CHAPTER

8 Memory

Genghis Khan got around. At the dawn of the 13th century, the Mongolian warrior conquered the largest empire the world had ever known: an expanse stretching from the Sea of Japan in the east to the Caspian Sea in the west, from Siberia in the north to India in the south. To conquer this territory and then maintain his domination, the emperor had to formulate complex plans. This created a problem: His soldiers were illiterate peasants, scattered over thousands of miles. How could he spread his complicated orders through the ranks quickly, simply, and without error?

His solution: Put the orders in a song. All the Khan's soldiers learned a small set of melodies, which they practiced as they traversed the mountains and steppes. Then, when the time for fighting arrived, commanders would set their orders to the tune of one of these melodies. The soldiers' task was simple: memorize a few new verses for an old song, rather than a series of entirely unfamiliar, abstract instructions. And if any one of the soldiers forgot the lyrics, hundreds of others could sing him the next line. Using this scheme, the soldiers crooned their battle instructions, and large segments of Eurasia fell.

Others in the ancient world also relied on deliberate memorization strategies. The Greeks of classical Athens, for example, put a high value on public speaking, much of which was done from memory. The Greeks therefore developed a number of memorization tricks (some of which we'll discuss later in the chapter) to help them in this endeavor.

Similar mnemonic tactics are used in the modern world. Medical students, for example, have developed strategies that help them memorize anatomy, drug names, and disease symptoms. Thus, they learn the 12 pairs of cranial nerves (olfactory, optic, oculomotor, trochlear, trigeminal, and so on) by taking the first letter of each word and forming a sentence built from new words that start with the same letters. The resulting sentence—"On Old Olympus' Towering Tops A Friendly Viking Grew Vines and Hops"—paints a vivid image that's far easier to remember than the original list.

These examples remind us that—with just a bit of work—we can get enormous amounts of information into our memories, and then recall that information, in detail, for a very long time. But there's also a darker side to memory: Sometimes we remember things that never happened at all. Indeed, far more often than we realize, our memories blend together separate incidents, introduce rogue details, and incorporate others' versions of events into our own recall. In this chapter, you'll learn how these memory errors arise and what they tell us about remembering.

How far off track can memory go? In one study, researchers planted in participants a memory of getting lost in the mall as a child, then being brought home safely by a friendly stranger. Nothing of the sort had happened to anyone in the study, but they came to vividly "remember" it anyhow. Other studies have planted false memories of vicious animal attacks, and even—in one remarkable study—a false memory of a hot-air balloon ride.

How should we put these pieces together? How does memory operate, so that we can easily remember countless episodes, thousands of facts, and the lyrics to hundreds of songs? Why does Genghis Kahn's lyrical trick, or the medical students' sentence-building strategy, help memory? More broadly, what can we do to learn more rapidly and hold on to the information longer? And why do our memories sometimes betray us, leading us to endorse large-scale fictions? We'll tackle all of these questions in this chapter.

ACQUISITION, STORAGE, RETRIEVAL

Each of us has a huge number of memories. We can recall what we did yesterday, or last summer. We can remember what the capital of France is, or what the chemical formula is for water. We remember how to ride a bicycle and how to throw a baseball. These examples—remembering *episodes,* remembering *general facts,* and remembering *skills* or *procedures*—actually draw on different memory systems; but it also turns out that the various types of memory have some things in common, so let's begin with the common elements.

Any act of remembering requires success at three aspects of the memory process. First, in order to remember, you must learn something—that is, you must put some information into your memory. This point seems obvious, but it deserves emphasis because many failures of memory are, in fact, failures in this initial stage of *acquisition.* For example, imagine meeting someone at a party, being told his name, and moments later realizing that you don't have a clue what his name is—even though you just heard it! This common (but embarrassing) experience is probably not the result of ultrarapid forgetting. Instead, it's likely to stem from a failure in acquisition. You were exposed to the name but barely paid attention to it and, as a result, never learned it in the first place.

The next aspect of remembering is *storage*. To be remembered, an experience must leave some record in the nervous system. This record—known as the memory trace—is squirreled away and held in some enduring form for later use. One question to be asked here is how permanent this storage is: Once information is in storage, does it stay there forever? Or does information in storage gradually fade away? We'll tackle these questions later in this chapter.

The final aspect of remembering is *retrieval*, the process through which you draw information from storage and use it. Sometimes, retrieval takes the form of **recall**—a process in which you retrieve information from memory in response to some cue or question (Figure 8.1A). Trying to answer a question like "What is Sue's boyfriend's name?" or "Can you remember the last time you were in California?" requires recall. A different way to retrieve information is through **recognition** (Figure 8.1B). In this kind of retrieval, you're presented with a name, fact, or situation and are asked if you have encountered it before. "Is this the man you saw at the bank robbery?" or "Was the movie you saw called *Memento*?" are questions requiring recognition. Recognition can also be tested with multiple items: "Which of these pictures shows the man you saw earlier?" This latter format obviously resembles a multiple-choice exam, and in fact multiple-choice testing in the classroom probes your ability to recognize previously learned material. In contrast, exams that rely on essays or short answers emphasize recall.

8.1 **Using memory** (A) In this card game, you need to recall which card is in which position; in this case, *position* is the memory cue, and *card identity* is what you're trying to recall. (B) Most standardized tests, in multiple-choice format, rely on recognition. The correct answer is in front of you, as one of your options, and you need to recognize it.

recall A type of retrieval that requires you to produce an item from memory in response to a cue or question.

recognition A type of retrieval that requires you to judge whether you have encountered a stimulus previously.

acquisition The processes of gaining new information and placing it in memory.

intentional learning Placing new information into memory in anticipation of being tested on it later.

incidental learning Learning without trying to learn, and often without awareness that learning is occurring.

ACQUISITION

People commonly speak of "memorizing" new facts or, more broadly, of "learning" new material. However, psychologists prefer the term memory **acquisition** and use it to include cases of deliberate memorization (**intentional learning**) as well as cases of **incidental learning**—learning that takes place without any intention to memorize and often without the awareness that learning is actually occurring. (You know that grass is green and the sky is blue, and you probably can easily recall what you had for dinner yesterday, but you didn't set out to memorize these facts; the learning, therefore, was incidental.)

Memory acquisition is not just a matter of "copying" an event or a fact into memory, the way a camera copies an image onto film. Instead, acquisition requires some intellectual engagement with the material—thinking about it in some way—and it's then the product of this engagement (i.e., what you thought about during the event) that's stored in memory. As we'll see, this simple point turns out to have crucial implications for what you will remember and for how accurate (i.e., true to the actual history) your memory will be.

Working Memory, Long-Term Memory

How does memory acquisition proceed? The answer has to begin with the fact that we have several types of memory, each with different properties, and each type plays its own role in the acquisition process. Historically, these different types have been described in terms of the *stage theory of memory,* which proposed (among other points) that memory acquisition could be understood as dependent on three types of memory: When information first arrived, it was stored briefly in *sensory memory,* which held onto the input in "raw" sensory form—an *iconic memory* for visual inputs and an *echoic memory* for auditory inputs. A process of selection and interpretation then moved the information into *short-term memory*—the place you hold information while you're working on it. Some of the information was then transferred into *long-term memory*, a much larger and more permanent storage place (Atkinson & Shiffrin, 1968; Broadbent, 1958; Waugh & Norman, 1965).

working memory A term describing the status of thoughts in memory that are currently activated.

long-term memory The vast memory depository containing all of an individual's knowledge and beliefs—including all those not in use at any given time.

primacy effect In free recall, the tendency to recall the first items on the list more readily than those in the middle.

recency effect In free recall, the tendency to recall items at the end of the list more readily than those in the middle.

This early conception of memory captured some important truths—but needs to be updated in several ways. As one concern, the idea of "sensory memory" plays a much smaller role in modern theorizing and so, for example, many discussions of visual information processing (like our discussion in Chapters 4 and 5) make no mention of iconic memory. In addition, modern proposals use the term **working memory** rather than short-term memory to emphasize the function of this memory: Ideas or thoughts in this memory are currently activated, currently being thought about—and so they're the ideas you are currently *working on*. **Long-term memory,** in contrast, is the vast depository that contains all of your knowledge and all of your beliefs that you happen not to be thinking about at the moment, and this includes your beliefs about relatively recent events. Thus, if just a few minutes ago you were thinking about your weekend plans but now you're thinking about something else, these plans are gone from working memory (because you're no longer working on them); and so, if you can recall your plans, you must be drawing them from long-term memory.

Let's note, though, that what's at stake here is more than a shift in terminology, because the modern view also differs from the stage theory in how it conceptualizes memory. In the older view, working memory was understood broadly as a *storage place,* and it was often described as the "loading dock" just outside the long-term memory "warehouse." In the modern conception, working memory is not a "place" at all; instead, it's just the name we give to a *status.* When we say that ideas are "in working memory," this simply means—as we've already noted—that these ideas are currently activated. This focus on status is also the key to understanding the difference between working memory and long-term memory—the modern conception emphasizes whether the mental content is currently active (working memory) or not (long-term memory), in contrast to older theory's emphasis on time frame ("short term" or "long").

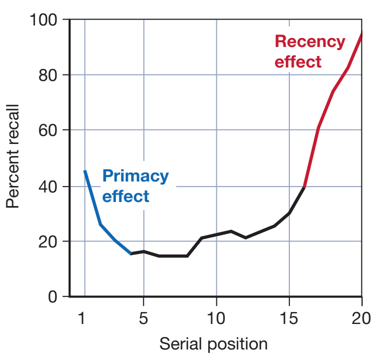

8.2 Primacy and recency effects in free recall Research participants heard a list of 20 common words presented at a rate of 1 word per second. Immediately after hearing the list, participants were asked to write down as many of the words on the list as they could recall. The results show that position in the series strongly affected recall—participants had better recall for words at the beginning of the list (a pattern called the primacy effect) and for words at the end of the list (the recency effect).

PRIMACY AND RECENCY

Why should we take this broad proposal seriously? Why should we make any distinction between working memory and long-term memory, and why should we think about working memory in the way we've just described? As a first step toward answering these questions, consider the results of studies in which participants hear a series of unrelated words—perhaps 15 words in total, or 20, presented one word at a time. At the end of the list, the participants are asked to recall the items in any order they choose (this is why the participants' task is called *free recall*—they're free to recall the items in any sequence).

In this task, there's a reliable pattern for which words the participants recall and which ones they don't. Words presented at the beginning of the list are very likely to be recalled; this memory advantage for early-presented words is called the **primacy effect.** Likewise, the last few words presented are also likely to be recalled; this is the **recency effect.** The likelihood of recall is appreciably poorer for words in the middle of the list (Figure 8.2).

What creates this pattern? As the to-be-remembered words are presented, the participants pay attention to them, and this ensures the activated status that we call "working memory." There's a limit, however, on how many things someone can think about at once, and so there's a limit on how many items can be maintained in working memory. According to many authors, this limit is seven items, give or take one or two; the capacity of working memory is therefore said to be *seven plus or minus two* items (G. Miller, 1956). As a result, it's just not possible for the participants to maintain all of the list words in their current thoughts. Instead, they'll just do their best to "keep up" with the list as they hear it. Thus, at each moment during the list presentation, their working memories will contain only the half-dozen or so words that arrived most recently.

Notice that, in this situation, new words entering working memory will "bump out" the words that were there a moment ago. The only words that don't get bumped out are the last few words on the list, because obviously no further input arrives to displace them. Hence, when the list presentation ends, these few words are still in working memory—still in the participants' thoughts—so are easy to recall. This is why the participants reliably remember the end of the list; they are producing the result we call the recency effect.

The primacy effect comes from a different source. We know that these early words are not being recalled from working memory, because they were—as we've already noted—bumped from working memory by later-arriving words. It seems, therefore, that the primacy effect must involve long-term memory—and so, to explain why these early words are so well recalled, we need to ask how these words became well established in long-term storage in the first place.

The explanation lies in how participants allocate their attention during the list presentation. To put this in concrete terms, let's say that the first word on the list is *camera*. When research participants hear this word, they can focus their full attention on it, silently rehearsing "*camera, camera, camera,* . . ." When the second word arrives, they'll rehearse that one too; but now they'll have to divide their attention between the first word and the second ("*camera, boat, camera, boat,* . . ."). Attention will be divided still further after participants hear the third word ("*camera, boat, zebra, camera, boat, zebra,* . . ."), and so on through the list.

Notice that earlier words on the list get more attention than later ones. At the list's start, participants can lavish attention on the few words they've heard so far. As they hear more and more of the list, though, they must divide their attention more thinly, simply because they have more words to keep track of. Let's now make one more assumption: that the extra attention given to the list's first few words makes it more likely that these words will be well established in long-term memory. On this basis, participants will be more likely to recall these early words than words in the middle of the list—exactly the pattern of the data.

Support for these interpretations comes from various manipulations that affect the primacy and recency effects. For example, what happens if we require research participants to do some other task immediately after hearing the words but before recalling them? This other task will briefly divert the participants' attention from rehearsing or thinking about the list words—and so the words will be bumped out of working memory. Working memory, in turn, was the hypothesized source of the recency effect, and so, according to our hypothesis, this other task—even if it lasts just a few seconds—should disrupt the recency effect. And indeed it does. If participants are required to count backward for just 30 seconds between hearing the words and recalling them, the recency effect is eliminated (Figure 8.3).

Other manipulations produce a different pattern—they alter the primacy effect but have no effect on recency. For example, if we present the list items more slowly, participants have time to devote more attention to each word. But we've just proposed that attention helps to establish words in long-term memory. We should therefore expect that a slower

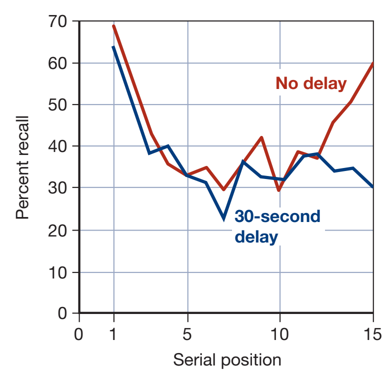

8.3 The recency effect and working memory Research participants heard several 15-word lists. In one condition (red), free recall was tested immediately after hearing the list. In the other condition (blue), the recall test was given after a 30-second delay during which rehearsal was prevented. The delay left the primacy effect unaffected but abolished the recency effect, confirming that this effect is based on retrieval from working memory.

