

8

Law and Eyewitness Accuracy

Do you swear to tell the truth, the whole truth, and nothing but the truth?

Such a simple question that people so often “swear” to, yet how easy is it to live up to telling the “whole truth, and nothing but the truth”? Certainly, for eyewitnesses who are reporting on their memory of a crime, telling the whole truth (and nothing but the truth) requires accurate metacognitive monitoring and control—to sift out incorrect memories or fabricated ones to get to just the valid ones. From what you’ve already learned from this textbook, you know that it is likely impossible for anyone to meet this standard—telling nothing but the truth—in any but the simplest of scenarios. Moreover, the perception that people can achieve this standard in the courtroom is just one of many examples of how metacognition is relevant to law and eyewitness testimony. People’s beliefs about their own and other people’s memories—their confidence that others are remembering correctly, their beliefs with respect to suppressing or disregarding evidence and their ability to do so, and their evaluations about the veracity of other people’s reports—are crucial metacognitive phenomena of particular relevance to courtroom proceedings and the judicial system.

In Chapter 8, we review some findings on the relation of people’s confidence to the correctness of their memories, on people’s assessment of the truth of witnesses’ memories based on their expressed confidence, and on people’s

abilities to detect lies. Finally, we will explore the pernicious influence of a phenomenon called hindsight bias or the “I knew it all along” effect. Experiments show that even though people have a good ability, for example, to assess whether testimony should be admissible, this “knew it all along” effect thwarts their ability to control their decisions about guilt or innocence accordingly. As we shall see, this particular kind of mental time travel—back to a naive state—is extraordinarily difficult for people to do. Before we get to that, however, let us begin our exploration of the role of metacognition and the law with people’s feelings and expressions of confidence.

Confidence and False Memories

If people are given misinformation about an event that they experienced or witnessed, they may later believe that the misinformation is true and actually happened. For example, imagine watching a film in which a car goes through a stop sign, and then you receive the suggestion that it actually had been a yield sign. Later, when you are asked whether it was a stop sign or a yield sign that you saw, you may well say it was a yield sign (Loftus, Miller, & Burns, 1978). Even *how* you are asked about what happened during an event can influence your memory of it. After witnessing a car accident on video, people can be asked how fast the cars were going when they “smashed” into each other, or instead, how fast the cars were going when they “hit” each other. When the question is framed in terms of “smashing,” people are likely to remember the speed as being faster than if the question is framed as “hit.” They also may report that they saw broken glass at the scene of the accident, even though no windows were broken (see Loftus & Hoffman, 1989, for these and other examples). Moreover, when an inaccurate event (e.g., you saw a yield sign) is suggested repeatedly (versus just a single time), people are even more likely to claim that the event actually happened, and their confidence in this false memory is greatly enhanced (Zaragoza & Mitchell, 1996). In all of these situations, people come to inaccurately judge that an untrue memory about an event is actually true. People’s metacognitive monitoring is failing, and unknowingly, they do not tell “nothing but the truth.”

Perhaps the most famous research involving completely implanted misinformation is a study by Loftus and Pickrell (1995) in which the experimenters implanted a person with the “memory” of having been lost in a shopping mall. We describe this experiment in detail below, but an important point here is that the implanted memory had emotional content. By some views, emotional memories should be difficult to mistake. Indeed,

some therapists and lay people believe that emotional memories simply cannot, by their very nature, be false, and that reported emotional memories of abuse, such as those that frequently surface during therapy, are necessarily true. In contrast to this view, the possibility that some disturbing memories might be attributable to suggestion, to imagination, or to having been inadvertently implanted by therapists trying to help the victim, need not indicate that the victims are not reporting what they remember in good faith (or, indeed, that the therapists were not trying to help). It is conceivable that a person might genuinely believe that a false memory is true. Whether or not this happens, however, is an empirical question.

High confidence is frequently taken to mean that a memory is true. But, as we saw in the chapter on confidence, people, at least in laboratory situations, are frequently overconfident. Are people in the real world also overconfident? When people swear they are telling the truth, can they be mistaken? The answer to both questions is “yes.” And, in the next sections, we’ll consider some answers to more specific questions that arise in the literature: First, could a person show high confidence about remembering an event that never happened? Second, what underlies people’s confidence, and does this confidence normally provide a reliable index of which memories are true and false? Third, can confidence be manipulated? If so, how? And, fourth, are there any special circumstances—such as when a memory is of a traumatic event—in which a person’s confidence in his or her memory is infallible?

Can We Have High Confidence, Even for an Event That Never Happened?

The study introduced above by Loftus and Pickrell (1995) addressed our first question. They conducted a case study in which they tried to implant a nonexistent memory into the mind of a 14-year-old boy named Chris. He was given descriptions of three true events that had happened in his childhood along with one false event that had never happened. The false memory was introduced in a paragraph reminding Chris that when he was 5 he had been lost in the University City shopping mall in Spokane, Washington, where the family often shopped. He was crying loudly when rescued by an elderly man, and reunited with his family. Chris wrote about the four events every day for 5 days, providing all of the descriptions he could remember of each, but being told that if nothing more came to mind he could write “I don’t remember.” On the final day, it was evident that the false event that had been initially suggested to Chris had become a vivid memory. He remembered that the man who rescued him was “cool,” that he was scared he would never see his family again, and that his mother scolded him.

He remembered that the man wore a blue flannel shirt, that he was “kind of bald on top,” and that he had glasses.

But surely, even though he could imagine such details, he did not believe that this completely confabulated event had actually happened? When asked to give his confidence about these four memories on a scale from 1 (low confidence) to 11 (high confidence), the ratings for the true memories were 1, 5, and 10, for a mean of 5.3. He gave an 8 to the false memory, and provided rich details about his thoughts at the time it had occurred. Clearly, his confidence was misplaced, because the event never happened (for other provocative examples, see Loftus, Coan, & Pickrell, 1979/1996).

An equally famous experimental study—the so-called Sam Stone study—was conducted on young children by Leichtman and Ceci (1995), who investigated how 3- to 4-year-olds and 5- to 6-year-olds remembered events that had never happened and how such “memories” might play out in a courtroom trial. Their study, like many others that asked the question of whether we should invariably believe children’s testimony, was conducted hard on the heels of a number of real-life incidents in which young children had falsely accused adults of wrongdoing. Before we describe the experimental evidence, we’ll highlight two—of many—cases in which children’s false accusations have had a devastating effect on everyone involved.

