine. (Data from Zikmund-Fisher et al., 2006).

are more likely to recommend that others receive the shot than they are to choose the shot for themselves.

Apparently, the decision a person makes can be influenced by the person or group for whom they are making the decision. But why does this occur? Zikmund-Fisher and coworkers propose that when making decisions for others, people take into account the possibility that they will be held responsible if something bad happens. Looked at from this point of view, it is easy to understand why a medical director would be prone to recommend that hospital patients receive the vaccine, because it is easy to justify a decision that maximizes survival chances for a group of people.

The most important implication of these results may be what it suggests about how physicians should present choices to their patients. Physicians often feel that they should simply present the information and let their patients deal with making the decision. But perhaps physicians should be sensitive to some of the emotional factors facing patients who are being asked to make decisions about their own treatment. Zikmund-Fisher and coworkers suggest that physicians should consider asking patients to “re-frame” their decision by thinking about it as if it were a decision they were making for someone else. The idea behind doing this would be to help the patient gain a better understanding of the trade-offs they face.



Test Yourself 12.2

1. What is inductive reasoning, and how is it different from deductive reasoning?
2. How is inductive reasoning involved in the practice of science? How do inductive and deductive reasoning work together in scientific research?
3. How is inductive reasoning involved in everyday experience?
4. How do the following cause errors in reasoning: availability heuristic; illusory correlations; representativeness heuristic; confirmation bias?
5. How can failure to take into account base rates and small sample sizes cause errors in reasoning?
6. What is the cross-cultural evidence regarding how U.S. and Itza participants use typicality and diversity when making inferences about categories? How are the behaviors of these two groups similar? Different? What is an explanation for the differences?
7. What is the utility approach to decisions? What are some problems with the utility approach? As you consider this, take into account mental simulations and the focusing illusion.
8. How does the way a problem is stated and the need to justify decisions affect the decisions people make?
9. How is the prefrontal cortex involved in problem solving and reasoning? Cite evidence from both neuropsychology and brain imaging experiments to support your answer.
10. What is neuroeconomics? Describe Sanfey and coworkers' (2003) experiment, and indicate what it adds to our understanding of decision making.

11. How are people's decisions about treatment options influenced by the person or group for whom they are making the decision?

Chapter Summary

1. Reasoning is a cognitive process by which people start with information and come to conclusions that go beyond that information. Deductive reasoning involves syllogisms and can result in definite conclusions. Inductive reasoning is based on evidence and results in conclusions that are *probably* true.
2. Categorical syllogisms have two premises and a conclusion that describe the relation between two categories by using statements that begin with *All*, *No*, or *Some*.
3. A syllogism is valid if its conclusion follows logically from its premises. If the premises of a valid syllogism are true, then the conclusion must be true. People are able to correctly judge the validity of Aristotle's "perfect" syllogism, but make errors in all other forms of categorical syllogisms. Two reasons for errors are the atmosphere effect and the belief bias.
4. One way to determine whether a categorical syllogism is valid or invalid is to use diagrams. Another method is to create mental models representing the premises.
5. Conditional syllogisms have two premises and a conclusion, like categorical syllogisms, but the first premise has the form "If . . . then . . ." The four basic types of conditional syllogism are (a) affirming the antecedent and (b) denying the consequent (both valid), and (c) affirming the consequent and (d) denying the antecedent (both invalid).
6. The Wason four-card problem has been used to study how people think when evaluating conditional syllogisms. People make errors in the abstract version, but perform better for versions of the problem that are restated in real-world terms, such as the "drinking age" and "postal" versions. The key to solving the problem is to apply the falsification principle.
7. Based on experiments using different versions of the Wason problem, a number of mechanisms have been proposed to explain people's performance. These mechanisms include using permission schemas, and the evolutionary approach, which explains performance in terms of social-exchange theory. Many experiments have provided evidence for and against these explanations, leaving the controversy about how to explain the Wason problem still unresolved.
8. In inductive reasoning, conclusions follow not from logically constructed syllogisms, but from evidence. Conclusions are *suggested* in inductive reasoning, with varying degrees of certainty. The strength of an inductive argument depends on the representativeness, number, and quality of observations on which the argument is based.
9. Inductive reasoning is one of the basic mechanisms for developing scientific theories and evaluating scientific evidence. It also plays a major role in everyday life because we often make predictions about what we think will happen based on our observations about what has happened in the past.