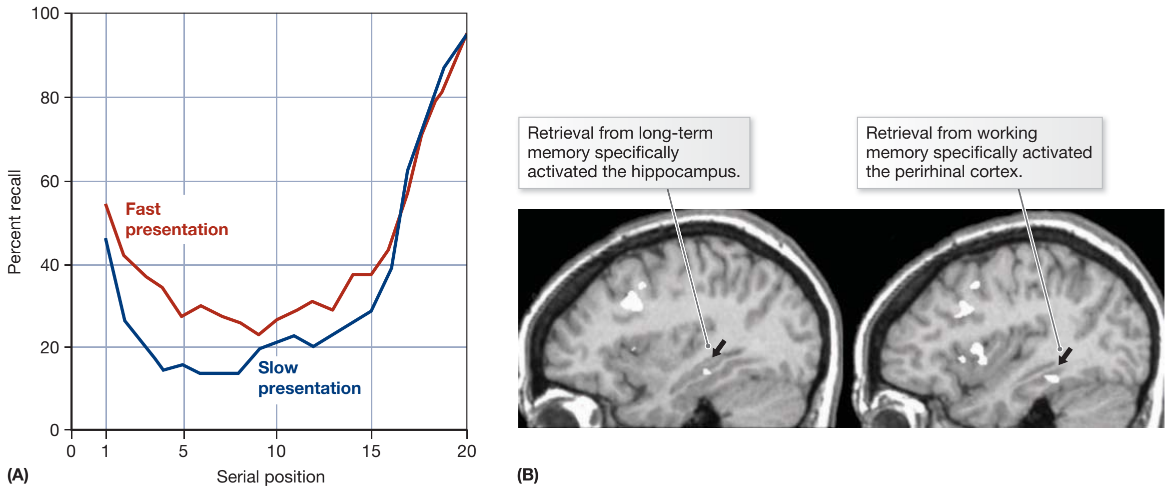

8.4 The primacy effect and long-term storage (A) The graph compares free-recall performance when item presentation is relatively slow (2 seconds per item) and fast (1 second per item). Slow presentation enhances the primacy effect but leaves the recency effect unaltered. (B) We can also confirm the distinction between working memory and long-term memory with fMRI scans. These suggest that memory for early items on a list depends on brain areas (in and around the hippocampus) that are associated with long-term memory; memory for later items on the list do not show this pattern (Talmi, Grady, Goshen-Gottstein & Moscovitch, 2005). This obviously provides confirmation that the recency items are coming from a different source than items heard earlier in the list.

presentation will lead to a stronger primacy effect (since primacy depends on retrieval from long-term memory) but no change in the recency effect (because the recency items aren't being retrieved from long-term memory). This is exactly what happens (Figure 8.4).

RECODING TO EXPAND THE CAPACITY OF WORKING MEMORY

As we've mentioned, working memory has a limited capacity. There is, however, enormous flexibility in how we use that capacity—and so, if we can pack the input more efficiently, we can increase the amount of information maintained in this memory.

For example, consider an individual who tries to recall a series of digits that she heard only once:

177620001066

If she treats this as a series of 16 unrelated digits, she'll surely fail in her attempt to remember the series. But if she thinks of the digits as years (i.e., the year the U.S. Declaration of Independence was signed; the year of the new millennium; and the year the Normans invaded England), the task becomes much easier because now she has just three items to remember.

Cases like this one make it plain that working memory's capacity can't be measured in *digits*, or *words*, or *kilobytes*. Instead, the capacity is measured in **chunks.** This unscientific-sounding word helps us remember that this is a flexible sort of measurement, because what's in a chunk depends on how the person thinks about, and organizes, the information. Thus, if a person thinks of each digit as a chunk, working memory

chunking A process of reorganizing (or recoding) materials in working memory by combining a number of items into a single, larger unit.

can hold (roughly) seven digits. If *pairs* of digits are chunked together, working memory's capacity will be more than a dozen digits.

To see how important chunking can be, consider a remarkable individual studied by Chase and Ericsson (Chase & Ericsson, 1978, 1979, 1982; Ericsson, 2003). This fellow happens to be a fan of track events, and when he hears numbers, he thinks of them as finishing times for races. The sequence "3, 4, 9, 2," for example, becomes "3 minutes and 49 point 2 seconds, near world-record mile time." In this way, four digits become one chunk of information. The man can then retain seven finishing times (seven chunks) in memory, and this can involve 20 or 30 digits. Better still, these chunks can be grouped into larger chunks, and these into even larger ones. For example, finishing times for individual racers can be chunked together into heats within a track meet, so that, now, 4 or 5 finishing times (more than 12 digits) become one chunk. With strategies like this and with a lot of practice, this man has increased his apparent memory capacity from the "normal" 7 digits to 79 digits!

Let's be clear, though, that what has changed through practice is merely the man's chunking strategy, not the holding capacity of working memory itself. This is evident in the fact that, when tested with sequences of letters rather than numbers—so he can't use his chunking strategy—his memory capacity drops to a perfectly normal six consonants. Thus, the seven-chunk limit is still in place for this fellow, even though (with numbers) he's able to make extraordinary use of these seven slots.

Establishing Long-Term Memories

So far, we've argued for a separation between working memory and long-term memory, and we're starting to see indications of each memory's attributes. Working memory has a small capacity—although it's flexible in what it can hold, thanks to the process of chunking. Long-term memory, in contrast, is vast. After all, the average college student remembers the meanings of 80,000 words, thousands of autobiographical episodes, millions of facts, hundreds of skills, the taste of vanilla and the smell of lemon. All these things and more are stored in long-term memory.

Working memory and long-term memory also differ in how they're "loaded" and "unloaded." To get information into working memory, all you need to do is pay attention to the material; that's built into the definition of working memory. Getting information into long-term storage, in contrast, seems to take some time and effort; that was essential for our discussion of the primacy effect.

We need to fill in some of the details, though, about how this "loading" of long-term memory works. With that, we'll get a clearer picture of why working memory is defined the way it is—as an active process rather than as a mere storage box.

THE IMPORTANCE OF ACTIVE ENGAGEMENT

In explaining primacy, we made a key assumption—namely, that *paying attention* to words on a list helps you establish those words in long-term memory. Presumably the same would be true for other contents, so that, no matter what you're memorizing, attention plays a key role in establishing memories. But is this assumption correct?

Consider people's memory for ordinary coins. Adults in the United States have probably seen pennies, for example, tens of thousands of times; adults in other countries have seen their own coins just as often. But, of course, most people have little reason to pay attention to the penny. Pennies are a different color and size from the other coins, so we can identify them at a fast glance and with no need for further

8.5 An ordinary penny Despite having seen the U.S. penny thousands and thousands of times, people seem to have little recollection of its layout. Test yourself. Which of these drawings is most accurate?

scrutiny. And, if attention is what matters for memory—or, more broadly, if we remember what we pay attention to and think about—then memory for the coin should be quite poor.

In one study, participants were asked whether Lincoln's profile, shown on the heads side of the penny, is facing to the right or the left. Only half of the participants got this question right—exactly what we'd expect if they were just guessing. Other participants were shown drawings of the penny, and had to choose the "right one" (Figure 8.5). Their performance was quite poor. These results—participants' remarkably poor memory for this coin despite countless opportunities to view it, provides striking confirmation that memory does require attention—it requires mental engagement with a target, not mere exposure (Nickerson & Adams, 1979; Rinck, 1999; for some complications, see Martin & Jones, 2006; in Figure 8.5, the top left drawing shows the correct layout).

But we need to be more precise about what *paying attention* means, and what it accomplishes. To make the issue clear, imagine you want to order a pizza. You look up the pizza restaurant's phone number on the Web or in a phone book, and then you walk across the room to pick up your phone and make the call. In this setting, you need to retain the number long enough to complete the dialing—and so, presumably, you're paying attention to the number for that span of time. But you have no need to memorize the number for later use, and so you're likely to think about the number in a limited way. Specifically, you're likely to employ what's called **maintenance rehearsal**—a mechanical process of repeating the memory items over and over, giving little thought to what the items are or whether they form any pattern.

maintenance rehearsal Mechanical repetition of material without thinking about its meaning or patterns.

This maintenance is easy and effective: It keeps the digits in your thoughts, and so you remember them long enough to place your call. But what happens if the line is busy when you call, and so you need to try again a moment later? In this setting, it's quite likely that you'll have forgotten the number and will need to look it up again! Apparently, maintenance rehearsal kept the number in working memory long enough for you to dial it the first time but utterly failed to establish it in long-term memory. As a result, you forget the number after just a few seconds.

THE LINK BETWEEN LONG-TERM MEMORY AND UNDERSTANDING

Apparently, establishing information in long-term storage is not an automatic process that is triggered merely by having the stimulus in front of your eyes or ears, or by having an idea mechanically maintained in working memory for a few seconds. Instead, some sort of work is involved so that, to put the matter simply, whether you'll remember something or not depends on how—and how fully—you thought about that information when you first met it.

As we've seen, we can confirm these claims by documenting how poor memory is for material that you've encountered but not paid much attention to. Further confirmation comes from studies that examine people's brain activity during learning. In brief, these studies show that during the learning process, some sort of effort is crucial for establishing long-term memories. Specifically, the studies show that greater levels of activity during the initial memory acquisition are reliably associated with greater probabilities of recall later on. This is especially true for brain activity in the hippocampus and regions of the prefrontal cortex (Brewer, Zhao, Desmond, Glover, & Gabrieli, 1998; A. Wagner, Koutstaal, & Schacter, 1999; A. Wagner et al., 1998), but it may also include brain activity in the parietal cortex (A. Wagner, Shannon, Kahn, & Buckner, 2005).

But what exactly is this brain activity accomplishing? Crucial information comes from studies that compare the memory effects of different types of engagement at the time of learning. In one study, participants were shown 48 words. As each word was presented, the participants were asked a question about it. For some words, they were asked about the word's physical appearance ("Is it printed in capital letters?"); this kind of question should produce **shallow processing**—an approach emphasizing the superficial characteristics of the stimulus. For other words, the participants were asked about the word's sound ("Does it rhyme with *train*?"); this should encourage an intermediate level of processing. For the remainder, they were asked about the word's meaning ("Would it fit into the sentence: The girl placed the ______ on the table?"); this presumably would lead to **deep processing**—an approach to the material that emphasizes what the stimulus means.

shallow processing An approach to memorization that involves focusing on the superficial characteristics of the stimulus, such as the sound of a word or the typeface in which it's printed.

deep processing An approach to memorization that involves focusing on the meaning of the stimulus.

After the participants had gone through the entire list of words, they were given an unexpected task: They were asked to write down as many of the words as they could remember. The results were clear-cut: Participants recalled very few of the words that called for shallow processing (capitalization). Words that required an intermediary level (sound) were recalled a bit better; and words that demanded the deepest level (meaning), were recalled best of all (Craik & Tulving, 1975).

Attention to a word's sound, therefore, is better for establishing memories than thoughtless and mechanical rehearsal; but attention to a word's *meaning* is better still and, across many studies, attention to meaning is reliably associated with high levels of subsequent recall. And it's not just the *search* for meaning that helps long-term memory. Instead, memory is promoted by *finding* the meaning—that is, by gaining an understanding of the to-be-remembered materials. In some studies, for example, experimenters have given participants material to read that was difficult to understand; then, immediately afterward, they probed the participants to see whether (or how well) they understood the material. Some time later, the experimenters tested the participants' memory for this material. The result was straightforward: the better the understanding at the time the material was presented, the better the memory later on (e.g., Bransford, 1979).

Other studies have manipulated the to-be-remembered material itself. For example, in one experiment, investigators presented this (tape-recorded) passage:

> The procedure is actually quite simple. First you arrange things into different groups depending on their makeup. Of course, one pile may be sufficient depending on how much there is to do. If you have to go somewhere else due to lack of facilities that is the next step; otherwise you are pretty well set. It is important not to overdo any particular endeavor. That is, it is better to do too few things at once than too many. In the

short run this may not seem important, but complications from doing too many can easily arise. A mistake can be expensive as well. The manipulation of the appropriate mechanisms should be self-explanatory, and we need not dwell on it here. At first, the whole procedure will seem complicated. Soon, however, it will become just another facet of life. It is difficult to foresee any end to the necessity for this task in the immediate future, but then one never can tell. (Bransford & Johnson, 1972, p. 722)

Half of the people heard this passage without any further information as to what it was about, and, when tested later, their memory for the passage was poor. The other participants, though, were given a clue that helped them to understand the passage—they were told, "The paragraph you will hear will be about washing clothes." This clue allowed that group to make sense of the material and dramatically improved their later recall (Bransford & Johnson, 1972; for a related example with a nonverbal stimulus, see Figure 8.6).

There's a powerful message here for anyone hoping to remember some body of material—for example, a student trying to learn material for the next quiz. Study techniques that emphasize efforts toward *understanding* the material are likely to pay off with good memory later on. Memory strategies that don't emphasize meaning will provide much more limited effects. Mechanical memory strategies—such as repeating the items over and over without much thought—may produce no benefits at all!

mnemonics Deliberate techniques people use to memorize new materials.

THE KEY ROLE FOR MEMORY CONNECTIONS

Attention to meaning is an effective way to establish long-term memories. Still, it's not the only way to establish memories, and we'll need to accommodate this point in our theorizing. What other memory acquisition procedures are effective? We can draw our answer from the study of **mnemonics**—deliberate techniques that people use to help them memorize new materials. Mnemonics come in many varieties, but all build on the same base: To remember well, it pays to establish memory connections. In some cases, the connections link the new material to ideas already in memory. In other cases, the connections link the various elements of the new material *to each other,* so that the mnemonic helps organize complex information into a small number of memory chunks.

The role of connections is clear, for example, in the various mnemonics that rely on *verse* in which a fixed rhythm or rhyme scheme links each element being memorized to the other elements within the poem. Thus, young children find it easier to memorize the calendar's layout if they cast the target information as a rhyme: "Thirty days hath September, April, June, and November," and high-school students have an easier time memorizing the fates of Henry VIII's wives by summarizing the history in a little verse: "divorced, beheaded, died; divorced, beheaded, survived."

Connections are also the key in other mnemonics, including ones that organize material by linking the first letters of the words in the sequence that's being memorized. Thus, students rely on ROY G. BIV to memorize the sequence of colors in the rainbow (*red, orange, yellow . . .*), and learn the lines in music's treble clef via "Every Good Boy Deserves Fudge" (the lines indicate the musical notes *E, G, B, D,* and *F*). Various first-letter mnemonics are also available for memorizing the taxonomic categories ("King Philip Crossed the Ocean to Find Gold and Silver," to memorize *kingdom,*

8.6 Nonverbal stimulus In general, we easily remember things that are meaningful but don't remember things that seem to have no meaning. This picture can be used to demonstrate this point with a nonverbal stimulus. At first the picture looks like a collection of meaningless blotches, and it's very hard to remember. But if viewers discover the pattern, the picture becomes meaningful and is then effortlessly remembered.

phylum, class, order, family, genus, and *species*). And so on for other memory tasks (Figure 8.7).

Still other mnemonics involve the use of mental imagery. One such technique, developed by the ancient Greeks, is the *method of loci,* which requires the learner to visualize each of the items she wants to remember in a different spatial location ("locus"). In recall, the learner mentally inspects each location and retrieves the item that she placed there in imagination. Does this work? In one study, college students had to learn lists of 40 unrelated nouns. Each list was presented once for about 10 minutes, during which the students tried to visualize each of the 40 objects in a specific location on their college campus. When tested immediately afterward, the students recalled an average of 38 of the 40 items; when tested one day later, they still managed to recall 34 (Bower, 1970; also see Bower, 1972; Higbee, 1977; Roediger, 1980; J. Ross & Lawrence, 1968). In other studies, participants using the method of loci were able to retain seven times more than their counterparts who learned by rote.

It's also worth mentioning that visualization is, on its own, an effective memorization tool. If you're trying to remember a list of words, for example, it's helpful to form a mental picture of each item on the list (a mental picture of a hammer, for example, and then a mental picture of a puppy, and so on.) Visualization is far more effective, though, if it serves to link the to-be-remembered words to each other—and so here, once again, we see the importance of memory connections. To make this idea concrete, consider a student trying to memorize a list of word pairs. He might decide just to visualize the items side by side—and so (for example), after hearing the pair *eagle-train,* he might visualize an eagle and then, separately, he might visualize a train. Alternatively, he might try to form mental pictures that bring the items into some kind of relationship—so he might, for example, imagine the eagle winging to its nest with a locomotive in its beak. Evidence indicates that images of the second (interacting) sort produce much better recall than nonunifying images do (Wollen, Weber, & Lowry, 1972; also Figure 8.8).