In the Kelly Michaels case, a minor incident (one of the children, after having had his temperature taken at a pediatrician’s office, said that his teacher had done that as well) led to repeated and very leading interrogations of all of the children in the Wee Care Nursery School. Of course, the parents were alarmed and outraged. The “memories” of the children suddenly started revealing all manner of abuse by Kelly Michaels. The escalation of the hysteria over what was apparently completely confabulated memory on the part of the children eventuated in Michaels being charged with 131 counts of sexual abuse against 20 children. All of these charges were made solely on the children’s alleged memory of abuse. Michaels evidently committed these 131 heinous acts without provoking the notice of any other teacher (even though the memories of the children included things like Michaels dancing naked on top of a piano). No physical evidence was available. Michaels was convicted of 115 counts of sexual abuse of preschoolers, and was sentenced to 47 years in prison. Eventually, with the help of two investigative reporters (Dorothy Rabinowitz and Debbie Nathan), an attorney who took up her case (Morton Stavis), and a brief by 45 cognitive and clinical psychologists, the verdict was overturned and Michaels was released. But not before she had served 5 years in prison, 18 months of which were in solitary confinement (for her own safety). In a similar and equally outrageous case (the “Christchurch Crèche case” in New Zealand), Peter Ellis

served his full sentence for alleged satanic abuse for which there was absolutely no evidence except the children's often wildly impossible testimony. Ellis refused to be released early on parole because the condition for doing so was that he admit his guilt.

Leichtman and Ceci's (1995) experimental study was an attempt to investigate such extreme examples of memory contagion in a controlled manner. Unlike most real-world events where it isn't possible to know (for sure) what really happened, the original events in the experiment were documented and allowed careful analyses of the boundary conditions of children's memory and confidence. The experiment takes its name from its central character, Sam Stone. Before he visited the nursery school, the children were told, repeatedly, that Sam was clumsy and tended to break things. When he actually visited, Sam stayed for only 2 minutes, and interacted amicably with the children. He broke nothing. The next day, however, the children were shown a soiled teddy bear and a torn book, and were asked whether, perhaps, Sam had done it. About 25% of the children said that perhaps he had done it, but none claimed to have seen him do it. Over the next 10 weeks, however, they were repeatedly asked things like, "I wonder whether Sam Stone was angry when he tore the book?" "I wonder whether he got the teddy bear dirty on purpose or by accident?" This questioning was considerably more benign, by the way, than the interrogations endured by the children in the Kelly Michaels case. Nevertheless, by the end of the 10-week period, when the children were asked to describe Sam Stone's visit to an outsider who hadn't been there at the time, 72% of the 3- to 4-year-olds said that Sam Stone had ruined one of the items, and fully 45% of them said they had actually seen him do it—indicating a high-confidence memory. The older children were a little better than the younger ones, with only 11% claiming to have actually seen him do it (but 11% is enough to convict!). The children who claimed to have seen him do it embellished these false memories with perceptual details that, according to Johnson and Raye's (1981) framework of source monitoring (see Chapter 7), would be convincing to both the children and to outside onlookers. So, if you confidently remember seeing the event (versus just believing you saw it), and you can describe some detail as to how it happened, it seems natural to believe the event could have happened.

False memories might be thought to be a phenomenon specific to children. Children have difficulty with reality monitoring, as is illustrated in Box 10.1 in the developmental chapter. Indeed, children have, until recently, been excluded from testifying in the U.S. legal system, in part, because of the notorious false accusations and memories of the children of Salem, Massachusetts, in 1692. Twenty-nine people were convicted and 19 were hanged because children remembered seeing them doing acts of witchcraft.

Although false memories may, in fact, be more common in children, false memories are certainly not the exclusive domain of children. Susan Clancy (2005) has interviewed more than 50 people who believe, with great conviction, that they have been abducted by aliens. This is a claim that most of us would agree is a false memory, despite the interviewees' high confidence. It would seem, then, that it is entirely possible for someone to have a high-confidence memory about something that never happened.

So, does high confidence in a memory count for nothing? We will explore the mechanisms underlying confidence in the next section.

What Underlies People's Confidence?

Given that people's confidence is based on the amount, speed, or clarity of what comes to mind, as well as the surrounding details indicating that the event was real and happened to them, one might expect that confidence would be a good indicator of correctness. Indeed, there are a host of confidence studies that begin with people being presented with a list of words. On a later test, they are shown each word again, along with new words that they did not study, and they indicate whether each word was presented earlier. They also make a confidence judgment about whether each recognition decision is correct. Most relevant here, the relative accuracy of people's confidence judgments is high. Higher confidence ratings almost inevitably mean that the item had been previously presented. Low ratings correlate very well with the items being new.

And, in the real world, too, when a person says, "I saw the person clearly and have high confidence that this is the culprit" as compared to someone who says, "Well, it was kind of fast and pretty much a blur and I think that maybe it was this person but I'm not sure," most people readily believe the person with high confidence and not the one with low confidence, and with good reason. Indeed, when people have been given factual questions to answer, followed by their confidence judgments about the correctness of their answers, their confidence judgments are highly accurate. Data from Butterfield and Metcalfe (2006), for example, showed that when people expressed confidence in the highest third of the confidence range, they were correct 70% of the time. In contrast, they were only correct 16% of the time when their confidence was low. And when people find out they have made a high-confidence error, they are surprised. Butterfield and Mangels (2003) have used event-related potential (ERP) techniques—in which electrodes measuring brain activity are attached to the scalp—to show that a particular "surprise"-related brain wave called the p300 occurs in these cases. Thus, under conditions in which confidence is not manipulated, an individual's

confidence is likely to be a good marker of whether he or she is right or wrong. The problem is, as we will see shortly, that confidence can be readily manipulated, so that it does not accurately discriminate between true and false memories.

Even so, several studies have shown that people's confidence can be used as a barometer for determining whether a memory is true or false, at least to a limited extent. Read (1996; also see Roediger & McDermott, 1995), for example, conducted a study in which people saw the words "slumber, tired, rest, night, dark, comfort, sounds, eat, bed, snore, dream, and awake." When asked to recall the items from the list, 65.9% of the participants falsely recalled that "sleep" had been on the list. But was it? Look back. It's not there. Moreover, the same percentage recalled words that actually had been presented, so percentage of recall did not indicate whether the memory was a true memory or a false one (for a demonstration of this powerful effect, see the Minds-On Activity at the end of this chapter). People's confidence, however, does. In particular, in the study by Read (1996), retrospective confidence judgments were higher for words that actually had been presented (4.55, where 5 = extremely confident that the word was presented to 1 = no confidence) than for the seductive lures (e.g., "sleep") that had not been presented (3.0). Furthermore, when participants were asked whether they actually *remembered* hearing the presented words (e.g., slumber, tired, and rest), as compared to sleep, the proportion of "remember" responses was greater in the former case than in the latter (.73 and .46, respectively). But is the glass half empty or half full? That is, people believed in a false memory 46% of the time! So, although the confidence judgments do show some accuracy, people still do make metacognitive errors.