10. The availability heuristic states that events that are more easily remembered are judged as being more probable than events that are less easily remembered. This heuristic can sometimes lead to correct judgments, and sometimes not. Errors due to the availability heuristic have been demonstrated by having people estimate the relative prevalence of various causes of death.
11. Illusory correlations and stereotypes, which can lead to incorrect conclusions about relationships between things, are related to the availability heuristic, because they draw attention to specific relationships and therefore make them more “available.”
12. The representativeness heuristic is based on the idea that people often make judgments based on how much one event resembles another event. Errors due to this heuristic have been demonstrated by asking participants to judge a person’s occupation based on descriptive information. Errors occur when the representativeness heuristic leads people to ignore base-rate information. In other situations, judgment errors occur when people ignore the conjunction rule and the law of large numbers.
13. The confirmation bias is the tendency to selectively look for information that conforms to a hypothesis and to overlook information that argues against it. Operation of this bias was demonstrated by Wason’s number-sequence task. This bias also operates in real life when people’s attitudes influence the way they evaluate evidence.
14. Cultural differences in inductive reasoning have been demonstrated by comparing how Itza and American participants apply the similarity-coverage model and the associated typicality and diversity principles to their conceptions of different categories. Experiments have demonstrated that these groups apply the typicality principle similarly but differ in their application of the diversity principle.
15. The utility approach to decision making states that people are basically rational, so when they have all of the relevant information, they will make a decision that results in outcomes that are in their best interest. Evidence that people do not always act in accordance with this approach includes failure to act in a way to maximize monetary value, errors caused by application of mental simulations, and errors caused by the focusing illusion.
16. Decisions can depend on how choices are presented, or *framed*. For example, when a choice is framed in terms of gains, people use a risk-aversion strategy, and when a choice is framed in terms of losses, people use a risk-taking strategy. Decision making is also influenced by people’s tendency to want to justify their decision and state a rationale for the decision.
17. The prefrontal cortex (PFC) is one of the major areas of the brain involved in thinking. Damage to the PFC can cause perseveration and poor planning ability, which results in poor performance on everyday tasks, problems such as the Tower of Hanoi and water-jug problem, and some basic problems in reasoning.
18. Neuroeconomics studies decision making by combining approaches from psychology, neuroscience, and economics. The results of a neuroeconomics experiment using the ultimatum game have shown that people’s emotions can interfere with their ability to make rational decisions. Brain imaging indicates that the anterior insula is

associated with the emotions that occur during the ultimatum game and also suggests that the PFC may be involved in the cognitive demands of the task.

19. An experiment that involved asking people to make a risky decision about being vaccinated against a deadly disease has shown that people are more likely to recommend that others receive the vaccination than they are to choose to receive the vaccination themselves. This result has implications about how physicians talk about treatment options with their patients.



Think About It

1. Astrology is popular with many people because they perceive a close connection between astrological predictions and events in their lives. Explain factors that might lead to this perception, even if a close connection does not, in fact, exist.
2. Think about a decision you have made recently. It can be a minor one, such as deciding which restaurant to go to on Saturday evening, or a more major one, such as picking an apartment or deciding which college to attend. Analyze this decision, taking into account the processes you went through to arrive at it, and how you justified it in your mind as being a good decision.
3. Create deductive syllogisms and inductive arguments that apply to the decision you analyzed in the previous question.
4. Johanna has a reputation for being extremely good at justifying her behavior by a process that is often called “rationalization.” For example, she justifies the fact that she eats anything she wants by saying “Ten years ago this food was supposed to be bad for you, and now they are saying it may even have some beneficial effects, so what’s the point of listening to the so-called health experts?” or “That movie actor who was really into red meat lived to be 95.” Analyze Johanna’s arguments by stating them as inductive or deductive arguments or, better yet, do that for one of your own rationalizations.
5. From watching the news or reading the paper, what can you conclude about how the availability heuristic can influence our conceptions of the nature of the lives of different groups of people (for example, movie stars, rich people, various racial, ethnic, or cultural groups) and how accurate these conceptions might actually be?
6. Describe cases in which you made a poor decision because your judgment was clouded by emotion or some other factor.



If You Want to Know More

1. *Using diagrams to determine the validity of syllogisms.* There are a number of ways to use diagrams to determine whether syllogisms are valid. One method uses Venn diagrams; another uses Euler circles.

Edwards, A. W. F. (2004). *Cogwheels of the mind: The story of Venn diagrams*. Baltimore: Johns Hopkins University Press.

Shin, S.-J. (1994). *The logical status of diagrams*. Cambridge, UK: Cambridge University Press.

2. Culture and cognition. We have seen that culture can affect inductive reasoning (see page 463 in this text). Other studies in which reasoning has been tested in different cultural groups (and usually compared with a sample of participants from the United States) have investigated how people from different cultures solve syllogisms, reason inductively, and solve math problems.

Dehaene, S., Izard, V., Pica, P., & Spelke, E. (2006). Core knowledge of geometry in an Amazonian indigene group. *Science*, *311*, 381–384.

Nisbett, R. E. (2003). *The geography of thought*. New York: Free Press.