Whether mnemonics are based on imagery or some other system, though, there's no question that they are enormously useful in memorizing, say, a list of foreign vocabulary words or the names of various parts of the brain. But before we move on, we should note that there's also a downside to using mnemonics: During learning, someone trying to memorize via a mnemonic is likely to focus all their attention on just a narrow set of connections—the fact that the locomotive is in the eagle's beak, or that *September*

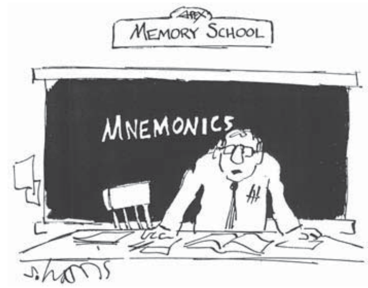

You simply associate each number with a word, such as "table" and 3,476,029.

8.7 Memory school Some mnemonics are more successful than others.

(A) Interactive depiction **(B) Noninteractive depiction**

8.8 Interacting and noninteracting depictions Research participants shown related elements, such as a doll sitting on a chair and holding a flag (A), are more likely to recall the trio of words *doll*, *flag*, and *chair* than are participants shown the three objects next to each other but not interacting (B).

rhymes with *November*. This strategy guarantees that these connections will be well established; and that's great if, later on, those connections are just the ones you need. But at the same time, if you focus on just these few connections, you're putting little effort into developing other possible connections—so you're not doing much to promote your *understanding* of the material you're memorizing. On this basis, mnemonics—as effective as they are for memorization—are an unwise strategy if understanding is your goal.

STORAGE

memory trace The physical record in the nervous system that preserves a memory.

memory consolidation The biological process through which memories are transformed from a transient and fragile status to a more permanent and robust state; according to most researchers, consolidation occurs over the course of several hours.

retrograde amnesia A memory deficit, often suffered after a head injury, in which the patient loses memory for events that occurred *before* the injury.

We've been focusing on the first step involved in memory—namely memory acquisition. Once a memory is acquired, though, it must be held in storage—i.e., held in long-term memory until it's needed. The mental representation of this new information is referred to as the **memory trace**—and, surprisingly, we know relatively little about exactly how traces are lodged in the brain. At a microscopic level, it seems certain that traces are created through the three forms of neural plasticity described in Chapter 7: Presynaptic neurons can become more effective in sending signals; postsynaptic neurons can become more sensitive to the signals they receive; and new synapses can be created.

On a larger scale, evidence suggests that the trace for a particular past experience is not recorded in a single location within the brain. Instead, different aspects of an event are likely to be stored in distinct brain regions—one region containing the visual elements of the episode, another containing a record of our emotional reaction, a third area containing a record of our conceptual understanding of the event, and so on (e.g., A. Damasio & H. Damasio, 1994). But, within these broad outlines, we know very little about how the information content of a memory is translated into a pattern of neural connections. Thus, to be blunt, we are many decades away from the science-fiction notion of being able to inspect the wiring of someone's brain in order to discover what he remembers, or being able to "inject" a memory into someone by a suitable rearrangement of her neurons. (For a recent hint about exactly how a specific memory might be encoded in the neurons, see Han et al., 2009.)

One fact about memory storage, however, is well established: Memory traces aren't created instantly. Instead, a period of time is needed, after each new experience, for the record of that experience to become established in memory. During that time, **memory consolidation** is taking place; this is a process, spread over several hours, in which memories are transformed from a transient and fragile status to a more permanent and robust state (Hasselmo, 1999; McGaugh, 2000, 2003; Meeter & Murre, 2004; Wixted, 2004).

What exactly does consolidation accomplish? Evidence suggests that this time period allows adjustments in neural connections, so that a new pattern of communication among neurons can be created to represent the newly acquired memory. This process seems to require the creation of new proteins, so it is disrupted by chemical manipulations that block protein synthesis (H. Davis & Squire, 1984; Santini, Ge, Ren, deOrtiz, & Quirk, 2004; Schafe, Nader, Blair, & LeDoux, 2001).

The importance of consolidation is evident in the memory loss sometimes produced by head injuries. Specifically, people who have experienced blows to the head can develop **retrograde amnesia** (*retrograde* means "in a backward direction"), in which they suffer a loss of memory for events that occurred before the brain injury (Figure 8.9). This form of amnesia can also be caused by brain tumors, diseases, or strokes (Cipolotti, 2001; M. Conway & Fthenaki, 1999; Kapur, 1999; Mayes, 1988; Nadel & Moscovitch, 2001).

Retrograde amnesia usually involves *recent* memories. In fact, the older the memory, the less likely it is to be affected by the amnesia—a pattern referred to as Ribot's law, in honor of the 19th-century scholar who first discussed it (Ribot, 1882). What produces this pattern? Older memories have presumably had enough time to consolidate, so they are less vulnerable to disruption. Newer memories are not yet consolidated, so they're more liable to disruption (A. Brown, 2002; Weingartner & Parker, 1984).

There is, however, a complication here: Retrograde amnesia sometimes disrupts a person's memory for events that took place months or even years before the brain injury. In these cases, interrupted consolidation couldn't explain the deficit unless one assumes—as some authors do—that consolidation is an exceedingly long, drawn-out process. (For discussion of when consolidation takes place, and how long it takes, see Hupbach et al., 2008; McGaugh, 2000.) However, this issue remains a point of debate, making it clear that we haven't heard the last word on how consolidation proceeds.

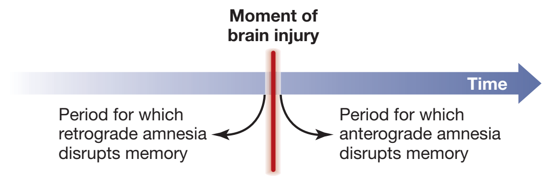

8.9 Retrograde and anterograde amnesia Retrograde amnesia disrupts memory for experiences *before* the injury, accident, or disease that triggered the amnesia. Anterograde amnesia disrupts memory for experiences *after* the injury or disease.

RETRIEVAL

When we learn, we transfer new information into our long-term store of knowledge, and then we consolidate this newly acquired information. But we still need one more step in this sequence, because memories provide no benefit for us if we can't retrieve them when we need them. Hence **retrieval**—the step of locating and activating information in memory—is crucial. Moreover, the success of retrieval is far from guaranteed, and many cases of apparent "forgetting" can be understood as *retrieval failures*—cases in which the information is in your memory, but you fail to locate it.

retrieval The process of searching for a memory and finding it.

tip-of-the-tongue (TOT) effect The condition in which one remains on the verge of retrieving a word or name but continues to be unsuccessful.

Partial Retrieval

Retrieval failure can be documented in many ways—including the fact that sometimes we remember *part* of the information we're seeking, but we can't recall the rest. This pattern can arise in many circumstances, but it's most clearly evident in the phenomenon psychologists call the **tip-of-the-tongue (TOT) effect.**

Try to think of the word that means "to formally renounce the throne." Try to think of the name of the Russian sled drawn by three horses. Try to think of the word that describes someone who, in general, does not like other people. Chances are that, in at least one of these cases, you found yourself in a frustrated state: certain you know the word but unable to come up with it. The word was, as people say, right on the "tip of your tongue."

People who are in the so-called TOT state can often remember roughly what the word sounds like—and so, when they're struggling to recall *abdicate,* they might remember *abrogate* or *annotate* instead. Likewise, they can often recall what letter the word begins with, and how many syllables it has, even though they can't recall the word itself (A. Brown, 1991; R. Brown & McNeill, 1966; Harley & Bown, 1998; L. James & Burke, 2000; B. Schwartz, 1999).

Similar results have been obtained when people try to recall specific names—for example, what is the capital of Nicaragua? Who was the main character in the movie *The*

Matrix? In response to these questions, people can often recall the number of syllables in the target name and the name's initial letter, but not the name itself (Brennen, Baguley, Bright, & Bruce, 1990; Yarmey, 1973). They also can often recall related material, even if they can't remember the target information. (Thus, they might remember Morpheus, but not the main character, from *The Matrix;* the main character, of course, was *Neo*. And the Russian sled is a *troika;* it's a *misanthrope* who doesn't like other people; Nicaragua's capital is *Managua*.)

People in the TOT state cannot recall the target word, but the word is certainly in their memory. If it weren't, they wouldn't be able to remember the word's sound, or its starting letter and syllable count. What's more, people often *recognize* the word when it's offered to them ("Yes! That's it!"). This is, therefore, unmistakably a case of retrieval failure—the information is preserved in storage, but for various reasons it is inaccessible.

Effective Retrieval Cues

retrieval cue A hint or signal that helps one to recall a memory.

retrieval paths The mental connections linking one idea to the next that people use to locate a bit of information in memory.

Retrieval failure is also clearly the problem whenever you seem to have forgotten something, but then recall it once you're given an adequate **retrieval cue.** A clear illustration of this pattern often arises when someone returns to his hometown after a long absence. This return can unleash a flood of recollection, including the recall of many details the person thought he'd forgotten long ago. Since these memories do surface, triggered by the sights and sounds of the hometown, there's no doubt about whether the memories were established in the first place (obviously, they were) or lost from storage (obviously, they weren't). Only one explanation is possible, therefore, for why the memories had been unavailable for so many years prior to the person's return to his hometown. They were in memory, but not findable—exactly the pattern we call retrieval failure.

Why do some retrieval cues (but not others) allow us to locate seemingly long-lost memories? One important factor is whether the cue re-creates the context in which the original learning occurred. This is obviously the case in returning to your hometown—you're back in the context in which you had the experiences you're now remembering. But the same broad point can be documented in the lab; and so, for example, if an individual focused on the sounds of words while learning them, then she would be well served by reminders that focus on sound ("Was there a word on the list that rhymes with *log?*"); if she focused on meaning while learning, then the best reminder would be one that again draws her attention toward meaning ("Was one of the words a type of fruit?"; R. Fisher & Craik, 1977).

The explanation for this pattern lies in our earlier discussion of memory connections. Learning, we suggested, is essentially a process of creating (or strengthening) connections that link the to-be-remembered material to other things you already know. But what function do these connections serve? When the time comes to recall something, the connections serve as **retrieval paths**—routes that lead you back to the desired information. Thus, if you noticed in a movie that Jane's smile caused Tarzan to howl, this will create a link between your memory of the smile and your memory of the howl. Later on, thinking about the smile will bring Tarzan's howl into your thoughts—and so your retrieval is being guided by the connection you established earlier.

On this basis, let's think through what would happen if a person studied a list of words and focused, say, on the sound of the words. This focus would establish certain connections—perhaps one between *dog* and *log,* and one between *paper* and *caper.* These connections will be useful if, later, this person is asked questions about rhymes.

If she's asked, "Was there a word on the list that rhymes with *log?*" the connection now in place will guide her thoughts to the target word *dog.* But the same connection will play little role in other situations. If she's asked, "Did any of the words on the list name animals with sharp teeth?" the path that was established during learning—from *log* to *dog*—is much less helpful; what she needs with this cue is a retrieval path leading from *sharp teeth* to the target.

The impact of these same retrieval cues would be different, though, if the person had thought about meaning during learning. This focus would have created a different set of connections—perhaps one between *dog* and *wolf.* In this case, the "rhymes with *log?*" cue would likely be ineffective, because the person has established no connection with *log.* A cue that focused on meaning, however, might trigger the target word.

Overall, then, an effective retrieval cue is generally one that takes advantage of an already established connection in memory. We've worked through this issue by pointing to the difference between meaning-based connections and sound-based connections, but the same point can be made in other ways. In one experiment, the researchers asked deep-sea divers to learn various materials. Some of the divers learned the material while sitting on land by the edge of the water. Others learned the material while 20 feet underwater, hearing the material via a special communication set. Within each of these two groups, half of the divers were then tested while above water, and half were tested below (Godden & Baddeley, 1975).

Imagine that you're a diver in the group that learned while underwater. In this setting, the world has a different look and feel than it does above water: The sound of your breathing is quite prominent; so is the temperature. As a result, you might end up thinking about your breathing (say) during learning, and this will likely create memory connections between these breathing thoughts and the materials you're learning. If you are then back underwater at the time of the memory test, the sound of your breathing will again be prominent, and this may lead you back into the same thoughts. Once thinking these thoughts, you will benefit from the memory connection linking the thoughts to the target materials—and so you'll remember the materials. In contrast, if you're on land during the memory test, then the sound of breathing is absent, and so these thoughts won't be triggered and the connections you established earlier will have no influence.

We might therefore expect the divers who learned underwater to remember best if tested underwater; this setting increases their chances of benefiting from the memory connections they established during learning. Likewise, the divers who learned on land should do best if tested on land. And that's exactly what the data show (Figure 8.10).

Related examples are easy to find. Participants in one study were asked to read an article similar to those they routinely read in their college classes; half read the article in a quiet setting, and half read it in a noisy environment. When tested later, those who read the article in quiet did best if they were tested in quiet; those who read it in a noisy environment did best if tested in a noisy setting (Grant et al., 1998). In both cases, participants showed the benefit of being able to use, at time of retrieval, the specific connections established during learning.

In case after case, then, it's helpful, at the time of memory retrieval, to return to the context of learning. Doing this will encourage some of the same thoughts that were in place during learning, and so will allow you to take advantage of the connections linking those thoughts to the target material. This broad pattern is referred to as a benefit of **context reinstatement**—a benefit of re-creating the state of mind you were in during learning.

context reinstatement A way of improving retrieval by re-creating the state of mind that accompanied the initial learning.

Let's also note that, in these experiments, the physical setting (noisy or not; underwater or above) seems to have a powerful influence on memory. However,

8.10 SCIENTIFIC METHOD: Is memory enhanced when the recall situation is similar to the encoding situation?

Method

1. One group of divers learned a word list on land. Another group of divers learned a word list underwater.

2. Each group was tested in both environments for recall of the list items.

Results

Divers who learned underwater recalled more words underwater, and those who studied on land tested better on land.

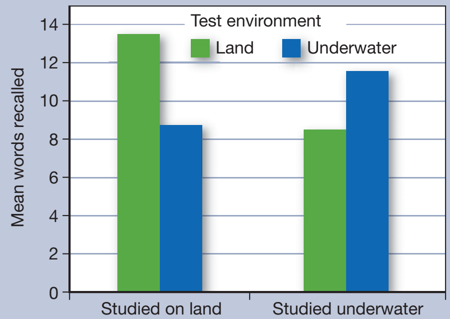

CONCLUSION: Information is best recalled in the environment where it is learned.

SOURCE STUDY: Godden & Baddeley, 1975

evidence suggests that the physical setting matters only indirectly: A return to the physical circumstances of learning does improve recollection, but only because this return helps re-create the mental context of learning—and it's the mental context that matters. This was evident, for example, in a study in which participants were presented with a long list of words. One day later, the experimenter brought the participants back for an unexpected recall test that took place in either the same room or a different one (one that differed in size, furnishings, and so on, from the context of learning). Not surprisingly, recall was better for those who were tested in the same physical environment—documenting, once again, the benefit of context reinstatement. Crucially, though, the investigator found a straightforward way of eliminating the difficulty caused by an environmental change: A different group of participants were brought to the new room; but just before the test, they were asked to think about the room in which they had learned the lists—what it looked like, how it made them feel. By doing so, they mentally re-created the old environment for themselves; on the subsequent recall test, these participants performed just as well as those who were tested in their original room (S. Smith, 1979; S. Smith & Vela, 2001; Figure 8.11). Apparently, then, what matters for retrieval is your mental perspective, not the room you're sitting in. If you change the physical context without changing your mental perspective, the physical relocation has no effect.