Some data indicate that a different metacognitive measure—people's source judgments—can sometimes be used to determine whether a memory is real, due to actually witnessing an event, or merely the result of a suggestion (Hicks & Marsh, 1999). Although people's fast familiarity recognition judgments may lead them to wrong conclusions about past experiences, being asked for more refined source judgments allows the elimination of some of the errors. For example, Lindsay and Johnson (1989) showed that sometimes people initially claim to have seen something that had only been suggested to them. If they were pressed further, however, to determine the source of the memory, people were better able to identify those items that were actually seen as compared to those only suggested. Use of this careful kind of source monitoring of memories could be invaluable in extracting accurate testimony from witnesses (see Box 8.1, The Cognitive Interview). Later, in Chapter 11, we will discuss how older adults who were trained to adopt more stringent standards to evaluate their memories were able to

overcome one kind of memory illusion (Multhaup, 1995). It would appear that if metamemory processes are adequately engaged, people may be in a better position to sort out real memories from confabulations or implanted memories (Lane, Roussel, Villa, & Morita, 2007).

BOX 8.1 **The Cognitive Interview**

In the United States, more than 200 people have been convicted for committing a crime they did not commit and later have been exonerated by DNA evidence—after serving an average 12 years in prison. Cassel (2000) pointed out the great concern over such wrongful convictions. These, and other cases, led the U.S. Department of Justice to issue national Guidelines for Eyewitness Evidence, and to a nationwide effort to improve the methods of eliciting information from witnesses to avoid the production of false memories.

The Virginia case of Tommy David Strickler is an example of why such guidelines are needed. Strickler was convicted of capital murder, abduction, and robbery in the death of Leanne Whitlock. In 1991, he was sentenced to death. The case was appealed on the grounds that police evidence that might have discredited the testimony of the central eyewitness—who, at trial, convincingly gave a compelling and confident narrative filled with vivid detail—had not been made available to the defense. The prosecutor had said in his closing argument, “We are lucky enough to have an eyewitness who saw [what] happened out there in that parking lot. [In a] lot of cases you don’t. A lot of cases you can just theorize what happened in the actual abduction. But Mrs. Stoltzfus was there, she saw [what] happened.” The eyewitness, Anne Stoltzfus, publicly credited her vivid trial account to her “exceptionally good memory” and to her “very close contact with [the petitioner that] . . . made an emotional impression” on her. She said she had “absolutely no doubt” about seeing Strickler abduct Whitlock.

There was much, however, that the jury had not heard. For example, during her first interview with police two weeks after the crime, Stoltzfus had been unable to identify the victim, Leanne Whitlock. She identified Whitlock two weeks after her first meeting with Detective Daniel Claytor, and only after spending time with Whitlock’s boyfriend “looking at current photos” of Whitlock. Early in the investigation, she also could not identify the alleged perpetrators. In a letter written to the detective three days after their first interview, she admitted that she had not remembered even being at the mall, but that her daughter had helped “jog” her “vague memory.” In another early note to the detective, she only vaguely described the victim’s car and she failed even to mention the license plate number that she claimed she dictated to her daughter, a passenger

in her car at the scene of the abduction. At trial, however, she recalled both the car and the license number in detail. Finally, in a letter to Detective Claytor, she thanked him for his "patience with her sometimes muddled memories" and noted that "it would have been nice if I had remembered all this at the time and simply gone to the police with the information. But I totally wrote this off as a trivial episode of college kids carrying on." Further, as Cassel (2000) noted, in a letter to the Harrisonburg *Daily News-Record*, Stoltzfus said (all unbeknownst to the jury), "It never occurred to me that I was witnessing an abduction. In fact, if it hadn't been for the intelligent, persistent, professional work of Detective Daniel Claytor, I still wouldn't realize it. What sounded like a coherent story at the trial was the result of an incredible effort by the police to fit a zillion little puzzle pieces into one big picture." Strickler eventually lost his appeal on the basis that there was sufficient evidence, other than Stoltzfus's testimony, to justify his conviction and the imposition of the death penalty.

Police departments don't want to present questionable testimony that could discredit them or have verdicts overturned (or have questionable verdicts upheld). Accordingly, even many years before the Department of Justice guidelines were issued, police departments throughout the United States had been collaborating with psychologists to improve the way they interviewed witnesses to decrease the chances that they would implant false information. Ideally, interrogators want to extract more information from witnesses, without increasing errors. In response to this need, Ronald Fisher, a professor at Florida International University, along with Ed Geiselman, a professor at UCLA, developed the *cognitive interview*, a procedure now used by many police departments in the United States (Fisher & Geiselman, 1992).

The cognitive interview guides a witness through four memory-enhancing procedures: (a) thinking about the context, including the physical surroundings and their personal emotional reactions at the time of the event; (b) reporting everything that comes to mind, regardless of how fragmentary or inconsequential it seems; (c) remembering the events in several different sequences—from beginning to end, but also in reverse order, and from highly memorable points as well as from points of low salience; and (d) recalling the events from different perspectives, such as from a different person's point of view, or from above the scene. In its modified form, the cognitive interview also concentrates on building rapport with the witness, and getting the witness to actively participate in the interview rather than just responding to questions. It encourages the witness to tell the story in his or her own words without interruption before asking the witness to return to the memories from different perspectives and to examine the memory images in as much detail as possible. What the cognitive interview does not do is make suggestions of any sort about the content of what is remembered. Most important, studies have shown that this method increases the amount of correct information the witness recalls (often by 30 to 40%), and it does not increase the amount of incorrect information.

Can Confidence Be Manipulated?

