It is well known that changes in attitude can lead to changes in behavior, but research on cognitive dissonance shows that changes in attitude can also *follow* changes in behavior. According to the theory of cognitive dissonance, the pressure to feel consistent will often lead people to bring their beliefs in line with their behavior. In Chapter 3, we will see that, in many cases, people also distort or forget what their initial beliefs were.

CHAPTER 3

MEMORY AND HINDSIGHT BIASES

"To-day isn't any other *day, you know."*
"I don't understand you," said Alice. "It's dreadfully confusing!"
"That's the effect of living backwards," the Queen said kindly: "it always makes one a little giddy at first—"
"Living backwards!" Alice repeated in great astonishment. "I never heard of such a thing!"
"—but there's one great advantage in it, that one's memory works both ways. . . . It's a poor sort of memory that only works backwards," the Queen remarked.

—Lewis Carroll, *Through the Looking-Glass*

Take a moment to reflect on whether the following statement is true or false: "Memory can be likened to a storage chest in the brain, into which we deposit material and from which we can withdraw it later if needed. Occasionally, something gets lost from the 'chest,' and then we say we have forgotten."

What do you think—true or false? (See Item #19 of the Reader Survey for your answer.) Roughly 85 percent of the college students in a study by P. A. Lamal (1979, October) agreed with this statement, yet something is terribly wrong with the way it characterizes memory (aside from the question of whether material is ever truly lost from memory).

Memories are *not* like copies of our past experiences on deposit in a memory bank. Instead, they are constructed at the time of withdrawal (Loftus, 1980; Myers, 1990). The "materials" used in this split-second reconstruction are logical inferences that fill in missing detail, associated memories that blend in with the original memory, and other relevant information. To verify that memory is reconstructive, try an exercise suggested by Myers (1990): Close your eyes and recall a scene in which you experienced something pleasurable. Don't read any further until you have finished replaying your experience.

Did you see yourself in the scene? Most people do. But if you saw yourself, then you must have reconstructed the scene (unless, of course, you were looking at yourself during the original experience).

SHATTERED MEMORIES

One of the best demonstrations that memory is reconstructive was provided in two experiments by Beth Loftus and John Palmer (1974). In the first experiment, 45 students were asked to view seven different film clips depicting a traffic accident. The clips ranged from five to thirty seconds in length and were borrowed from longer driver's education films.

After each film clip, students answered a series of questions, including one on how fast the cars had been traveling. One-fifth of the students answered the question: "About how fast were the cars going when they contacted each other?" Equal numbers of the remaining students answered the same question, except that the word "contacted" was replaced with "hit," "bumped," "collided," or "smashed."

As you can see in Table 3.1, students who were asked how fast the cars were going when they "smashed" gave a mean estimate that was 9 miles faster than the average estimate given by students who were asked how fast the cars were going when they "contacted" each other. Thus, Loftus and Palmer concluded that the form of a question—even when changed by only one word—can markedly affect how people reconstruct their memory of an event.

If anything, results from the second experiment were even more dramatic. This time, Loftus and Palmer had 150 students watch a one-minute film that included a four-second, multiple-car crash. Fifty students were asked: "About how fast were the cars going when they smashed into each other?" Another 50 students were asked: "About how fast were the cars going when they hit each other?" And the last 50 students were not asked to judge car speed. Then the students returned one week later and, without viewing the film again, answered a series of questions. The key question Loftus and Palmer were interested in was whether students remembered having seen any shattered glass during the car crash.

Loftus and Palmer found that asking students how fast the cars were going when they "smashed" not only led to faster estimates, but that one

TABLE 3.1
HOW FAST WERE THE CARS GOING WHEN THEY . . .

Verb	Mean Speed
Smashed	40.8
Collided	39.3
Bumped	38.1
Hit	34.0
Contacted	31.8

Note: These are average speed estimates from Experiment 1 of a study by Elizabeth Loftus and John Palmer (1974).

TABLE 3.2
DID YOU SEE ANY BROKEN GLASS?

	EXPERIMENTAL CONDITION		
Response	"Smashed"	"Hit"	Control Group
Yes	16	7	6
No	34	43	44

Note: This is the distribution of "yes" and "no" answers found in Experiment 2 of a study by Elizabeth Loftus and John Palmer (1974). Fifty subjects were assigned to each of the three experimental conditions.

week later, a greater proportion of the students remembered the accident as having involved broken glass. The results, which show statistically reliable differences among the experimental conditions, are shown in Table 3.2. What is interesting about these results is that the accident never involved broken glass—subjects who estimated the speed of smashing cars reconstructed the accident so that it involved broken glass!

SWEET REMEMBRANCES

As the experiments by Loftus and Palmer show, memories are not fixed in storage. In 1971, John Bransford and Jeffrey Franks further showed that memories are not stored separately from one another. Bransford and Franks (1971) initially presented college students with a list of sentences about an event. For example, one of the lists—reprinted on page 7 of the Reader Survey—went like this:

- The ants ate the sweet jelly which was on the table.
- The ants were in the kitchen.
- The ants ate the sweet jelly.
- The ants in the kitchen ate the jelly which was on the table.
- The jelly was on the table.
- The ants in the kitchen ate the jelly.