8.11 Context reinstatement for students These students are probably forming connections between the material they're learning and library-related cues. To help themselves recall this material later on, they'll want to think about what the library looked like and how they felt in that environment.

Encoding Specificity

The effectiveness of context reinstatement also tells us something important about how materials are recorded in memory in the first place. When people encounter some stimulus or event, they *think about* this experience in one way or another; and as we've described, this intellectual engagement serves to connect the new experience to other thoughts and knowledge. We've been discussing how these connections serve as retrieval paths, helping people to recall the target information, but let's now add that this is possible only because those connections are themselves part of the memory record. Thus, continuing an earlier example, if people see the word *dog* and think about what it rhymes with, what ends up being stored in memory is not just the word. What's stored must be the word plus some record of the connections made to rhyming words—otherwise, how could these connections influence retrieval? Likewise, if people see a picture and think about what it means, what's stored in memory is not just the picture, but a memory of the picture together with some record of the connections to other, related ideas.

In short, what's placed in memory is not some neutral transcription of an event. Instead, what's in memory is a record of the event *as understood from a particular perspective* or perceived within a particular context. Psychologists refer to this broad pattern as **encoding specificity**—based on the idea that what's recorded in memory is not just a "copy" of the original, but is instead *encoded from* the original (in other words, it's translated into some other form) and is also quite *specific* (and so represents the material *plus* your thoughts and understanding of the material; Tulving & Osler, 1968; Tulving & Thomson, 1973; also Hintzman, 1990).

encoding specificity The hypothesis that when information is stored in memory, it is not recorded in its original form but translated ("encoded") into a form that includes the thoughts and understanding of the learner.

This specificity, in turn, has powerful effects on retrieval—that is, on how (or whether) the past is remembered. For example, participants in one study read target words (e.g., *piano*) in either of two contexts: "The man lifted the piano" or "The man tuned the piano." These sentences led the participants to think about the target word in a particular way, and it was then this line of thinking that was encoded into each person's memory. Thus, continuing the example, what was recorded in memory was the idea of "piano as something heavy" or "piano as a musical instrument." This difference in memory content became clear when participants were later asked to recall the target words. If they had earlier seen the "lifted" sentence, then they were quite likely to recall the target word if given the hint "something heavy." The hint "something with a nice sound" was much less effective. But if participants had seen the "tuned" sentence, the result reversed: Now the "nice sound" hint was effective, but the "heavy" hint was not (Barcklay, Bransford, Franks, McCarrell, & Nitsch, 1974). In both cases, the memory hint was effective only if it was congruent with what was stored in memory—just as the encoding specificity proposal predicts.

This notion of encoding specificity is crucial in many contexts. For example, imagine two friends who have an argument. Each person is likely to interpret the argument in a way that's guided by his own position—and so he'll probably perceive his own remarks to be clear and persuasive, and his friend's comments to be muddy and evasive. Later on, how will each friend remember the event? Thanks to encoding specificity, what each person places in memory is the argument *as he understood it*. As a result, we really can't hope for a fully objective, impartial memory, one that might allow either of the friends to think back on the argument and perhaps reevaluate his position. Instead, each will, inevitably, recall the argument in a way that's heavily colored by his initial leaning.

MEMORY GAPS, MEMORY ERRORS

The processes we've been discussing—acquisition, storage, and retrieval—function extremely well in a huge range of circumstances. As a result, each of us can learn an enormous quantity of information, store that information for a long time, and then swiftly retrieve the information when we need it. But of course there are times when remembering is less successful. Sometimes we try to remember an episode but simply draw a blank. Sometimes we recall something, but with no conviction that we're correct: "I think it happened on Tuesday, but I'm not sure." And sometimes our memories fail us in another way: We recall a past episode, but it turns out that our memory is mistaken. Perhaps details of the event were different from the way we recall them; perhaps our memory is altogether wrong, misrepresenting large elements of the original episode. Why, and how often, do these memory failures occur?

Forgetting

There are many reasons why we sometimes cannot recall past events. In many cases, as we've noted, the problem arises because we didn't learn the relevant information in the first place! In other cases, though, we learn something—a friend's name, the lyrics to a song, the content of the Intro Bio course—and can remember the information for a while; but then, sometime later, we're unable to recall the information we once knew. What produces this pattern?

retention interval The time that elapses between learning and retrieval.

forgetting curve The graphic pattern representing the relationship between measures of learning and the length of the retention interval: As the retention interval gets longer, memory decreases.

One clue comes from the fact that it's almost always easier to recall recent events (e.g., yesterday's lecture or this morning's breakfast) than it is to recall more distant events (a lecture or a breakfast 6 months ago). In technical terms, recall decreases, and forgetting increases, as the **retention interval** (the time that elapses between learning and retrieval) grows longer and longer.

This simple fact has been documented in many studies; indeed, the passage of time seems to work against our memory for things as diverse as past hospital stays, our eating or smoking habits in past years, car accidents we experienced, our consumer purchases, and so on (Jobe, Tourangeau, & Smith, 1993). The classic demonstration of this pattern, though, was offered more than a century ago by Hermann Ebbinghaus (1850–1909). Ebbinghaus systematically studied his own memory in a series of careful experiments, examining his ability to retain lists of nonsense syllables, such as *zup* and *rif.* (Ebbinghaus relied on these odd stimuli as a way of making sure he came to the memory materials with no prior associations or links; that way, he could study how learning proceeded when there was no chance of influence from prior knowledge.) Ebbinghaus plotted a **forgetting curve** by testing himself at various intervals after learning (using different lists for each interval). As expected, he found that memory did decline with the passage of time. However, the decline was uneven; it was sharpest soon after the learning and then became more gradual (Ebbinghaus, 1885; Figure 8.12).

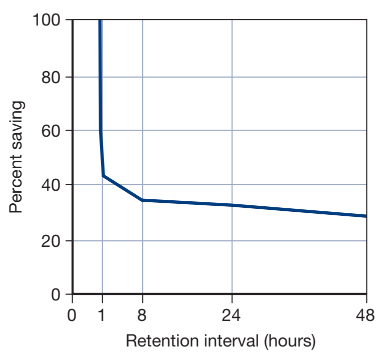

8.12 Forgetting curve The figure shows retention after various intervals since learning. Retention is here measured in percent saving—that is, the percentage decrease in the number of trials required to relearn the list after an interval of no practice. If the saving is 100%, then retention is perfect; no trials to relearn are necessary. If the saving is 0%, there's no retention at all; it takes just as many trials to relearn the list as it took to learn it initially.

There are two broad ways to think about the effect of retention interval. One perspective emphasizes the passage of time itself—based on the idea that memories *decay* as time passes, perhaps because normal metabolic processes wear down the memory traces until they fade and finally disintegrate. A different perspective suggests that time itself isn't the culprit. What matters instead is *new learning*—based on the idea that new information getting added to long-term memory somehow disrupts the old information that was already in storage. We'll need to sort through why this disruption might happen; but notice that this perspective, too, predicts that longer retention intervals will lead to more forgetting—because longer intervals provide more opportunity for new learning and thus more disruption from the new learning.

Which perspective is correct? Is forgetting ultimately a product of the passage of time, or a product of new learning? The answer is *both.* The passage of time, by itself, does seem to erode memories (e.g., E. Altmann & Gray, 2002; C. Bailey & Chen, 1989; Wixted, 2004); but the effect of new learning seems larger. For example, Baddeley and Hitch (1977) asked rugby players to recall the names of the other teams they had played against over the course of a season; the researchers then systematically compared the effect of *time* with the effects of *new learning*. To examine the effects of time, Baddeley and Hitch capitalized on the fact that not all players made it to all games (because of illness, injuries, or schedule conflicts). These differences allowed them to compare players for whom "two games back" means 2 weeks ago, to players for whom "two games back" means 4 weeks ago. Thus, they were able to look at the effects of time (2 weeks vs. 4) with the number of more recent games held constant. Likewise, to examine the effects of new learning, these researchers compared (say) players for whom the game a month ago was "three games back" to players for whom a month ago means "one game back." Now we have the retention interval held constant, and we can look at the effects of intervening events. In this setting, Baddeley and Hitch report that the mere passage of time accounts for very little; what really matters is the number of intervening events—just as we'd expect if intervening learning, and not decay, is the major contributor to forgetting (Figure 8.13). (For other—classic—data on this issue, see Jenkins & Dallenbach, 1924; for a more recent review, see Wixted, 2004.)

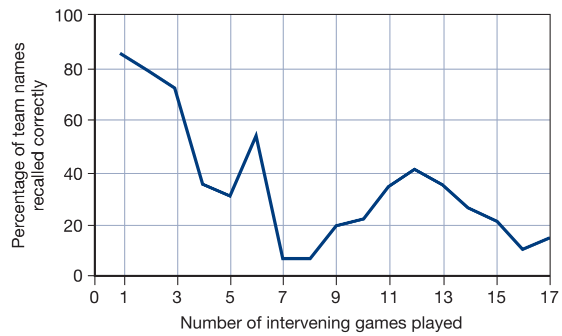

8.13 Forgetting from interfering events Members of a rugby team were asked to recall the names of teams they had played against. Their performance was heavily influenced by the number of games that intervened between the game to be recalled and the attempt to remember.

An effect of new learning undoing old learning can also be demonstrated in the laboratory. In a typical study, a control group learns the items on a list (A) and then is tested after a specified interval. The experimental group learns the same list (A), but they must also learn the items on a second list (B) during the same retention interval. The result is a marked inferiority in the performance of the experimental group. List B seems to interfere with the recall of list A (Crowder, 1976; McGeoch & Irion, 1952).

Of course, not all new learning produces this disruption. No interference is observed, for example, between dissimilar sorts of material—and so learning to skate doesn't undo your memory for irregular French verbs. In addition, if the new learning is *consistent* with the old, then it certainly doesn't cause disruption; instead, the new learning actually *helps* memory. Thus, learning more algebra helps you remember the algebra you mastered last year; learning more psychology helps you remember the psychology you've already covered.

Memory Intrusions

We still need to ask *why* new learning seems sometimes to disrupt old. Why can't the newly acquired information peacefully coexist with older memories? In fact, there are several reasons. In some cases, the new information simply sits side by side with old memories, creating a danger that you'll get mixed up about which is which—recalling the newer information when you're trying to come up with the older. In other cases, the new information may literally *replace* the old memory, much as you delete an old version of a paper from your computer's hard drive once you've created a newer, updated version.

In most experiments, it's difficult to distinguish these two possibilities—that is, to tell whether the new information is merely competing with the old information, or whether it has literally replaced the old information. In either case, though, the new material will lead to **intrusion errors**—mistakes about the past in which other information is mixed into (*intrudes into*) your recall. These intrusions are often small (so that you recall having a cheese sandwich when you really had a salad) but can sometimes be quite large: People may confidently, vividly recall a past event that never took place at all.

intrusion errors Memory mistakes in which elements that were not part of the original information get mixed into ("intrude" into) someone's recall.

8.14 Eyewitness memory Considerable research has been done on the question of how accurately and how completely witnesses (or victims) to crimes will remember what they've experienced.

misinformation effect The result of a procedure in which, after an experience, people are exposed to questions or suggestions that misrepresent what happened. The term refers to people's tendency to include the misinformation as part of their recall of the original experience.

THE MISINFORMATION EFFECT

Intrusion errors can arise in many ways. Often, though, the intrusion involves information about an event that you learned *only after the event was over*. For example, imagine that you witness a crime and see the thief flee in a blue car. The next day, you read a newspaper account of the same crime and learn that another witness has reported that the thief fled in a *green* car (Figure 8.14). How will this experience influence your memory? A number of experiments have examined this issue by exposing participants to an event and then giving them some misinformation about the event. In some studies, the misinformation comes from another person's report ("Here's the way another witness described the event . . ."). In other studies, the misinformation is contained within a leading question: Participants might be asked, for example, "Did you see the children getting on the school bus?" after seeing a video that showed no bus. In all cases, though, the effect is the same: This misinformation is often incorporated into the participants' memory, so that they end up misremembering the original event, mistakenly including the bits suggested after the fact by the experimenter.

The errors produced by the **misinformation effect** can actually be quite large. Participants have in fact been led to remember buses that weren't actually present in an event as well as whole buildings that didn't exist (Loftus, 2003). Indeed, with slight variations of this technique, participants have been led to recall entire events that never occurred. In one study, participants were asked to recall a time they had been at an outdoor wedding and had accidentally knocked over the punchbowl, spilling it onto the bride's parents. With suggestive questioning, the researcher led 25% of the participants to "remember" this nonexistent episode (I. Hyman, Husband, & Billings, 1995). In similar experiments, participants have been led to recall a nonexistent episode in which they were hospitalized, or a hot-air balloon ride that really never happened (Figure 8.15), or a (fictitious) event in which they were the victim of a vicious animal attack (Loftus, 2003, 2004; also Chrobak & Zaragoza, 2008; Geraerts et al., 2006; Geraerts et al., 2009; Geraerts et al., 2007; Laney et al., 2008; and many more).

Errors like these are easily documented in the laboratory, but can also be observed in real-life settings. We are, after all, often exposed to alternate versions of events we've experienced—for example, whenever we discuss a shared experience with a friend, and the friend recalls things differently from the way we do. Moreover, the leading questions examined in much of this research are modeled directly on the questions sometimes asked in law enforcement investigations involving adults ("Did you see the gun?") as well as children ("When did Uncle Seth touch you?"; cf. Bruck & Ceci, 1999; Ceci &

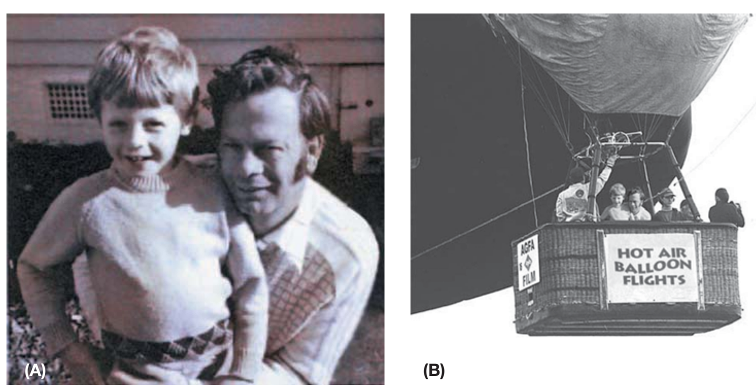

8.15 The balloon ride that never was In this study, participants were shown a faked photo (as in B) created from a real childhood snapshot (as in A). With this prompt, many participants were led to a vivid, detailed recollection of the balloon ride—even though it never occurred!

Bruck, 1995; Melnyk, Crossman, & Scullin, 2007; Westcott, Davies, & Bull, 2002). As a consequence, we can readily find examples in which the memory errors made in the laboratory are mirrored by errors outside of the lab—including settings (like law enforcement) in which the memory mistakes are deeply troubling and potentially very costly.