Recall the Howard Haupt case, from Chapter 6. Haupt was mistakenly accused of kidnapping. He was not one of the original suspects in the case, but he had been staying at the hotel where the victim was kidnapped, and he had appeared in a photo lineup shown to the witness who later identified him. It was after the photo lineup, and after the witness saw Haupt in person, that he identified Haupt in another photo lineup as the kidnapper, with high confidence. That is, after multiple exposures to Haupt, the witness became familiar with him and began to believe that in fact he had committed the crime. Presumably, the witness had mistaken the origin of his familiarity with Haupt: Instead of saying “Haupt seems familiar because I recently saw him in a photo lineup,” he instead attributed his familiarity to Haupt’s having been the kidnapper. In this example, increased confidence appears to have resulted from exposure.

There is now good evidence that confidence can be influenced by the mere presentation or repetition of information, whether that information is diagnostic or not, and even whether the information is true or not. In a seminal study on spurious confidence, Oskamp (1965) had three groups of participants—licensed clinical psychologists, psychology graduate students in training, and undergraduates—who were briefed about the case of Joseph Kidd (a pseudonym). The experimental participants received four installments of material about Kidd in an effort to simulate the increased knowledge that therapists might acquire over sessions, as they got to know a client better. The first stage contained only demographic information as follows: “Joseph Kidd is a 29-year-old man. He is white, unmarried, and a veteran of World War II. He is a college graduate, and works as a business assistant in a floral decorating studio.” The second stage consisted of 1½, single-spaced, typed pages about Kidd’s childhood; the third stage was 2 pages about his high school and college years. The fourth stage, 1⅓ pages, went through his army experience and his life up to age 29.

After reading each stage of Joseph’s case, the participants answered 25 questions, such as

During college, when Kidd was in a familiar and congenial social situation, he often

- (a) Tried to direct the group and impose his wishes on it.
- (b) Stayed aloof and withdrawn from the group.
- (c) Was quite unconcerned about how people reacted to him.
- (d) Took an active part in the group but in a quiet and modest way.
- (e) Acted the clown and showed off.

The correct answer is (e).

The other questions were similar to this one in that they tapped into Kidd's customary behavior patterns, attitudes, interests, and typical reactions. They were based on either factual data or well-documented conclusions from the actual case. The experimental participants were expected to follow the usual procedure in clinical judgment in making their assessments by forming a personality picture of Kidd on which to base their conclusions—in much the way a juror is expected to amass knowledge and draw inferences. None of the questions, however, were rote memory questions from the materials presented. The judge-participants made a confidence judgment on each of their answers, with a value of 20% indicating chance performance.

The data were similar across groups except that the more experienced clinicians were, perhaps surprisingly, somewhat less confident by stage four than were the other groups. The data (collapsed across the various groups of participants) are shown in Figure 8.1. Test performance did not improve as the evaluators were given more and more information. Mostly, the questions were answered incorrectly—none of the performance data are above chance levels (which was 20% given that each multiple-choice question had five alternatives). What did increase, and dramatically so, was people's confidence in their answers. Whereas the participants were only slightly overconfident after reading the demographic information alone, their ratings were wildly overconfident by the end of stage four. It seems that even information that is unhelpful in increasing accuracy (in this case, additional information provided about Kidd across the four stages) has a large and unwarranted effect on people's confidence.

Along similar lines, Shaw and McClure (1996) showed that merely asking questions several times increased people's confidence in their answers, even though the answers themselves did not improve. In their experiment, they staged an interruption in a classroom, and then questioned the students about the event 5 times over the next 5 weeks, with some of the questions being repeated and some not. Although the accuracy of their answers did not improve over the 5-week interval, their confidence in those answers did. Studies in which people's confidence increases with the sheer volume of information, or from mere repetition, rather than with the quality of the information have implications for courtroom confidence. A standard practice in preparing witnesses for trial is to rehearse the witnesses in the story that they will tell. The grooming that goes on to prepare witnesses in high-stakes trials approaches that done for presidential candidates, and the result, when successful, is the same. The witnesses become highly confident, but sometimes for the wrong reasons. The pernicious aspect of this drilling

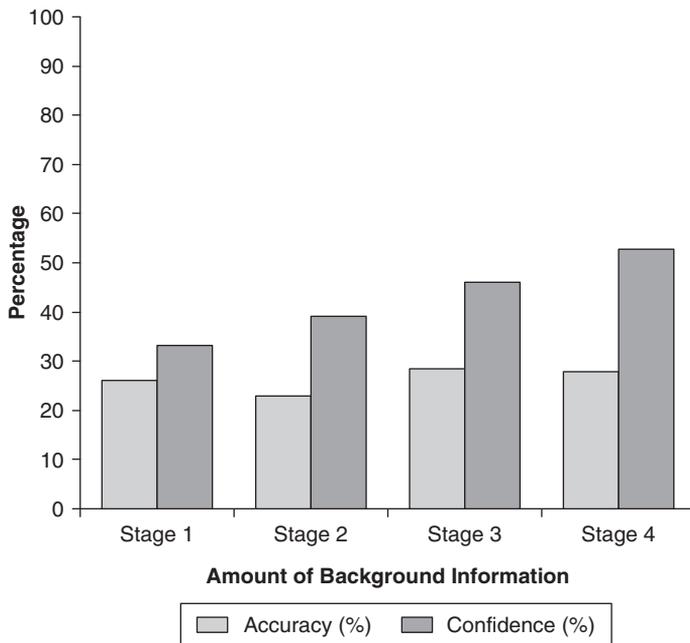


Figure 8.1 Percentage correct (accuracy) on the tests and the percentage confidence in answers on the test across four stages in the spurious confidence study.

SOURCE: Oskamp (1965).

procedure is that the witnesses become confident, and exude that assured aura, whether what they are remembering is true or not.

Is People's Confidence in Their Memories Ever Infallible?

Surely the inaccurate confidence judgments given in the above cases are mere quirks. Perhaps we do need to be wary of what young children say with high confidence—their frontal lobes are not yet well developed, and their metacognitive judgments of all kinds may be fragile. And laboratory studies with adults may be questionable as well, because the stakes are typically low. Perhaps Chris's high confidence in his false memory of being lost in the shopping mall resulted, in part, because it didn't really matter whether or not he had been lost in the shopping mall. Perhaps Oskamp's participants also knew that it really didn't matter, so why work that hard to be accurate?

What if it did matter? What if the events were emotional? Life threatening? What about cases where the memory is one of a trauma, as, unhappily, is often the case in criminal investigations? When a person is tortured or abused, our folk knowledge is that the victim's confidence about the identity of the perpetrator will be high and that he or she will be right. These are the excruciatingly painful situations that can't be ousted from memory, in which people could *never* forget a face, much as they might wish to. But is this folk knowledge correct?