Scribner, S. (1977). Modes of thinking and ways of speaking: Culture and logic reconsidered. In P. N. Johnson-Laird & P. C. Wason (Eds.), *Thinking: Readings in cognitive science* (pp. 483–500). Cambridge: Cambridge University Press.

Tang, Y., Zhang, W., Chen, K., Feng, S., Ji, Y., Shen, J., Reiman, E. M., & Liu, Y. (2006). Arithmetic processing in the brain shaped by cultures. *Proceedings of the National Academy of Sciences*, *103*, 10775–10780.

3. Making decisions rapidly. The deliberation-without-attention effect states that under certain conditions, people make better decisions when they act quickly than when they consciously deliberate over the decisions.

Dijksterhuis, A., Bos, M. W., Nordgren, L. F., & van Baaren, R. B. (2006). On making the right choice: The deliberation-without-attention effect. *Science*, *311*, 1005–1007.

4. Reasoning and the law. Juries are asked to come to a conclusion by evaluating evidence. New research indicates that a person's beliefs can affect his or her decision making. Also, the brain's response depends on whether the evidence is consistent or inconsistent with the person's beliefs.

Fugelsang, J. A., & Dunbar, K. N. (2004). A cognitive neuroscience framework for understanding causal reasoning and the law. *Philosophical Transactions of the Royal Society of London, B: Biological Sciences*, *359*, 1749–1754.

5. Regret and decision making. The results of the ultimatum game experiments show that human decision making can be influenced by emotions. Other experiments have shown that people's decision making is influenced by a desire to avoid the regret they would experience if they made the wrong decision.

Coricelli, G., Critchley, H. D., Joffily, M., O'Doherty, J. P., Sirigu, A., & Dolan, R. J. (2005). Regret and its avoidance: A neuroimaging study of choice behavior. *Nature Neuroscience*, *8*, 1255–1262.

6. Another view of rationality. We described the idea that people's reasoning and decision making can be negatively affected by bias and the use of heuristics. Some researchers have proposed another approach that sees people as behaving more rationally than this view gives them credit for.

Chase, V. M., Hertwig, R., & Gigerenzer, G. (1998). Views of rationality. *Trends in Cognitive Sciences*, 2, 206–214.

7. Embodied cognition. Some cognitive psychologists propose that to understand cognition, we need to take into account how a person's body interacts with the environment. This idea has implications for understanding problem solving, reasoning, and language.

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin and Review*, 9, 625–636.

8. Neurons that respond to abstract rules. There are neurons in the monkey PFC that respond to abstract rules. This adds to the evidence that the PFC is important for problem solving.

Wallis, J. D., Anderson, K. C., & Miller, E. K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature*, 411, 953–956.

Key Terms

Affirming the antecedent, 444

Affirming the consequent, 444

Antecedent, 443

Atmosphere effect, 439

Availability heuristic, 456

Base rate, 460

Belief bias, 439

Categorical syllogism, 437

Conclusion, 437

Conditional syllogism, 443

Confirmation bias, 462

Conjunction rule, 461

Consequent, 443

Decisions, 437

Deductive reasoning, 436

Denying the antecedent, 444

Denying the consequent, 444

Diversity principle, 464

Economic utility theory, 466

Evolutionary perspective on cognition, 450

Falsification principle, 446

Focusing illusion, 467

Framing effect, 471

Illusory correlation, 458

Inductive reasoning, 436

Law of large numbers, 462

Mental model, 441

Mental simulation, 467

Neuroeconomics, 475

Omission bias, 477

Opt-in procedure, 469

Opt-out procedure, 469

Permission schema, 448

Perseveration, 473

Pragmatic reasoning schema, 448

Premise, 437

Reasoning, 435

Representativeness heuristic, 459

Risk-aversion strategy, 470

Risk-taking strategy, 470

Similarity-coverage model, 464
Social-exchange theory, 450
Stereotype, 459
Syllogism, 436
Typicality principle, 464

Ultimatum game, 475
Utility, 466
Valid, 438
Wason four-card problem, 445

CogLab

To experience these experiments for yourself, go to <http://coglab.wadsworth.com>. Be sure to read each experiment's setup instructions before you go to the experiment itself. Otherwise, you won't know which keys to press.

Primary Labs

Wason selection task Two versions of the Wason four-card problem (p. 445).


Typical reasoning How the representativeness heuristic can lead to errors of judgment (p. 459).

Risky decisions How decision making is influenced by framing effects (p. 469).

Decision making An experiment that demonstrates how decisions can be affected by the context within which the decision is made (p. 469).

Related Lab

Monty Hall A simulation of the Monty Hall three-door problem, which involves an understanding of probability.



The cue ball is to the left of your line of sight.

■ **Figure 12.21** Result of visualization of pool balls (see page 441).