Then, after five minutes or so, students were presented with another list of sentences and asked to indicate which sentences were in the first list. They were also asked to rate their confidence in each answer on a scale from 1 to 5. Item #34 in the Reader Survey contains a second list of sentences, along with blanks for confidence ratings as to whether these sentences appeared in the original list.

As it happens, the only sentence that appeared in the first set was Item #34c: "The ants ate the sweet jelly." If you are similar to most people in the study by Bransford and Franks, you were moderately confident (2 to 4 on the confidence scale) that this sentence had appeared before.

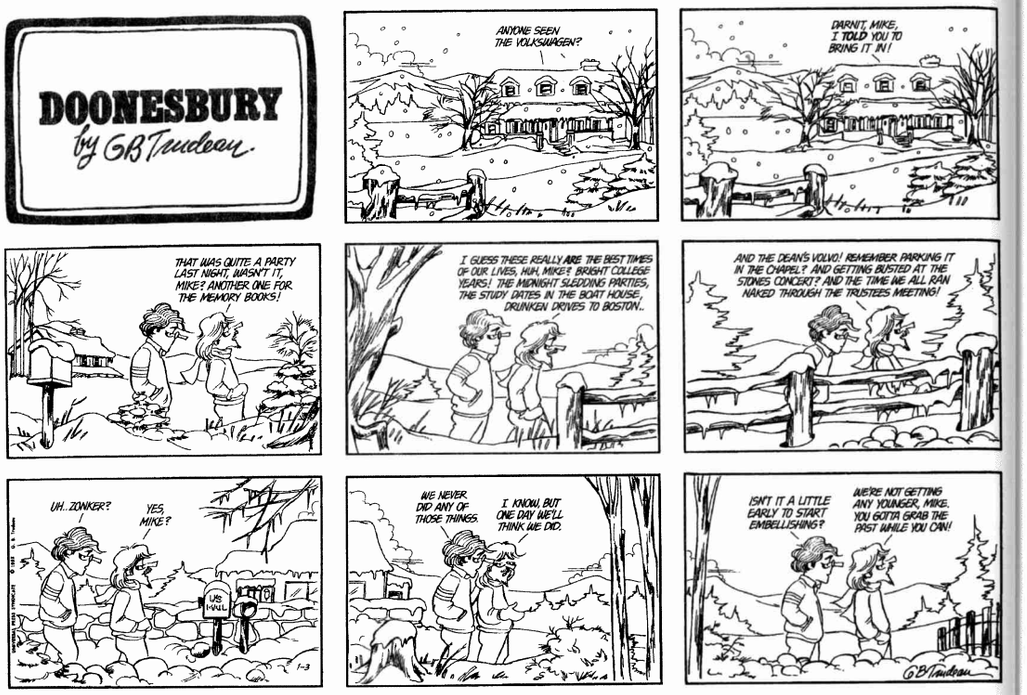

FIGURE 3.1
Reconstructive memory. (Doonesbury *copyright 1992 G. B. Trudeau. Reprinted with permission of Universal Press Syndicates. All rights reserved.*)

More interesting is your response to Item #34b: "The ants in the kitchen ate the sweet jelly which was on the table." Even though this sentence did not appear in the first set, students in the study by Bransford and Franks tended to be very confident that they had seen it before. Were you?

What is significant about the sentence in Item #34b is that it contains combinations of relations that are not contained in any individual sentence from the first set. The original sentences never explicitly stated that the jelly in the kitchen was sweet, or that the ants in the kitchen ate the sweet jelly. The sentence in Item #34b can only be derived by combining separate sentences from the first set.

Thus, people do not simply memorize sentences; they construct and memorize a general scenario. Once one piece of information is integrated with others, it is sometimes difficult to remember which information was new and which was already known.

I KNEW IT ALL ALONG

People also have difficulty telling how they are affected by information about an outcome. For example, if they learn about the results of a psychological experiment, they tend to regard the findings as having been fairly predictable all along—or at least more predictable than they would have judged before learning of the results (one of the reasons why this book has a Reader Survey!). Moreover, if people are asked to behave as though they know nothing about the outcome of an experiment, they still respond more like people who know about the results than people who do not. That is, if they are asked to estimate in retrospect how likely they once thought the results were to occur, they assign higher probabilities than do people predicting the experimental outcome in advance.

This tendency is known as "hindsight bias," or the "I-knew-it-all-along" effect. Hindsight bias is the tendency to view what has already happened as relatively inevitable and obvious—without realizing that retrospective knowledge of the outcome is influencing one's judgments. Hindsight biases have been documented in elections (Leary, 1982; Synodinos, 1986), medical diagnoses (Arkes, Wortmann, Saville, & Harkness, 1981), pregnancy tests (Pennington, Rutter, McKenna, & Morley, 1980), buying decisions (Walster, 1967), games (Leary, 1981), and a number of other areas. They have also been shown using a variety of experimental techniques, response instructions, and groups of people (for reviews of hindsight bias and related effects, see Campbell & Tesser, 1983; Christensen-Szalanski & Willham, 1991; Hawkins & Hastie, 1990; Verplanken & Pieters, 1988).