Few studies of eyewitness memory have been conducted under conditions of traumatic stress, for obvious reasons—extreme trauma is not something that can be easily produced in the lab. Studies in which people are shown violent films are often used to study whether stress enhances memory or not, and whether people's confidence is appropriately enhanced. The American Academy of Pediatrics (2001) has estimated that by age 18, the average person in the United States will have viewed 200,000 acts of violence on television alone. It seems likely that we have habituated to filmed violence, and at least for some people, it is not really that stressful. Indeed, one of the most violent films to date (*The Shining*) to which a hormonal stress response has been measured showed that stress levels, while somewhat elevated, were well within normal diurnal variation (Hubert & de Jong-Meyer, 1989). So films are unlikely to tap into situations where people would “never forget that face.” But it is extremely difficult to test people in real-world traumatic situations.

Unfortunately, there is now plenty of real-life evidence that extremely confident witnesses have been mistaken, with dire consequences. These cases are invariably ones in which there is no physical evidence, but a witness confidently “tells the truth” about who committed some heinous crime, such as assault, robbery, or murder. Based on the witnesses' confidence in their memory, the jury votes to convict and the alleged criminal wastes away in jail, on death row, or worse. Later, DNA evidence saved from the scene of the crime has shown—without any doubt—that the alleged culprit was innocent. In a recent case, Charles Chatman served nearly 27 years of a 99-year sentence for aggravated sexual assault before he was released based on the outcome of DNA tests. In the past decade, over 30 cases like this one led to the exoneration of inmates based on the outcome of DNA tests—in the state of Texas alone. For more facts about misplaced confidence that led to wrongful convictions (that were later overturned by DNA tests), go to the Innocence Project on the World Wide Web.

We were able to find one study that experimentally investigated eyewitness identification of the perpetrator in a truly traumatic situation, where it was known with certainty what had actually happened and who the “aggressor” was. This study tapped a very special situation—soldiers' training in a

simulated prisoner-of-war camp. Morgan et al.'s (2004) investigation of eyewitness identification and people's confidence in that identification was made under conditions that, as evidenced by the extremely elevated levels of cortisol, were at trauma levels equivalent to those experienced by people undergoing open heart surgery or combat. Cortisol is a stress hormone that is released in response to either a physiological stressor, or, in the case of humans, a psychological stressor. Levels of cortisol give some grounding to our assessments of how stressful a situation is. Cortisol levels are normally quite low. A scary movie may cause the levels to triple—but this is still fairly low. The same increase occurs if people are asked to do spontaneous public speaking during which their intelligence will be evaluated. (This frightening task makes many people's stomachs turn over just to think about it, and is called the Trier Social Stress Test, or TSST.) Such situations are about the limit for normal psychological investigations, but these do not come close to producing the levels of cortisol seen in real trauma. The levels of cortisol in Morgan's study, however, were roughly 10 times higher than these already inflated levels, and so give a very good approximation of the condition one would experience, physiologically, during trauma.

The study site was a mock prisoner-of-war camp designed to train U.S. soldiers how not to capitulate should they be captured. The army trainers were realistic in the stresses to which they exposed the men, including subjecting them to situations modeled after those that American prisoners of war reported after the Korean War. According to Morgan et al. (2004), after some classroom style training,

participants are confined in a mock prisoner of war camp (POWC). This phase is designed to offer one of the most challenging training experiences that active duty participants will ever experience while in the military. In the POWC, each participant is placed in isolation and then subjected to various types of interrogation. These interrogations are designed to test the limits and abilities of the participants to withstand "exploitation by the enemy." (p. 266)

The exact procedures used are classified information, so the authors were unable to offer more details, except that people undergoing this training exhibited levels of stress hormone beyond anything seen in everyday life. They also showed astonishingly high levels of dissociative symptoms, such as out-of-body experiences and psychotic-like behavior.

During the effort to extract information from the trainee in the prisoner-of-war situation, there were two interrogators: the high-stress interrogator, whose job, in part, was "physical confrontation" with the participant, and the low-stress interrogator, who did not threaten at all, and indeed,

“befriended” the participant (but who did, at the same time, try to trick the participant into giving away information). The duration of exposure between these people was more than just a brief glimpse, as sometimes happens in experiments on eyewitness identification. Each of the two interrogators had 40 minutes of face-to-face contact with each participant.

One might suppose that the 509 participants would never forget the high-stress, violent interrogator (though they might possibly forget the unthreatening one). Just the opposite was found. Three methods of eyewitness identification were used: (a) a live lineup, (b) a photo lineup with all potential interrogators shown simultaneously (a simultaneous lineup) and the participant choosing from the group, and (c) a photo lineup with each potential interrogator shown individually (a sequential lineup) and the participant stating whether each photo was an actual interrogator or not. The likelihood of correctly identifying the interrogator (i.e., the proportion of hits) is presented in Figure 8.2, which clearly shows that participants had better memories for the low-stress interrogator than for the high-stress interrogator, with hit rates for the latter being astonishingly low. Thus, the victims identified the interrogator very poorly when the event was traumatic. But when they did correctly identify him, did they do so with extremely high confidence, as folk knowledge suggests? The answer is no. The mean confidence ratings about the accuracy of their reports (made on a scale from 1 to 10, with 10 being highest) were 6.2 for the high-stress interrogator and 7.9 for the low-stress interrogator.

Does Witness Confidence Matter to Jurors?

Perhaps the jury ultimately doesn’t care about the witnesses’ confidence. Unfortunately, given the evidence that a clever attorney can easily manipulate confidence and that it is undiagnostic, that does not appear to be the case. As in daily life, other people’s confidence is taken by jurors to be of prime importance in evaluating the reliability of a witness and in making their decisions. A number of studies have investigated people’s responses to other people’s confidence.