INTRUSIONS FROM SCHEMATIC KNOWLEDGE

Intrusion errors, interfering with our memory of the past, can also come from another source—because sometimes we blur together our recollection of an episode with our broader knowledge about the world. Classic data on this topic come from studies performed by the British psychologist Frederic Bartlett more than 75 years ago. Bartlett presented British research participants with stories drawn from Native American folklore; and for these participants, many elements of these stories seemed strange. In the participants' recollection of these stories, though, the tales became less strange. Parts of the tales that had made no sense to them (such as the supernatural elements) either were left out of their recall or were reinterpreted along more familiar lines. Similarly, participants often added elements so that plot events that had initially seemed inexplicable now made sense to them (Bartlett, 1932).

What happened here? Bartlett's participants quite naturally tried to understand these stories by relating them to other things they knew and understood. In the process, they ended up creating connections in their memories, weaving together the story elements with various aspects of their own knowledge about the world. This weaving together helped the participants comprehend the materials they were hearing by linking the unfamiliar materials to a more familiar framework. But, at the same time, this weaving caused problems later on, because it made it difficult for participants to keep track of which elements were actually in the stories and which were merely associated with the story via their understanding of it. This is what produced the memory errors.

Other studies have replicated Bartlett's findings, showing in many contexts that memory is strongly affected by an individual's conceptual framework. For example, participants in one study were told about a person's visit to the dentist and then later asked to recall what they had heard. Many participants remembered being told about the patient checking in with the receptionist and looking at a magazine in the waiting room, even though these details were not mentioned in the original account (G. Bower, Black, & Turner, 1979). In a different experiment (which we first met in the Prologue), participants waited briefly in a professor's office and, seconds later, were asked to recall the contents of the office. One-third of the individuals "remembered" seeing books in the office, even though none were present (Brewer & Treyens, 1981). In this case the error is a substantial one (bookshelves are large; the participants were actually in the office; the recollection took place just moments after leaving the office); but, again, it is entirely in line with participants' expectations of what "should" be in a professor's office.

In all these examples, memory is strongly affected by the research participants' broad knowledge of the world and by the conceptual framework they bring to the situation. Following Bartlett, many psychologists describe these frameworks as **schemas**—mental representations that summarize what we know about a certain type of event or situation. Schemas reflect the simple fact that many aspects of our experience are redundant—professors' offices do tend to contain many books, patients visiting the dentist do generally check in with a receptionist—and schemas provide a convenient summary of this redundancy.

schema An individual's mental representation that summarizes her knowledge about a certain type of event or situation.

Let's also be clear that a reliance on schematic knowledge is generally a good thing. When we encounter an event—whether it's a trip to the dentist or a story from another culture—we seek to understand it by relating it to a schematic frame. This helps us find meaning in our experience, and it also fills in the "gaps" that result from our failing to notice this or that detail. Then, when we try to remember the event, we rely on the same schema. And here, too, this strategy can help us by allowing us to make reasonable assumptions about what probably occurred, thus filling any gaps in what we recall. Even so, this reliance on schematic knowledge can lead to substantial memory errors. In particular, it can lead us to remember the past as being more regular and more orderly than it actually was.

INTRUSIONS FROM SEMANTIC ASSOCIATIONS

DRM paradigm A common procedure for studying memory, in which participants read and then immediately recall a list of related words, but the word providing the "theme" for the list is not included.

Intrusion errors can be documented in many settings, including settings that seem designed to encourage memory accuracy. For example, intrusion errors can be observed even with simple stimuli, short retention intervals, and instructions that warn participants about the kinds of memory errors they're likely to make. Even in these settings, participants make a surprising number of memory mistakes.

Evidence for these points comes from many studies that draw on the so-called **DRM paradigm**—named in honor of Deese, Roediger, and McDermott, the researchers who developed it (Deese, 1957; Roediger & McDermott, 1995, 2000; also Blair, Lenton, & Hastie, 2002; Brainerd, Yang, Reyna, Howe, & Mills, 2008; Sugrue & Hayne, 2006; etc.). In this procedure, participants hear a list of words, such as *bed, rest, awake, tired, dream, wake, snooze, blanket, doze, slumber, snore, nap, peace, yawn, drowsy*. Then, immediately after hearing the list, participants are asked to recall as many of the words as they can.

All of the words in this list are semantically associated with the word *sleep*, and the presence of this theme helps memory and makes the words on the list easy to recall. But as it turns out, *sleep*—the root of the list—isn't included in the presentation. Still, participants spontaneously make the connection between the list words and this associated word—and this almost invariably leads to a memory error. When the time comes for recall, participants are extremely likely to remember that they heard *sleep*. In fact, they're just as likely to recall this word as they are to recall the actual words on the list (Figure 8.16)!

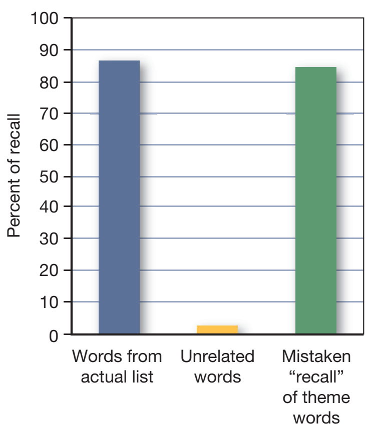

8.16 The effects of the DRM paradigm Because of the theme uniting the list, participants can remember almost 90% of the words they encountered. However, they're just as likely to "recall" the list's theme word—even though it was not presented.

Notice once again that participants' background knowledge both helps and hurts them. In the DRM task, the participants' knowledge helps them link together the list words according to a theme, and this strongly aids recall of the words. But the same knowledge leads to a remarkably high level of false recall—by powerfully encouraging recall for words that were never presented.

MISPLACED FAMILIARITY

familiarity A general sense that a certain stimulus has been encountered before.

recollection Recall of the context in which a certain stimulus was encountered.

One more mechanism plays a key role in producing intrusion errors in memory, and, with that, in encouraging misremembering of the past. To understand this mechanism, let's start with the fact that the memory processes that make a stimulus seem *familiar* are different from those that help us figure out *why* the stimulus feels familiar. As a result, sometimes the first process succeeds, so we correctly realize that a stimulus is familiar, but the second process fails, so we make a mistake about the source of that familiarity.

Familiarity (a general sense that a stimulus has been encountered before) and **recollection** (recall of the context in which a stimulus was encountered) can be distinguished in many ways. For a start, there's an obvious subjective difference—in other words, these two types of memory *feel* different from each other—and people can reliably tell whether they "remember" a prior event (and so have some recollection) or whether they don't remember the event, but just "know" that it happened (and so, apparently, are relying on familiarity—Rajaram, 1993; Tulving, 1985; also Aggleton & Brown, 2006). These two types of memory are also promoted by different types of strategies—so that some approaches to a stimulus or event are especially helpful for establishing a sense of familiarity; different strategies are needed for establishing the sort of memory that will later on lead to recollection.

Familiarity and recollection can also be distinguished biologically. During learning, activity in the rhinal cortex seems crucial for establishing a sense of familiarity; and so higher levels of activity in this brain area, during the initial encounter with a stimulus, are associated with greater likelihood of familiarity later on. In contrast, areas in and around the hippocampus seem essential for establishing a basis for recollection; higher levels of activity in these regions, during learning, are associated with greater likelihood of subsequent recollection (e.g., Davachi, Mitchell, & Wagner, 2003; Davachi & Dobbins, 2008; Ranganath et al., 2003; Figure 8.17).

Then, during retrieval, familiarity and recollection both rely on the prefrontal cortex; but they depend on clearly different areas within this cortex (e.g., Diana, Yonelinas, & Ranganath, 2007; Dobbins, Foley, Schacter, & Wagner, 2002; Kahn, Davachi, & Wagner, 2004; Rugg & Curran, 2007; A. Wagner et al., 2005). Therefore, the brain state of someone who remembers seeing a stimulus earlier is distinct from the brain state of someone who doesn't remember the earlier encounter but still feels that the stimulus is familiar.

This distinction between familiarity and recollection has many consequences, including the possibility of one process *succeeding* while the other *fails*. In fact, this pattern is easily detectable in everyday experience—when, for example, people have the frustrating experience of seeing someone, immediately knowing that the person is familiar, but not being able to figure out *why* the person is familiar. In that situation, you ask yourself: "Where do I know that guy from? Does he maybe work at the grocery store? At the shoe store? Where?" Here the familiarity process has succeeded—and so

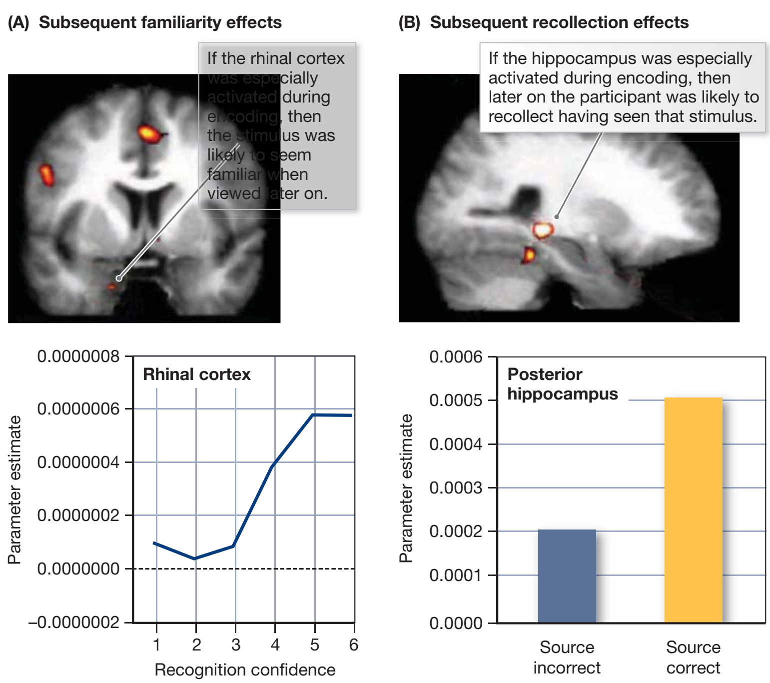

8.17 Familiarity vs. source memory In this study, researchers tracked participants' brain activity during encoding and then analyzed the data according to what happened later, when the time came for retrieval.

8.18 A photo lineup On TV, crime victims view a live lineup, but it's far more common in the United States for the victim (or witness) to see a "photo lineup" like this one. The victim (or witness) is told that the perpetrator may or may not be present and is asked to pick out the perpetrator if he's there. Unfortunately, victims sometimes pick the wrong fellow, and this error is more likely if the suspect is familiar to the victim for some reason other than the crime.

you're sure the person is indeed familiar—but the recollection process, allowing you to attribute the familiarity to a specific source, has let you down.

A different problem can also arise, and brings us back to our main agenda—the ways in which new learning can intrude on (and thus disrupt) older learning. Here's the situation: You notice that a stimulus (perhaps a face, or a place) seems familiar, and you think you know why—but you're simply mistaken, and you attribute the familiarity to the wrong source. This phenomenon can arise in a variety of circumstances outside the lab, and researchers have re-created one of those circumstances in their studies: In the experimental procedure, participants witness a staged crime. Two days later, they're shown "mug shots" of individuals who supposedly had participated in the crime. But, as it turns out, the people in these photos are different from the people who were actually involved in the crime; no mug shots are shown for the truly "guilty" individuals. Finally, after a few more days, the participants are shown a lineup like the one in Figure 8.18 and asked to select from this lineup the individuals seen in step one—namely, the original staged crime.

In this procedure, participants correctly realize that one of the people in the lineup looks familiar, but they're often confused about the source of the familiarity. They falsely believe they had seen his face in the original "crime," when in truth they'd seen it only in a subsequent photograph. In fact, the likelihood of this error is quite high in some experiments, and sometimes more than a quarter of the participants (falsely) select from the lineup an individual they'd seen only in the mug shots (E. Brown, Deffenbacher, & Sturgill, 1977; also D. Davis, Loftus, Vanous, & Cucciare, 2008).

TRYING TO AVOID MEMORY ERRORS

We've now discussed many experiments in which participants made substantial memory errors. Similar errors, as we've frequently mentioned, can also be documented outside the lab—including errors in highly consequential settings. In fact, evidence suggests that eyewitness errors in the American court system may account for more false convictions than all other causes combined (Connors, Lundregan, Miller, & McEwan, 1996).

These points invite two questions: First, is there anything we can do to avoid the errors—and thus to improve memory accuracy? Second, what do these errors imply for our overall assessment of memory? Should we perhaps put less trust in our memories than we generally do? Let's tackle these questions in turn.

Some steps do seem helpful in avoiding (or at least diminishing) memory error. For example, we noted earlier in the chapter that some memory problems involve retrieval failures, and so we can improve memory by means of instructions or strategies that promote retrieval—such as trying, at the time of recall, to reinstate the psychological context of learning (for a real-world use of this procedure, see Fisher & Schreiber, 2007). Likewise, since attention to meaning seems an effective way to memorize, we can improve memory by encouraging people to think more deeply about the materials they're encountering; this will promote both understanding and memory.

People have also suggested more exotic means of improving memory—but these seem less helpful. For example, some people have proposed the use of *hypnosis* as an aid to memory, based on the idea that someone—for example, an eyewitness to a crime—

can be hypnotized, given the suggestion that she's back at a certain time and place, and asked to tell what she sees. On the surface, the results of this procedure—in a police station or in laboratory studies—are quite impressive. A hypnotized witness mentally returns to the scene of the crime and seems able to recall exactly what the various participants said; a hypnotized college student mentally returns to childhood and appears to relive his sixth birthday party with childlike glee.

Careful studies reveal, however, that hypnosis doesn't improve memory. Descriptions of crimes or childhood reports elicited under hypnosis often turn out to be false when checked against available records (Lynn, Lock, Myers, & Payne, 1997; Spanos, C. Burgess, M. Burgess, Samuels, & Blois, 1999; also Figure 8.19; for more on hypnosis, see Chapter 6).

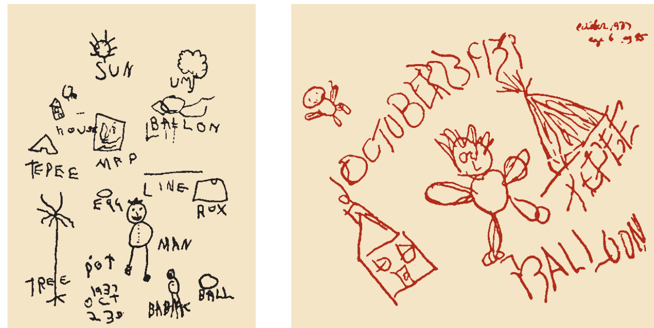

(A) Drawings done at age 6

(B) Drawings done by hypnotized adult told that he was 6 years old

8.19 Hypnotic age regression In one study, participants were asked to draw a picture while mentally "regressed" to age 6. At first glance, their drawings (an example is shown in A) looked remarkably childlike. But when compared to the participants' own drawings made at that age (an example is shown in B), it's clear that the hypnotized adults' drawings were much more sophisticated. They represent an adult's conception of what a childish drawing is rather than being the real thing.