One of the first studies on hindsight bias was published in 1975 by Baruch Fischhoff and Ruth Beyth. The main events Fischhoff and Beyth (1975) used in their study were President Nixon's trips to China and the Soviet Union in 1972. In the first phase of the experiment, several groups of Israeli students were asked to estimate the probability of 15 different outcomes for either the China trip or the Soviet trip—*before* the trip took place. For example, students who were asked about the China trip estimated the chances that the United States would establish a diplomatic mission in Peking, that President Nixon would meet Mao at least once, that President Nixon would announce the trip a success, and so forth. Similarly, students who were asked about Nixon's trip to the Soviet Union estimated outcomes such as the establishment of a joint space program, or the arrest of Soviet Jews trying to speak with Nixon.

In the second phase of the study—two weeks to six months after the trip had taken place—students were asked to recall what their earlier predictions had been. For instance, students who had answered questions about the China trip were told the following:

> As you remember, about two weeks ago, on the eve of President Nixon's trip to China, you completed a questionnaire by providing probabilities for the occurrence of a number of possible outcomes of the trip. We are presently interested in the relation between the quality of people's predictions and their ability to remember their predictions. For this reason, we would like to have you fill out once again the same questionnaire which you completed two weeks ago, giving the *same probabilities which you gave then*. If you cannot remember the probability which you then assigned, give the probability which you would have given to each of the various outcomes on the eve of President Nixon's trip to China.

Students were also asked to indicate whether, as far as they knew, each outcome had in fact occurred. Fischhoff and Beyth wanted to see if students would remember their predictions as having been more accurate than they actually were.

In general, this is just what Fischhoff and Beyth (1975) found. Three-quarters of the students tended to remember having assigned higher probabilities than they actually had to outcomes that they thought had occurred, and the majority of students remembered having assigned lower probabilities to outcomes they believed had not occurred. Hindsight biases were particularly strong when the initial predictions preceded the recall task by several months. When three to six months separated the prediction and recall tasks, 84 percent of the students showed hindsight biases—after learning the outcome of Nixon's trips, they viewed the outcome as having been more predictable than it actually was.

REDUCING HINDSIGHT BIAS

In 1977, Paul Slovic and Baruch Fischhoff published a study that showed how to reduce hindsight biases when learning the results of research (the feeling of having known the results all along). Slovic and Fischhoff found that hindsight biases diminished when people stopped to consider reasons why the results might have turned out differently.

Subjects in this research read four brief descriptions of studies drawn from biology, psychology, and meteorology. *Foresight* subjects were told that the four studies would be conducted soon, and *hindsight* subjects were told that the studies had already been conducted. After reading about each study, all subjects then estimated the probability of replicating an outcome obtained on the first experimental trial (each trial always had two possible outcomes). In other words, hindsight subjects were told that a particular outcome had already been observed, and foresight subjects were asked to *suppose* the outcome occurred.

Slovic and Fischhoff (1977) found that compared with foresight subjects, hindsight subjects gave higher probability estimates that all future trials would replicate the first one. This difference was substantially reduced, however, when hindsight subjects were asked to consider reasons why *either* experimental outcome might have occurred. Hindsight bias was still present in this case, but to a much lesser degree.

Thus, the moral of the story is as follows: If you want to reduce hindsight biases, you should explicitly consider how past events might have turned out differently. If you only consider the reasons why something turned out as it did, you run a good risk of overestimating how inevitable that outcome was and how likely similar outcomes are in the future. In fact, Fischhoff (1977) has found that informing people about hindsight bias and encouraging them to avoid it is not enough to remove the bias. To avoid the ravages of hindsight bias, it is important to consider how an alternative outcome might have occurred.

CONCLUSION

In his book *Memory*, Ian Hunter (1964) tells the story of two British psychologists who secretly recorded a discussion that took place after a meeting of the Cambridge Psychological Society. Two weeks later, the psychologists contacted all the participants and asked them to write down everything they could remember about the discussion. When these accounts were checked against the original recording, it turned out that respondents typically omitted more than 90 percent of the specific points that had been discussed. Moreover, of the points that were recalled, nearly half were substantially incorrect. Respondents remembered comments that were never made, they transformed casual remarks into lengthy orations, and they converted implicit meanings into explicit comments.

This story highlights the value of keeping accurate records. Even the most sophisticated decision maker is susceptible to biases in memory, and there is no better way to avoid these biases than maintaining careful notes and records of past events (e.g., meetings, important conversations, and agreements). As the research in this chapter shows, memory is, by its very nature, reconstructive and highly dependent upon contextual factors. Chapter 4 further explores the effects of context on judgment and decision making.