Cutler, Penrod, and Dexter (1990), for example, conducted a mock-jury study to examine juror sensitivity to eyewitness identification evidence. Participants were eligible and experienced jurors from Dane County, Wisconsin, who viewed a videotaped trial that involved an eyewitness identification. The responses of experienced jurors were also compared to those of undergraduates, but differences between the undergraduates and

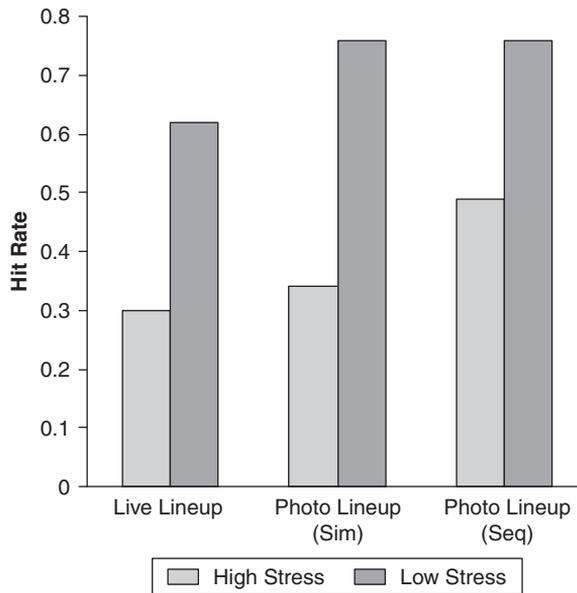


Figure 8.2 The probability of correctly identifying an interrogator during a live lineup, a photo lineup where all photos are presented simultaneously (Sim), or a photo lineup where all photos are presented sequentially (Seq).

SOURCE: Data adapted from Morgan et al. (2004).

the eligible jurors in their sensitivity to eyewitness evidence were negligible. Ten factors associated with the crime and the identification were manipulated—such as whether or not the perpetrator was disguised as well as what the witnesses' confidence was in their memory of the crime. The result of main interest was that the confidence of the eyewitness was a more powerful predictor of the verdicts people returned than any of the other nine factors!

Brewer and Burke (2002) constructed a mock-jury situation in which they manipulated both the consistency of the testimony given by a witness and his or her confidence. Surprisingly, consistency mattered little, but confidence, once again, had a strong influence on jury decisions. As disturbing, Fox and Walters (1986) showed that eyewitness confidence had a strong impact on jury decisions, even in the face of conflicting expert testimony.

There are, however, some boundary conditions. Tenney, MacCoun, Spellman, and Hastie (2007) showed that although confident witnesses are deemed more credible than unconfident ones in general, should such a confident witness be caught in a mistake, less-confident witnesses might appear more credible than more-confident ones. Thus, although mere confidence does have an effect, juries are not entirely insensitive to other factors that indicate the credibility of a witness. The problem, of course, is that in most courtroom situations, jurors do not know what is accurate. Thus, the many studies showing the overwhelming impact of confidence on the verdicts that juries return, makes the unreliability and the manipulability of this metacognitive judgment highly worrisome.

Lying

In all of the above discussion of metacognition in the courtroom, we have assumed that people were acting, remembering, and judging in good faith—that they were trying to be honest. But what if they weren't? What if they were lying? Despite perjury laws, there are, of course, many reasons that a person might lie. In this section, we will examine several aspects of lying relevant to the courts.

Deception is an act that is intended to foster in another person a belief or understanding that the deceiver considers false. It consists of both a communication of information and a metacommunication about the sincerity of the message, that is, the message is false, but the communicator intends to instill the belief that it is true. The recipient hears the message and also makes a judgment about its truth. Thus, lying and detecting people's lies will draw upon people's metacognitive abilities. The first question we address is whether people are able to detect when other people are lying.

Can People Detect Lying?

Ekman and Friesen (1969) refer to behaviors that give away a falsehood as *deception cues*, and those that reveal the true information, *leakage cues*. In some intermediate cases, it is not clear that the situation is one of deception, such as self-deception, intentionally transparent lies (where the sender wants the receiver to ascertain that the communication is a lie), and mistaken lies (where people think they are telling a lie but it really is the truth, or vice versa). Here, we will exclude consideration of these complicated cases and focus only on what is considered barefaced lying. When people engage

in such lying, are there clues that give them away, such that a jury, or an experienced investigator, might be able to tell?

It has turned out to be notoriously difficult to determine when a person is lying. The folk wisdom is that people can control their verbal but not their nonverbal behavior. If true, one should study body language and voice characteristics, not the words that are spoken. Many people think (and there is some empirical support) that lies can be detected with tics of the body, more so than the face, presumably because the body is less controllable. But there are also facial clues that provide a giveaway, although it is not the large expressions, which tend to be controllable by the liar. Rather it is the facial microexpressions—very fast muscle movements—that are inconsistent when a person is lying. Such microexpressions, however, can generally only be seen with a slow-motion replay (although some people can detect them). In addition, the voice is particularly difficult to control. Thus, a good lie detector may pay particular attention to the manner of speaking rather than the content of the speech.

Because lying is apparently more cognitively demanding than telling the truth, people often pause more when they are lying. In some studies, brow furrowing has been correlated with lying, as have speech errors, and shrugs (Zuckerman, DePaulo, & Rosenthal, 1981). Although all of these factors have some relation to lying, their diagnosticity is hardly overwhelming. An early study by Zuckerman et al. (1981) showed that the correlation between people's confidence in the accuracy of their own ability to detect lies and their actual accuracy had a median value of .06, which was not different from zero. The inescapable conclusion is that although liars may show subtle signs of lying and most people think they can detect these, they can't.

The folk view of lying also favors some blatant misconceptions. Surveys indicate that people (both lay people and police) believe that liars avert their gaze and fidget (Akehurst, Köhnken, Vrij, & Bull, 1996). This notion is promoted in Inbau, Reid, Buckley, and Jayne's (2001) interrogation manual, a source that has been widely used by police departments across the United States and by other interrogators. As Mann, Vrij, and Bull (2002) have shown, however, police officers who endorse this view are among the worst at detecting liars, because the cue is simply undiagnostic: Liars don't fidget any more than do truth tellers.

Some studies (e.g., DePaulo & Pfeifer, 1986; Köhnken, 1987; Kraut & Poe, 1980) have reported that people in general have no ability to detect deceit. But can *anybody* reliably detect when someone else is lying? Ekman and O'Sullivan (1991) investigated individual differences in lie detection. They looked at the performance and the confidence of secret service observers, federal polygraphers, robbery investigators, judges, psychiatrists, special

interest participants (who were taking a course on lie detection, but before they had explicit instruction on diagnostic factors), and college students. Although it might be thought that several of these groups would be especially good at this skill, only the secret service participants were accurately able to detect liars.

All of the groups in Ekman and O'Sullivan's (1991) study were asked twice to rate their ability to detect lies, first before seeing the videotapes that constituted the test, and then afterwards, when they were specifically asked about how well they had done. The correlation between metacognition of general ability and actual performance was not different from zero: Overall, their perception of how well they would do was not related to how well they did do. Similarly, overall retrospective confidence about performance also had a zero correlation with performance. Across all of the groups, then, people's metacognitive assessments of their ability were undiagnostic.