Likewise, certain drugs are sometimes proposed as improving memory—but here too the actual benefits are small. Some of the drugs used to promote memory (e.g., sodium amytal) are sedatives, so they put an individual in a less guarded, less cautious state of mind. This state does allow the person to report more about the past—but not because she remembers more. Instead, in the relaxed state, the person is just more willing to talk and less likely to discriminate between genuine memories and fantasy. As a result, the person who has taken the drug will spin out a mix of recollection and fiction that robs their "recall" of any value. What's more, this less guarded state leaves an individual more vulnerable to the effects of leading or misleading questions, which can further undermine memory accuracy.

This evidence is more encouraging for a different drug—ginkgo biloba—sometimes advertised as improving memory (and other aspects of intellectual functioning). Ginkgo has an entirely different effect from the sedatives just mentioned, and it does improve memory for certain groups of people. Specifically, ginkgo can help with some types of blood-circulation problems and can also reduce certain forms of inflammation. It can therefore help people whose mental functioning has been compromised by specific physical maladies; this includes patients suffering from Huntington's disease or Alzheimer's disease. There's little evidence, however, that ginkgo improves the memory of healthy individuals; that is, there's no reliable effect for people with no circulatory problems or inflammation (Gold, Cahill, & Wenk, 2002; McDaniel, Maier, & Einstein, 2002).

Finally, there's one more—less exotic—step that people have tried to improve their memories: They have tried simply *being careful* in their recollection. This effort begins with the fact that we feel sure about some of our memories ("I'm certain that he's the guy who robbed me!") but more tentative about others ("I think she said it was size 6, but I'm not sure"). The obvious strategy, therefore, is to rely only on the memories we feel sure about and to be more cautious otherwise.

Surprisingly, though, this commonsense strategy offers little benefit. Many studies have compared the accuracy of memories people are certain about with the accuracy of memories they're not sure of. These studies often find a relationship in which confident memories are slightly more likely to be correct than unconfident ones. But the relationship is weak, and some studies have found no relationship at all (e.g., Bernstein & E. Loftus, 2009; Douglas & Steblay, 2006; Reisberg, 2010; Semmler & Brewer, 2006; Wells, Olson, & Charman, 2002). As a result, if we rely on our confidence in deciding which memories to trust, we'll regularly accept false memories and reject true ones.

Memory: An Overall Assessment

It seems, therefore, that false memories are essentially undetectable and unavoidable. In addition, we've seen that the errors in our recollection can be large and consequential. Does all of this mean that we should lament the poor quality of human memory? The answer to this question is an emphatic *no.* It's certainly true that we sometimes remember less than we'd like (a common experience for students taking an exam). It's also true that our recollection is sometimes mistaken—so the past as it actually unfolded is rather different from the past we remember. Even so, there's reason to believe our memories function in just the way we want them to.

How could this be? One point to bear in mind here is that, even with the memory errors we've discussed, our memories are *correct* far more often than not—so we usually remember the past accurately, in detail, and for a very long time. It's also important to highlight a point that has come up already—namely, that the mechanisms leading to memory error are mechanisms that *help us* most of the time; and so, in a sense, the errors are just the price we pay to gain other advantages. For example, errors in the misinformation paradigm arise (in part) because our memories are densely interconnected with each other; this is what allows elements to be transplanted from one remembered episode to another. But the connections from one memory to the next are, of course, there for a purpose: They're the retrieval paths that make memory search possible. Thus, to avoid the errors, we would need to restrict the connections—but if we did that, we'd lose the ability to locate our own memories within long-term storage!

The memory connections that lead to error also help us in other ways. Our environment, after all, is in many ways predictable—and it's enormously useful for us to exploit that predictability. There's little point in scrutinizing a kitchen to make sure there's a stove in the room, because in the vast majority of cases there is. So why take the time to confirm the obvious? Likewise, there's little point in taking special note that, yes, this restaurant does have menus and that, yes, people in the restaurant are eating and not having their cars repaired. These too are obvious points, and it would be a waste of time to give them special notice.

On these grounds, a reliance on schematic knowledge is a good thing. Schemas guide our attention to what's informative in a situation, rather than what's self-evident (e.g., Gordon, 2006); they also guide our inferences at the time of recall. If this use of schemas sometimes leads us astray, this may be a small price to pay for the gain in efficiency that schemas allow.

Finally, what about forgetting? This too may be a blessing in disguise, because sometimes it's to our advantage to remember less and forget more. For example, think about all the times in your life when you've been with a particular friend. These episodes are related to each other in an obvious way, so they're likely to become interconnected in your memory. This will cause difficulties if you want to remember which episode is which, and whether you had a particular conversation last Tuesday or the day before. But rather than lamenting this as an example of forgetting, we may want to celebrate what's going on here. Because of the "interference," all of the episodes will merge in your thoughts, so that what resides in memory is one integrated package containing all of your knowledge about your friend. This is, in fact, the way that much of your general knowledge is created! In other words, the same blurring together that makes it difficult to remember episodes also makes it possible to think in a general way, with a focus on what diverse experiences have in common rather than on what makes each experience unique. Without this blurring together, our capacity for thinking in general terms might be dramatically impaired.

It seems, then, that our overall assessment of memory can be rather upbeat. We've discussed a wide range of memory errors, but these errors are the exception rather than the rule. In addition, we've now seen that in most cases the errors are a by-product of mechanisms that otherwise help us—to locate our memories within storage, to be efficient in our contact with the world, and to form general knowledge. Thus, even with the errors, even with forgetting, it seems that human memory functions in a fashion that serves us well.

VARIETIES OF MEMORY

So far in this chapter, we've been discussing how memory functions and have given little attention to *what* was being remembered—and, to a large extent, this approach works well: The principles we've described apply equally to memory for word lists in the laboratory, for movies you've seen or songs you've heard, and for complex events in your everyday life. At the same time, it's also possible to distinguish different types of memory—each with its own operating principles and its own neural basis. Let's look at some of the crucial distinctions.

A Hierarchy of Memory Types

Many psychologists distinguish memory types in terms of a hierarchy like the one in Figure 8.20. On the left side of the hierarchy are the various forms of **explicit memory.** These are conscious memories—memories that you can describe if you choose—and they can usually be triggered by a direct question, such as "Do you know whether . . . ?" or "Do you recall the time when . . . ?" In contrast, **implicit memories** are remnants of the past that we may not recall at all, but they are (unconsciously) still with us, and we can detect these memories by the influence they still have on us. We'll consider some examples of implicit memories in a moment—but, in general, these memories cannot be revealed by direct questions; instead, they're usually revealed by some sort of indirect test.

explicit memory Conscious memories that can be described at will and can be triggered by a direct question.

implicit memory Memories that we may not recall consciously, but that are still demonstrable through an indirect test.

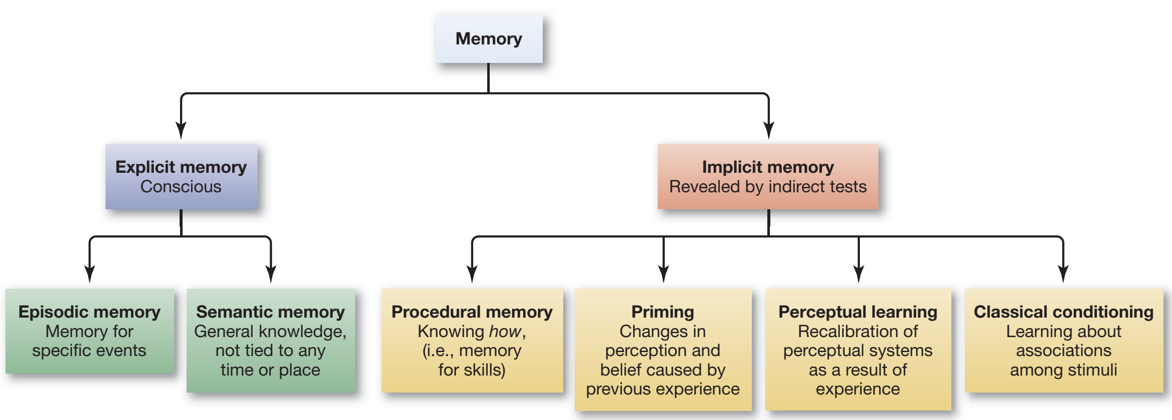

8.20 Hierarchy of memory types

Episodic and Semantic Memory

episodic memory Memory for specific events and experiences.

semantic memory Memory for facts (including word meanings); these memories are not tied to any specific time or place.

Both explicit and implicit memories can be subdivided further. Most of this chapter, in fact, has focused on just one type of explicit memory: **episodic memory.** This term refers to memory for specific events, including events outside the laboratory (e.g., the event of your 10th birthday) or inside (e.g., the event of memorizing a particular story). Just as important, though, is **semantic memory.** This is the memory that contains knowledge not tied to any time or place—your knowledge that London is the capital of England, that water is wet, that people become annoyed if you insult them. (Note that some information in semantic memory is concerned with semantics—including your memory for what the word *special* means, or what the opposite of *hot* is. Much of the information in this memory, however, is not specifically tied to semantics, and so some authors prefer to call it *generic memory,* or *generic knowledge*.)

Episodic and semantic memory can be distinguished on many grounds—including the specific brain areas that support each type of memory. This distinction is reflected in the fact that some forms of brain damage disrupt episodic memory but not semantic, and other forms do the reverse. For example, a patient known as Gene sustained a serious head injury in a motorcycle accident; the damage affected large areas of his frontal and temporal lobes, including his left hippocampus. As a result, he can recall no events at all from any time in his life. "Even when detailed descriptions of dramatic events in his life are given to him—such as the derailment, near his house, of a train carrying lethal chemicals that required 240,000 people to evacuate their homes for a week," Gene remembers nothing of this or any other event (D. Schacter, 1996, p. 150; Tulving, Schacter, McLachlan, & Moscovitch, 1988). But he does remember some things. He remembers that he owned two motorcycles and a car, he knows that his family has a summer cottage where he spent many weekends, and he recalls the names of classmates in a school photograph (D. Schacter, 1996). In short, Gene's episodic memory is massively disrupted, but his memory for generic information is largely intact.

Other patients show the reverse pattern. One woman, for example, suffered damage to the front portion of her temporal lobes as a result of encephalitis. As a consequence, she has lost her memory of many common words, important historical events, famous people, and even the fundamental traits of animate and inanimate objects. "However, when asked about her wedding and honeymoon, her father's illness and death, or other specific past episodes, she readily produced detailed and accurate recollections" (D. Schacter, 1996, p. 152).

Data like these make it clear that we need to distinguish between semantic and episodic memory. But these categories can themselves be subdivided. For example, people who have suffered brain damage sometimes lose the ability to name certain objects, or to answer simple questions about these objects (e.g., "Does a whale have legs?"). Often the problem is quite specific—and so some patients lose the ability to name living things but not nonliving things; other patients show the reverse pattern (Mahon & Caramazza, 2009). Indeed, sometimes the symptoms caused by brain damage are even more fine-grained: Some patients lose the ability to answer questions about fruits and vegetables, but they're still able to answer questions about other objects (living or nonliving). These data suggest that separate brain systems are responsible for different types of knowledge—and so damage to a particular brain area disrupts one type of knowledge but not others.

Possible Subdivisions of Episodic Memory

Plainly, then, we need to distinguish semantic memory from episodic, and we need to distinguish different types of semantic memory. Do we also need subdivisions within episodic memory? Some theorists believe we do, granting special status, for example, to *autobiographical memory*—the memory that defines, for each of us, who we are (e.g., Baddeley, Aggleton, & Conway, 2002; Cabeza & St. Jacques, 2007). Other theorists propose that *certain types of events* are stored in specialized memory systems—so that some authors argue for special status for *flashbulb memories*, and others suggest special status for memories for *traumatic events*.

FLASHBULB MEMORIES

Each one of us encounters a wide diversity of events in our lives. Many of these events are emotionally neutral (shopping for groceries, or buying a new Psychology textbook), while others trigger strong feelings (an especially romantic evening, or a death in the family). In general, emotional episodes tend to be better remembered—more vividly, more completely, and more accurately (e.g., Reisberg & Heuer, 2004)—and many mechanisms contribute to this effect. Among other points, emotional events are likely to be interesting to us, guaranteeing that we pay close attention to them; and we've already seen that attention promotes memory. Emotional events are also likely to involve issues or people we care about; this makes it likely that we'll readily connect the event to other knowledge (about the issues or the people)—and these connections, of course, also promote memory. In addition, the various biological changes that accompany emotion play a role—facilitating the process of memory consolidation (e.g., Buchanan & Adolphs, 2004; Dudai, 2004; Hamann, 2001; LaBar & Cabeza, 2006; LaBar, 2007).

Within the broad set of emotional memories, however, our memory for some events seems truly extraordinary for its longevity: People claim to remember these events, even decades later, "as if they happened yesterday." These especially vivid memories, called **flashbulb memories,** typically concern events that were highly distinctive and unexpected as well as strongly emotional. The most common examples involve emotionally *negative* events that triggered fear, or grief, or horror—such as the memory of an early morning phone call reporting a parent's death, or the memory of hearing about the attack on the World Trade Center in 2001 (Figure 8.21).

flashbulb memories Vivid, detailed memories said to be produced by unexpected and emotionally important events.

The clarity and longevity of flashbulb memories led psychologists R. Brown and Kulik (1977) many years ago to propose that we must have some special "flashbulb mechanism" distinct from the mechanisms that create other, more mundane memories.

8.21 Flashbulb memories (A) The classic example of a flashbulb memory is the assassination of John F. Kennedy in November 1963. Virtually all Americans (and most Europeans) who were at least 9 or 10 years old on that date still remember the day vividly. (B) The World Trade Center attack on September 11, 2001, is the sort of shocking and highly consequential event that seems very likely to create a flashbulb memory. Decades from now, people are likely to remember this day clearly.

The full pattern of evidence, however, suggests that there is no such special mechanism. As one concern, people usually talk with their friends about these remarkable events (i.e., you tell me your story, and I tell you mine), and it's probably this rehearsal, not some specialized mechanism, that makes these memories so long lasting. What's more, flashbulb memories—like other memories—are not immune to error: In fact, some flashbulb memories are filled with inaccuracies and represent the event in a fashion far from the truth (see, for example, M. Conway et al., 2009; Greenberg, 2004; Hirst et al., 2009; Luminet & Curci, 2009; Neisser, 1982a, 1986).

Moreover, the longevity of flashbulb memories may be less extraordinary than it seems, because other, more mundane, memories may also be extremely long lasting. One study tested people's memory for faces, asking in particular whether people could still identify photos of people they'd gone to high school with many years earlier (Bahrick et al., 1975; also see M. Conway, Cohen, & Stanhope, 1991; also see Bahrick & Hall, 1991; Bahrick, Hall, Goggin, Bahrick, & Berger, 1994). People who had graduated from high school 14 years earlier were still able to name roughly 90% of the faces; the success rate was roughly 75% for people who graduated 34 years earlier. A half-century after leaving high school, people could still name 60% of the faces (and it's unclear whether this small drop-off in accuracy reflects an erosion of memory or a more general decline in cognitive performance caused by normal aging). Clearly, therefore, even "non-flashbulb memories" can last for a very long time.

Flashbulb memories do seem remarkable—in their clarity, their durability, and (in some cases) their accuracy. But these attributes are likely to be the result of rehearsal plus the ordinary mechanisms associated with emotional remembering, and not a basis for claiming that flashbulb memories are somehow in a class by themselves.