Some group differences, however, were embedded in this overall zero correlation. The federal polygraphers' initial metacognitive ratings of their general ability to detect lies were positively correlated with their actual ability to do the task. Perhaps they had had enough feedback in their jobs to know whether they were good at detecting or not. Perhaps even more interesting, however, was a negative correlation in the secret service group. Within the one group of experts in this study who (as a group) could do the task, the metacognitions were backwards: secret service officers who thought they couldn't detect deception were better at detecting it than secret service officers who thought they could.

Lies Become Truth, or the Frequency-Validity Relation

Although people have great difficulty determining whether or not other people are lying, they might still be able to determine whether particular statements are true or not. And, of course, when a statement is entirely implausible, people will know that it is false, at least some of the time. Most research on this question has addressed statements in the murky middle, where most people do not know for sure whether statements are true or not. It was first empirically demonstrated by Hasher, Goldstein, and Toppino (1977) that the mere repetition of statements that were untrue could result in people coming to believe that they were true. This failure reflects a metacognitive deficit because when people encounter the false statement, their familiarity with it makes them judge it as true.

In the Hasher et al. (1977) experiment, participants rated the truthfulness of 60 statements. The statements themselves came from knowledge domains of politics, sports, religion, the arts, and so forth. The statements were along

the lines of “Lithium is the lightest metal” or “The total population of Greenland is about 50,000 people.” Some of them were true and some of them were false, but people were unlikely to know for sure. Validity ratings were made by participants after the tape-recorded presentation of each statement. There were three sessions in the experiment. The twist was a simple one: Over the three sessions, some of the statements were repeated whereas others were not. The results showed that the true statements were rated as more true, when they were repeated. Unfortunately, so too were the false statements. Whereas the false statements, on the first repetition, were close to the “uncertain” boundary, with a rating of 4.18 (on a 7-point scale, with 7 = definitely true, 4 = uncertain, and 1 = definitely false), by the third session (after multiple repetitions) they had crept up half a point to 4.67.

This result has now been replicated many times, with a variety of researchers showing what has come to be known as the “frequency-validity” relationship. In response to concerns about the external validity of laboratory findings using college students to “real people in the real world,” Gigerenzer (1984) telephoned people in Schwabing, a community on the outskirts of Munich, Germany, and asked them to answer questions, similar to those of Hasher et al. (1977)—giving their validity ratings—over the phone. The results were the same as other researchers have found with laboratory studies. Reber and Schwarz (1999) showed that by increasing the perceptual fluency of statements, by having them be easy or difficult to read against the background, they were also able to alter people’s judgments of their truth. Those statements that were easier to read were assessed as being truer. And there are many other examples of these kinds of surprising, and arguably distressing, effects—that is, where increasing the familiarity of a (false or true) statement enhances its believability. For instance, lies that are repeatedly told about political candidates during campaigns (e.g., she’s a flip-flopper or he’s in bed with corporate criminals) are eventually believed by many people who hear them but never attempt to evaluate whether they’re true.

Whereas many studies have shown that lies may become truth with repetition, the final study we mention in this section showed that liars may also come to be perceived as truth tellers, with repetitions. Brown, Brown, and Zoccoli (2002) showed that the more times people saw a photograph of a person, the more credible they found the person to be. The credibility increase was comparable for both judgments of honesty and judgments of sincerity at short (2-day) and long (14-day) intervals. The effect depended on repetition, but not explicit recognition that the faces had been seen before. So, the more often you see liars tell lies, the more likely it is you will believe them.

Hindsight Bias

The final topic that we will touch on in this chapter is *hindsight bias*, or what is sometimes called the “knew it all along” effect. This effect has profound implications for the criminal justice system, for issues as wide-ranging as whether jurors are able to disregard testimony that has been ruled inadmissible, or whether medical practitioners are liable for adverse outcomes in malpractice suits. Baruch Fischhoff was first to report this phenomenon in 1975. As we will see, one of the explanations for this phenomenon is considered metacognitive. But before we get into that, let’s first take a look at what Fischhoff did that created such a stir.

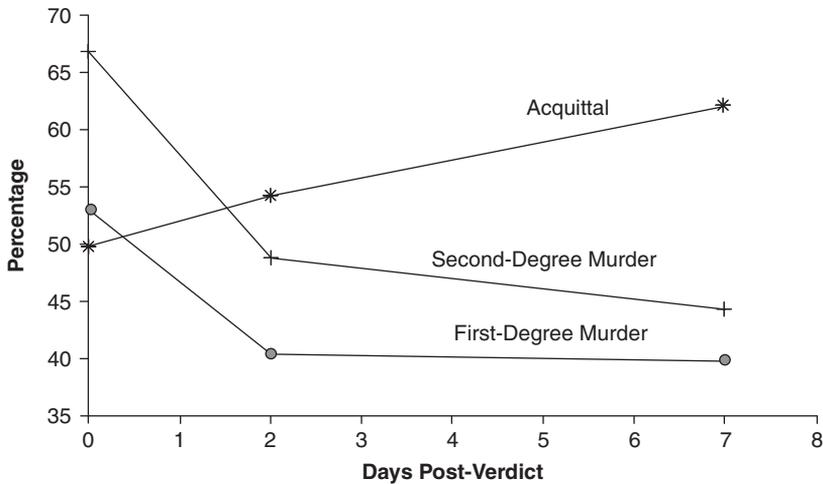
In a seminal experiment that has generated hundreds of subsequent studies in a wide variety of domains, Fischhoff simply informed his experimental group about the outcome of an event and then asked them to say what they thought the outcome would have been if they had not known already what the outcome was. The targeted event, in Fischhoff’s original experiment, was the 19th-century war between the British and the Gurkhas. Of course, people who actually knew the outcome had to be eliminated, but there were few such people. The subject was chosen for just that reason—to be something that people wouldn’t know much about. (Notice that this situation is also true for jurors hearing cases. People who know or think they know a lot about a case are excluded. If a crime occurs in your neighborhood, or if you know people involved, you will probably not be allowed on the jury.) Fischhoff developed four possible outcomes (e.g., Gurkhas won, or the British won), and—in a counterbalanced manner—he told different participants that each of these four possible outcomes had been the actual outcome. In the control condition, no outcome information was given. After giving (or in the case of the control condition, not giving) the outcome information, he asked all participants to estimate a probability for the likelihood of each of the four possible outcomes, but to do so retrospectively, *as if they did not know the outcome*. The result was that people were unable to keep from biasing their responses about the probabilities in favor of what they “knew” to be the actual outcome. That is, they judged the outcome they had been *told* was the actual outcome to be more likely than it actually was (as indicated by the control condition).