MEMORY FOR TRAUMATIC EVENTS

There has been considerable controversy over a different proposal—the notion that memory for *traumatic events* might follow its own rules, different from the principles governing other types of episodic memory. Certainly, many of the principles we've discussed apply to traumatic memory, just as they apply to memories of other sorts. Thus, traumatic memories become harder to recall as time goes by; and they sometimes contain errors, just like all memories do. Traumatic memories are also better retained if they're rehearsed (thought about) once in a while. In these regards, traumatic memories seem quite similar to other sorts of episodic memory. The debate, however, is focused on a further issue—whether trauma memories are governed by a separate set of principles tied to how people might *protect themselves* from the painful recollection of horrific events.

8.22 Memory for traumatic events
There has been considerable debate over whether people have special mechanisms that protect them from recalling traumatic events.

Overall, how well are traumatic events remembered? If someone has witnessed wartime atrocities or has been the victim of a brutal crime, how fully will he remember these horrific events (Figure 8.22)? If someone suffers through the horrors of a sexual assault, will she be left with a vivid memory as a terrible remnant of the experience? The evidence suggests that traumatic events tend to be remembered accurately, completely, and for many years. Indeed, the victims of some atrocities seem plagued by a cruel enhancement of memory, leaving them with extra-vivid recollections of the awful event (see, for example, K. Alexander et al., 2005; Brewin, 1998; Goodman et al., 2003; McNally, 2003b; Pope, Hudson, Bodkin, & Oliva, 1998; S. Porter & Peace, 2007; Thomsen & Berntsen, 2009). In fact, in many cases these recollections can become part of the package of difficulties that lead to a diagnosis of *post-traumatic stress disorder,* and so they may be horribly disruptive for the afflicted individual (see Chapter 16).

There are, however, some striking exceptions to this broad pattern. Researchers have documented many cases in which people have suffered through truly extreme events but seem to have little or no recall of the horrors (see, for example, Arrigo & Pezdek, 1997). Thus, someone might be in a terrible car crash but have absolutely no recollection of the accident. Someone might witness a brutal crime but be unable to recall it just a few hours later.

How should we think about this mixed pattern? Let's start with the outcome that's by far more common—the person whose trauma is remembered all too well. This outcome is probably best understood in terms of the biological process of consolidation, using the hypothesis that this process is promoted by the conditions that accompany bodily arousal (Buchanan & Adolphs, 2004; Hamann, 2001). But what about the cases of the other sort—the person with no memory of the trauma at all? In many of these cases, the traumatic events were accompanied by physical duress, such as sleep deprivation, head injuries, or alcohol abuse, each of which can disrupt memory (McNally, 2003b). In still other cases, the extreme stress associated with the event is likely to have disrupted the biological processes (the protein synthesis) needed for establishing the memory in the first place; as a result, no memory is ever established (Hasselmo, 1999; McGaugh, 2000; Payne, Nadel, Britton, & Jacobs, 2004).

Plainly, therefore, several factors are relevant to trauma memory. But the heated debate over these memories centers on a further claim: Some authors argue that highly painful memories will be *repressed*—that is, hidden from view by defense mechanisms designed to shield a person from psychological harm. In a related claim, some authors suggest that painful events will trigger the defense of *dissociation,* in which the person tries to create a sense of "psychological distance" between themselves and the horror. In either case, the proposal is that these memories are blocked by a specialized mechanism that's simply irrelevant to other, less painful, memories. According to this view, trauma memory is indeed a special subset within the broader domain of episodic memory.

Advocates for this special status point to several forms of evidence, including cases in which a memory seems to have been pushed out of consciousness and kept hidden for many years but is then "recovered" (brought back into consciousness) at some later point. This pattern is sometimes alleged in cases of child sexual abuse: The victim represses the memory (or dissociates) and so has no recollection of the abuse for years. Later in life, however, the victim recovers the memory, revealing at last the long-hidden crime.

Do these cases provide evidence for repression or dissociation, followed by a process of memory recovery? In answering, we need to start by acknowledging that incest and childhood sexual abuse are surely far more prevalent than many people suppose (Pipe, Lamb, Orbach, & Cederbork, 2007). It's also clear that some events—particularly emotionally significant events—can be held in memory for a very long time—years or even decades. We also know that it's possible for memories to be "lost" for years and then recovered (e.g., Geraerts et al., 2006; Geraerts et al., 2009; Geraerts et al., 2007). On all these grounds, then, it seems plausible that these memories of childhood abuse, hidden for decades, may be entirely accurate—and, indeed, provide evidence for horrid wrongdoing and criminal prosecution.

But there are some complications here. For one thing, the pattern just described—with memories lost from view and then recovered—may not involve repression or dissociation at all. As an alternative, these memories might just indicate a long-lasting retrieval failure that was eventually reversed, once the appropriate memory cue came

8.23 Recovered memories (A) Eileen Franklin (center) believed that she had repressed and then recovered memories of her father, George Franklin Sr., molesting and murdering her childhood friend 20 years earlier. (B) Based on his daughter's testimony about her recovered memories, George Franklin was found guilty and imprisoned until new evidence emerged showing that he could not have committed the crime.

along. This would take none of the importance away from these memories: They still would reveal serious crimes, worthy of substantial punishment. But, from this perspective, our explanation of these memories involves no mechanism beyond the memory processes we've already been discussing—and so, on this basis, the *content* of these memories would be distinctive in a tragic way; but the *processes* that create and maintain these memories would be the same as those operating in other cases.

In addition, and far more troubling, at least some of the "recovered memories" may in fact be false—created through mechanisms we discussed earlier in this chapter. After all, many of these recovered memories involve events that took place years ago, and we know that the risk of error is greater in remembering the distant past than it is in remembering recent events. Likewise, the evidence is clear that people can have detailed (false) recollection of entire episodes—including highly emotional episodes—that never happened at all. And we know that these false memories are even more likely if the person, from her current perspective, regards the "remembered" event as plausible (see, for example, Bernstein & Loftus, 2009; Chrobak & Zaragoza, 2008; Loftus, 2005; Ofshe, 1992; Pezdek, Blandon-Gitlin, & Gabbay, 2006; Principe, Kanaya, Ceci, & Singh, 2006). We also know that false memories, when they occur, can be recalled just as vividly, just as confidently, and with just as much distress as when recalling actual memories (Figure 8.23). All of these points remind us that we cannot take the veracity of the recovered memories for granted. Some recovered memories are very likely to be accurate, but some are likely to be false. And, sadly, we have no means of telling which memories are which—which provide a factually correct record of terrible misdeeds and which provide a vivid and compelling *fiction,* portraying events that never happened. (For discussion of this difficult issue, see Freyd, 1996, 1998; Geraerts et al., 2006; Geraerts et al., 2009; Geraerts et al., 2007; Ghetti et al., 2006; Giesbrecht et al., 2008; Kihlstrom & Schacter, 2000; Loftus & Guyer, 2002; McNally & Geraerts, 2009; Schooler, 2001.)

The debate over recovered memories is ongoing and has many implications—including implications for the legal system, because these memories are sometimes offered as evidence for criminal wrongdoing. But in the meantime, what about our initial question? Are traumatic memories in a special category, involving mechanisms irrelevant to other categories of episodic memory? In light of the continuing debate, the answer remains a matter of substantial disagreement.

Explicit and Implicit Memory

We've now considered several ways that explicit memory might be subdivided—into episodic memory and semantic, and then with each of those categories potentially divided further. But what about *implicit memory?* As we'll see, this memory provides an entirely different means through which we're influenced by past experience, and it is distinct from explicit memory in its functioning and in its biological basis.

DISTINGUISHING IMPLICIT FROM EXPLICIT MEMORY

Implicit memories are distinguishable from explicit memories in many ways. Perhaps the clearest evidence, however, comes from the study of the memory disruption caused by brain damage. In general, this disruption is referred to as *amnesia.* Earlier in the chapter, we mentioned one type of amnesia: *retrograde amnesia*—a loss of memories for events that took place before the brain injury that caused the amnesia. In other cases, though, brain damage produces **anterograde amnesia**—an apparent inability to form new memories (see Figure 8.9).

anterograde amnesia A memory deficit suffered after some kinds of brain damage, in which the patient seems unable to form new explicit memories; however, memories acquired before the injury are spared.

In general, anterograde amnesia is caused by damage to certain sites in the temporal cortex—specifically, in the hippocampus and nearby subcortical regions. In some cases, this damage is the result of illness—especially if the illness causes *encephalitis,* an inflammation in the brain tissue. In other cases, the damage is caused by stroke or physical trauma. One of the most common causes, though, is a type of malnutrition associated with chronic alcoholism; in this case, the amnesia is a central symptom of the illness called *Korsakoff's syndrome.*

One of the most carefully studied cases of anterograde amnesia, however, had an entirely different cause: A patient known as H.M. suffered from severe epilepsy. When all other treatments failed, the physicians tried (in the late 1950s) to treat H.M.'s disease with a neurosurgical procedure that deliberately removed most of his hippocampus, amygdala, and a considerable amount of nearby tissue (Figure 8.24). The procedure was, in a very narrow sense, a success: It did control his epilepsy. But the surgery also had a tragic and unanticipated side effect. Across the 55 years he lived after his surgery, H.M. seemed incapable of adding new information to his long-term memory. As his obituary put it: "each time he met a friend, each time he ate a meal, each time he walked in the woods, it was as if for the first time" (Carey, 2008). He remembered none of the episodes in his life after the surgery; he was entirely unable to recognize people he'd first met after the surgery—even if he saw them day after day (Milner, 1966, 1970; also see O'Kane, Kensinger, & Corkin, 2004; Skotko et al., 2004).

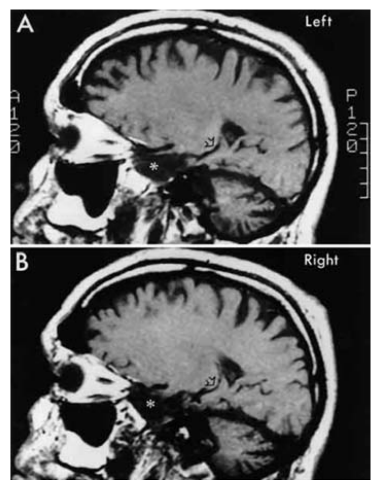

8.24 The brain of Henry Gustav Molaison Throughout his lifetime, Henry Molaison was identified in research papers only through his initials—H.M. His full name was released only after his death (in December 2008). Even after death, though, H.M. will contribute to our understanding of memory, because close analyses of his brain are underway. In these scans of his brain, we can see the space left by the surgical removal of tissue (marked with an asterisk). Note, though, that not all of H.M.'s hippocampus was destroyed; the remaining bit is marked with a small arrow.

Amnesia had devastating effects on H.M.'s life—including some effects that we might not think of as involving memory. For example, H.M. had an uncle he liked very much. When first told that his uncle had died, he was deeply distressed, but then he forgot all about this sad news. Some time later, he asked again when his uncle would come to visit, and he was told again of his uncle's death. His grief was as intense as before; indeed, each time he heard this sad news, he was hearing it for the first time—with all the shock and pain (Corkin, 1984; Hilts, 1995; Marslen-Wilson & Teuber, 1975; Milner, 1966; Milner, Corkin, & Teuber, 1968).

Crucially, and despite these remarkable problems, patients with anterograde amnesia—including H.M.—can acquire certain types of new memories which can be revealed with specialized testing. In some studies, for example, patients with anterograde amnesia have been given practice, day after day, in finding the correct path through a maze. Each time they're shown the maze, the patients insist they've never seen it before; this is simply a confirmation of their amnesia. Even so, they get faster and faster in solving the maze—and so apparently they do retain some information from each practice session.

Likewise, in another study, patients with Korsakoff's syndrome heard a series of brief melodies (Johnson, Kim, & Risse, 1985). A short time later, they listened to a new series and were told that some of the tunes in the second batch were repeats from the earlier presentation. As expected, these amnesic patients were completely unable to tell which tunes were the repeats and which were new; indeed, their memory responses were close to random. Remarkably, though, when asked which melodies they *preferred,* the patients uniformly preferred the familiar ones. The patients had no (explicit) memory for these tunes, but a memory did emerge with indirect testing—and emerged, in this case, as a preference.

In important ways, therefore, these patients can't remember their experiences. If we ask them directly about the past, they recall nothing. If we ask them which mazes they have solved before and which are novel, they can only guess. Thus it seems clear that these patients have no conscious recollection, no *explicit* memory, for the events in their lives. Still, we can find ways in which the patients' current skills and behaviors are shaped by their experiences—and so, apparently, the experiences have left some record, some residual imprint, in these patients. This lasting imprint, a demonstrable impact of the past, is what psychologists call *implicit memory*—an unnoticed "leftover" from life events that changes how someone now acts and thinks (Donaldson, Peterson, & Buckner, 2001; Fazio & Olson, 2003; Humphreys et al., 2003; Kinoshita, 2001; Yonelinas, 2002).

PROCEDURAL MEMORY

What exactly is implicit memory? In what circumstances does it influence us? And can implicit memory be demonstrated in people without amnesia, people whose brains are healthy and intact? To answer these questions, we need to distinguish different types of implicit memory, because each type influences us in its own way.

procedural knowledge Knowledge of how to do something, such as riding a bike; expressed in behaviors rather than in words.

declarative knowledge Knowledge of information that can be expressed in words.

Some cases of implicit memory involve **procedural knowledge** rather than **declarative knowledge.** Procedural knowledge is knowing *how*—knowing how to ride a bicycle, for example, or how to use chopsticks. Declarative knowledge, in contrast, is represented in explicit memory, not implicit, and it's knowing *that*: knowing that there are three outs in an inning, that automobiles run on gasoline, or that you woke up late this morning.

The earlier example we mentioned in our discussion of amnesia—that of patients learning how to get through a maze—involves procedural memory, and other examples are easy to find. In some procedures, patients have been shown a complex shape and asked to trace the outline of the shape with a stylus. What made this task difficult was that the patients couldn't see the shape directly; they could see it (and the stylus) only by looking into a mirror (Figure 8.25). This task is moderately difficult—but the patients got better with practice, all the while insisting that each try at the task was their very first time.

What about people who don't have amnesia—people with normal brains? In some studies, research participants are given four buttons and told to press button 1 if light 1 comes on, button 2 for light 2, and so on (Figure 8.26). The lights are then turned on in rapid succession, and participants do their best to keep up. As it turns out, the lights are turned on in a repetitive sequence—perhaps always 1-3-4-2-1-4-1-3-4-2-1-4-1-3-4-2-1-4. Participants seem to learn this sequence; and so, with a bit of practice, can respond more quickly if the lights follow this sequence than they can if the sequence is random. But when asked whether the sequence was random or not, participants are clueless. Thus, they seem to have procedural knowledge that allows them to respond more quickly to the patterned lights, but they don't have declarative knowledge—the same distinction we observe in patients with amnesia (Gazzaniga et al., 2009).

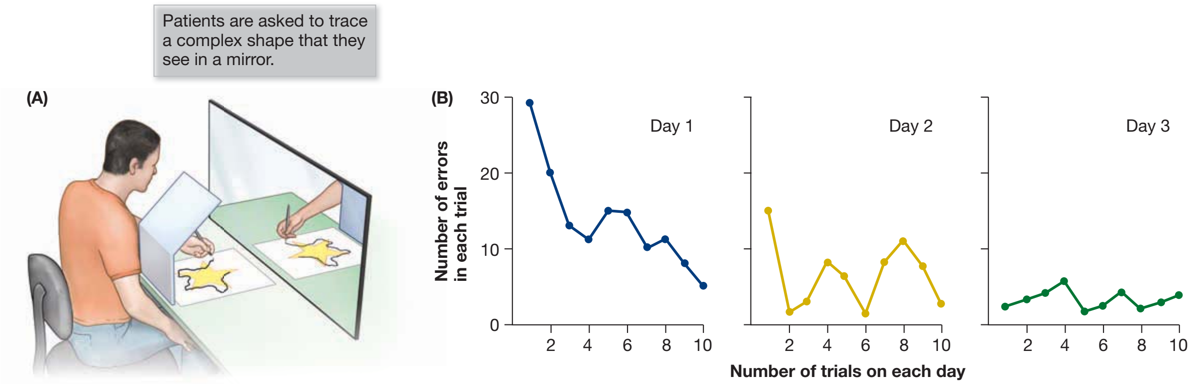

8.25 Mirror drawing (A) In mirror drawing, the research participant has to trace an outline of a figure while looking at his hand in a mirror. At first this task is very difficult, but after some practice the individual gets quite good at it. The same is true for amnesiacs. (B) The graphs show H.M.'s improvement on this task over a period of three days.

PRIMING EFFECTS

Procedural memories are typically concerned with behaviors—our actions and our skills. Other types of implicit memory, in contrast, influence our perceptions and our thoughts. Consider, for example, demonstrations of *priming.* Participants in one study were shown a number of words. Later, they were given a second task in which they simply had to identify words flashed briefly on a computer screen. Participants had no idea that many of the words in this second task were taken from the earlier list, but they still showed a pattern known as *repetition priming:* Words that had been on the original list were identified more readily than words that had not. This priming was observed even for words that the participants failed to recognize as familiar in a standard recognition task. Thus, the participants had no explicit memory for having seen these words, but they did have an implicit memory that showed up as priming. In other words, they were being influenced by a memory they didn't realize they had (Jacoby, 1983; Jacoby & Witherspoon, 1982).

Other procedures, with different tasks, show a similar pattern. In *fragment-completion tasks,* for example, participants are shown partial words (such as C__O__O__I__E) and asked to complete them to form actual words (CROCODILE). Success in this task is much more likely if the target word was encountered recently; this advantage is observed even when participants have no conscious recollection of the previous encounter (Graf & Mandler, 1984; Jacoby & Dallas, 1981; Tulving, Schacter, & Stark, 1982).

In another experiment, participants were asked to read sentences that were presented to them upside down. A *year* later, participants returned to the lab; there, they were shown a series of sentences and asked which ones they'd seen in their first visit to the lab and which ones were novel. Not surprisingly, after this long delay, participants couldn't tell which sentences they'd seen before. Still, when they were asked once again to read sentences presented upside down, they were faster with the sentences they'd seen before than they were with novel sentences—a case of priming that lasted across a full 12 months (Kolers & Roediger, 1984).

In each of these cases, it seems that an encounter with a stimulus leaves us better prepared for that stimulus the next time we meet it. This preparation can then influence

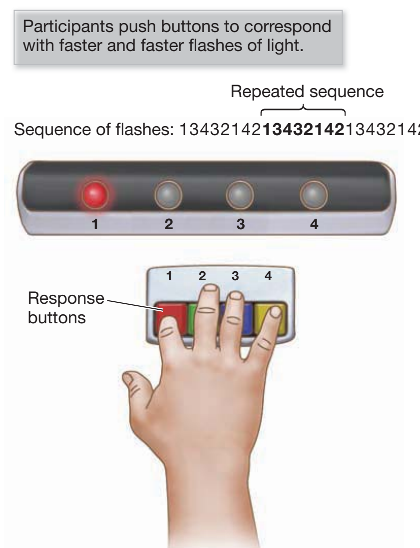

8.26 Procedural learning Participants have to press the appropriate key each time one of the numbered lights comes on. If there's a repeated sequence in the lights, participants seem to learn this and so get faster and faster in their task. However, participants have no declarative knowledge about the sequence—and may not even realize there was a repeated sequence.

us in many ways, quite independently of whether we can recall the earlier encounter with that stimulus. To illustrate how far this pattern can reach, consider the so-called illusion of truth. In the relevant studies, participants hear a series of statements like "The average person in Switzerland eats about 25 pounds of cheese each year," or "Henry Ford forgot to put a reverse gear in his first automobile."* Participants' task is to say how *interesting* each of these statements is. Later on, the same participants are presented with some more sentences but now have to rate the credibility of each one on a scale from "certainly true" to "certainly false." Needless to say, some of the sentences in this "truth test" are repeats from the earlier presentation; the question for us is how the judgments of sentence credibility are influenced by the earlier exposure.

The result of these studies is a propagandist's (or advertiser's) dream: Sentences heard before are more likely to be accepted as true, so that in essence familiarity increases credibility (Begg, Anas, & Farinacci, 1992). To make matters worse, the effect emerges even when participants are warned not to believe the sentences in the first list. That is, sentences plainly identified as false when they're first heard still create the illusion of truth, so that these sentences are subsequently judged to be more credible than sentences never heard before.

How could this be? Bear in mind that the participants in these procedures are shown a lot of sentences and that there's a delay between the first task (judging how interesting the sentences are) and the second (judging credibility). These steps make it difficult for participants to keep track of the sentences they hear; in other words, these steps work against explicit memory. As a result, participants have a hard time recalling which of the sentences in the truth test they encountered on the first list; so it doesn't help them to know that the sentences on that first list were all false. Thus, with no conscious memory of the earlier encounter, participants have no way to protect themselves from the illusion.

OTHER FORMS OF IMPLICIT MEMORY

We've now mentioned two forms of implicit memory—priming effects and procedural memory—but there are other forms as well. One plausible addition to this list is *perceptual learning*—the learning that you need to do whenever you "recalibrate" your perceptual systems. As an example, think of what happens when someone gets new eyeglasses, perhaps with a stronger prescription than they've had before. Across the next few days, he needs to "adjust" to the glasses—changing (among other things) how he interprets the degree of tension in his eye muscles as a cue to distance. This is surely a form of learning—and so places new information, or perhaps new skills, in memory. But it's learning that happens completely outside of awareness—and so involves implicit memory, not explicit.

A different example involves cases we considered in Chapter 7—including the learning called *classical conditioning.* This learning, too, creates new knowledge—knowledge about what follows what in the world—but can be done without conscious awareness. Indeed, classical conditioning can take place even if an organism is fully anesthetized during the learning (e.g., Cahoon, 2008).

These and other examples demonstrate the enormous breadth of implicit memory. We rely on *explicit* memory in many circumstances, and we're guided to an enormous extent by our conscious recollection of the past. But the reach of implicit memory may be even larger—so that in many situations, we're shaped in ways we do not notice by past experiences that we cannot recall.

---

*The first statement, by the way, is false; the average is closer to 18 pounds. The second statement is true.

SOME FINAL THOUGHTS: DIFFERENT TYPES, BUT COMMON PRINCIPLES

In the last few sections, we've seen that the term *memory* can be applied to a diverse set of phenomena. The term applies to the (episodic) memories that you draw on when you get together with friends and talk about what you did last summer. It also applies to your (semantic) memories that grass is green, "up" is the opposite of "down," and Beijing is the capital of China. The term also applies to your (procedural) memory of how to ride a bicycle, and your (classical conditioning) memory that makes you gag whenever you smell gin—a response rooted in the fact that you got horribly ill the day after you drank all those martinis.

Without question, these different types of memories obey different principles: Explicit memory, as we've seen, depends on active engagement with the materials at the time of learning. In contrast, many forms of implicit learning seem relatively passive and can perhaps be created by exposure alone. There has been debate over whether some forms of memory (episodic memory in particular) might be found only in a few species—perhaps those that are self-aware and capable of conscious reflection on the past. Other forms of memory (e.g., the memory that supports classical conditioning) can be found in a wide range of species—snails or worms, for example, as well as chimpanzees or humans.

We might worry that this broad usage of the term *memory* stretches things too far and might even be misleading in some cases. After all, if we say "the rat remembers that the tone was followed by food," this might imply that the rat is aware of this fact about the world, and can reflect on it and draw inferences from it. But this would be a mistake—misrepresenting the (unconscious, largely automatic) qualities of implicit memory in general and certainly misrepresenting rat capacities in particular.

How can we find a balance—so that we bring together the various memory achievements discussed in this chapter and the previous one, but so that we don't lose distinctions among the various types of memory? The answer may lie in careful emphasis on the distinction between implicit memory and explicit, keeping separate the types of memory that are conscious and allow reflection and the types of memory—crucial as they are for many purposes—that do not. This distinction, in turn, points us toward some new questions. For example, we know that humans have explicit memories; we see this in the simple fact that we can, if asked, report on our memories and describe the past as we recall it. Which other species also have explicit memories—and thus are aware of themselves and their own past? This question opens a window through which we might explore the intriguing issue of conscious experience and awareness in other creatures.

This distinction also highlights the key role that explicit memory plays for human experience—and, indeed, this is why we've devoted most of this chapter to this type of remembering. Explicit memories are—by definition—memories that we're aware of. We can therefore reflect on past experiences and discuss them with others—sometimes to instruct them, sometimes to foster social bonds. We can also report on our memories when someone else needs to learn what happened in a prior episode—whether it's a journalist trying to understand what happened in yesterday's storm or a police officer investigating how things unfolded in last night's robbery. And, finally, we can draw conclusions from these memories—and we often do so, because many of the decisions we reach, or judgments we make, are based on considerations drawn from memory. Plainly, then, explicit memories play important roles in many human functions, and we're obviously able to *think about* our explicit memories in ways that really matter for us. But what does this "thinking" involve? That will be the focus of our next chapter.

SUMMARY CHAPTER 8

ACQUISITION, STORAGE, RETRIEVAL

- Any act of remembering begins with *acquisition*, the process of gathering information and placing it into memory. The next aspect of memory is *storage*, the holding of information in some enduring form in the mind for later use. The final phase is *retrieval*, the point at which we draw information from storage and use it in some fashion.

ACQUISITION

- Memory *acquisition* includes cases of *intentional learning* and *incidental learning*. In either case, the person must pay attention to the material to be remembered, and it is the product of this intellectual engagement that is stored in memory.
- According to the stage theory of memory, information is held in *working memory* while one is thinking about it, but it's lodged in *long-term memory* for storage for longer intervals. This theory is supported by studies of free recall. In these studies, *primacy effects* reflect the fact that early items in a presentation receive more rehearsal and are more likely to be transferred to long-term storage. *Recency effects* reflect the fact that just-heard items can be retrieved directly from working memory.
- *Chunking* is the process through which items are recoded into a smaller number of larger units. The active nature of memory is also evident in the fact that mere *maintenance rehearsal* does little to promote long-term storage.
- According to many studies, how well someone remembers will depend on the depth at which he or she processed the incoming information; *shallow processing* refers to encoding that emphasizes the superficial characteristics of a stimulus, and *deep processing* refers to encoding that emphasizes the meaning of the material. Consistent with this perspective, we remember best the material that we've understood, thanks to the memory connections linking one memory to the next. At the time of recall, these connections serve as retrieval paths.
- *Mnemonics* help a person form memory connections, and these connections can dramatically improve memory. Many mnemonics utilize imagery, and imagery is most helpful if the visualized items are imagined in some interaction—linking the items to each other, as one would expect if imagery is a means of promoting memory connections.

STORAGE

- More research is needed to explore how the *memory trace* is actually represented in the brain. However, evidence suggests that different elements of a single memory (what things looked like, how one felt) may be stored in different brain sites.
- The establishment of a long-term memory depends on a *memory consolidation* process, during which new connections are formed among neurons. The need for consolidation is reflected in cases in which this process has been disrupted, resulting in *retrograde amnesia*.

RETRIEVAL

- The *retrieval* of memories is often easy, but it sometimes fails. The failure can be complete or can be partial, as in the *tip-of-the-tongue effect*. The retrieval of memories is often promoted by our having an appropriate *retrieval cue*. Whether a cue is useful depends on whether the cue re-creates the context in which the original learning occurred. This *context reinstatement* allows the person to use the connections they formed earlier as *retrieval paths*.
- What's stored in memory reflects how the person thought about or reacted to the object or event being remembered. This *encoding specificity* is reflected in the fact that remembering is more likely if one thinks about the target information during retrieval in the same way that one did during encoding.

MEMORY GAPS, MEMORY ERRORS

- Many cases of forgetting can be understood as the result of inadequate encoding. This is reflected in the fact that fMRI data, collected during encoding, show different patterns for later-remembered material and later-forgotten material.
- Forgetting generally increases as the *retention interval* gets longer, but the causes of forgetting are still being debated. One theory holds that traces gradually decay. Another view argues that the cause of forgetting is interference produced by other memories. In some cases, this is because the other memories promote retrieval failure—an inability to find information that's nonetheless still in storage. Retrieval fail-

ure is evident whenever some new cue allows us to recall previously forgotten materials.

- Interference can also result from the mixing together of memories. These *intrusion errors* are evident in the *misinformation effect,* in which specific episodes are blurred together. In other cases, intrusion errors are the result of schematic knowledge intruding into someone's memory of a particular event. This reflects a broader pattern of evidence indicating that events are usually understood (and remembered) with reference to knowledge structures called *schemas.*
- Intrusion errors can also be produced by semantic associations with the material being recalled. This is the source of the errors often observed in the *DRM paradigm.*
- Another category of memory errors involves cases in which someone correctly realizes that an idea (or face or stimulus) is *familiar,* but makes an error about why the idea is familiar. This pattern reflects the fact that separate memory systems are the bases for familiarity and *recollection.*
- Psychologists have searched unsuccessfully for ways of distinguishing correct memories from mistaken ones. The *confidence* expressed by the person remembering turns out to be of little value for this discrimination. Hypnosis also does nothing to improve memory and can actually increase the risk of memory error.

VARIETIES OF MEMORY

- Researchers find it useful to distinguish several types of memory. *Episodic memories* concern specific episodes; *semantic memories* concern broader knowledge, not tied to a particular episode. *Explicit memories* are consciously recalled; *implicit memories* are revealed when there is an effect of some past experience without the person being aware that she's remembering at all—or even that there was a relevant past experience.
- Some theorists subdivide episodic memory, distinguishing autobiographical memories from memories for other episodes, and placing *flashbulb memories* or *traumatic* memories into their own category. However, current evidence suggests that flashbulb memories are governed by the same principles as other memories, and the same is true for traumatic memories—although debate continues over the possible role of "repression" or "dissociation" in memory for traumatic events.
- Certain injuries to the brain produce *anterograde amnesia,* in which the patient's ability to fix material in long-term memory is reduced. However, someone with amnesia may still have intact implicit memories. Implicit memories, in turn, can be divided into several types: procedural memories, involving changes in behavior, priming, changing our perceptions and thoughts, and perceptual learning.

ONLINE STUDY TOOLS

Go to StudySpace, **wwnorton.com/studyspace**, to access additional review and enrichment materials, including the following resources for each chapter:

Organize
- Study Plan
- Chapter Outline
- Quiz+ Assessment

Learn
- Ebook
- Chapter Review
- Vocabulary Flashcards
- Drag-and-Drop Labeling Exercises
- Audio Podcast Chapter Overview

Connect
- Critical Thinking Activity
- Studying the Mind Video Podcasts
- Video Exercises
- Animations
- **ZAPS** Psychology Labs