An example of hindsight bias in the real world comes from a study by Bryant and Brockway (1997), who made use of the notorious O. J. Simpson case. They asked the participants, 2 hours before the verdict, 48 hours after the verdict, and 1 week after the verdict, to indicate the chances that the jury would find Simpson (a) guilty of first-degree murder, (b) guilty of second-degree murder, or (c) not guilty. In the post-verdict sessions, participants

were asked about their knowledge of the verdict (100% of the participants knew that he had been acquitted) as well as whether they believed that Simpson was guilty (83% said yes) or innocent (17% said yes). Then, participants were asked for their retrospective judgments of the jury's vote. (The researchers did not constrain the participants to make these three outcomes add up to 1.) The results are shown in Figure 8.3. People were much more likely to judge that Simpson would be acquitted in the post-verdict sessions (and hence after they knew that he had been acquitted) than before they knew the verdict in the case.

Why Does Hindsight Bias Occur?

The hindsight-bias effect has been explained from three viewpoints: the personal needs view, the memory view, and the anchoring view. The first explanation—the personal needs view—is based on the idea that people like to be right. But for some people this need to be right, and to have been right, is stronger than it is for others. In an early study, Campbell and Tesser



Day 0 = 2 hours before the verdict.

Day 2 = 48 hours after the verdict.

Day 7 = 1 week after the verdict.

Figure 8.3 Estimates of the chances of conviction and acquittal in the O. J. Simpson criminal trial, over time.

SOURCE: Bryant and Brockway (1997).

(1983) found that people with certain personal needs and personality characteristics were more prone to claim that they “knew it all along.” The researchers first attained people’s scores on the Rokeach Dogmatism Scale—a questionnaire that measures people’s intolerance for ambiguity. The more participants could not tolerate ambiguity, the greater their hindsight bias.

Other researchers have shown similar effects. For example, Musch (2003; also see Musch & Wagner, 2007, for a review) showed that people who were high on field dependence showed greater hindsight bias than did field-independent people. He also showed that a tendency toward favorable self-presentation, as well as rigidity or the need for predictability—both variables similar to those measured by Campbell and Tesser (1983)—were also related to greater hindsight bias. There is also a tendency for children and older adults to show stronger hindsight bias than young adults (Bayen, Pohl, Erdfelder, & Auer, 2007). The age-related differences, however, may or may not be due to personal needs or personality variables, because the cognitive variables such as memory, which we will discuss below, are also different for children and older adults as compared to young adults. In any event, whereas personality differences do have an effect on hindsight bias, they cannot be the whole story. Given that the effect is found even in people lacking the kinds of personal needs that have been thought likely to bias their outcomes, cognitive factors are implicated in the explanation of the effect.

The second explanation is purely cognitive—based on the idea that different sources of information in memory may be blended together. As a number of researchers have pointed out, the structure of the hindsight-bias situation is similar to that of the memory-based misleading-information paradigm. People see or witness an event, and then, afterwards, are given information about that event. In a classic example of the memory paradigm, people see a blue car and are later given the suggestion that it was green. When they are asked to choose the color of the car, they choose blue-green (Loftus, 1977). In the hindsight situation, people are given a question such as, “How many countries are there in Africa?” Perhaps they think it is 45. Then they are told that the actual answer is whatever the experimenter chooses—say 59. When asked about their initial estimation, they give a compromise somewhere between 45 and 59—maybe 52. The memory explanation postulates that in both cases the memory for the original event is blended with that of the later information or that there is some sort of interference or distortion from the later information that may impair accurate retrieval of the earlier information.

The third explanation of the hindsight-bias effect is that it is an anchoring effect. The anchoring explanation of hindsight bias is simple—it says that the outcome (e.g., you are told that O. J. Simpson was acquitted) serves as an anchor, and later estimates of the probabilities of events are pulled

toward this anchor. Anchoring effects in general (i.e., the tendency of a given quantity or entity to become a beginning point for later judgments) have been shown to be pervasive in human judgment and decision making. This explanation is viewed as metacognitive in nature, because it emphasizes how anchors unduly influence people's evaluations of their memories about their original beliefs.

Tversky and Kahneman (1974) have shown how numbers exert anchoring effects even when participants see researchers spin a wheel of fortune, leaving no doubt that the number provided is random. Some savvy real estate brokers and car dealers use this ploy regularly—showing clients an expensive house or car at the outset to get the customer to pay a higher price; that is, enticing customers to opt to buy a more expensive house (or car) or pay more for the same one.

The second and third explanations are difficult to tease apart. Indeed, McCloskey and Zaragoza (1985) have proposed that the effects found in the misleading-information paradigm in memory might really be due to a bias or anchoring effect, rather than being a memory effect. Pohl and Gawlik (1995) have attempted to distinguish the two cognitive explanations of hindsight bias by using a Markov model that proposes different processes for the two explanations. Even they admit, however, that the data in the two paradigms look remarkably similar, and that the distinction between them may be model dependent and not robust. Thus, we do not know whether what happens is that people's memory is changed by the outcome information, or the outcome information exerts a bias or pull on a decision process. Nevertheless, the empirical result itself is rock solid: Knowing the outcome of an event has an irreversible effect on what people think they would have believed had they not known the outcome. This "knew it all along" effect has legal consequences, some of which will be reviewed in the next section.

Confessions, Inadmissible Evidence, and Hindsight Bias

One of the most interesting hindsight situations occurs in cases of coerced confession. A famous case, relevant to this issue, was the murder trial *Arizona v. Fulminante*, in which Fulminante confessed. The confession was entered into testimony. But it was later ruled that the confession had been coerced. A coerced confession is, of course, inadmissible. Circumstantial evidence suggested that he had committed the murder, but, even so, confession has an enormous impact upon a jury. The issue was whether the circumstantial evidence alone would have resulted in a conviction had the confession not been presented. This case went to the Supreme Court, which ruled

that the introduction of the confession constituted a “harmless error,” and let the verdict stand.

In some states, rather than having a judge in a pretrial hearing assess whether or not a confession was coerced, the jury members are asked to hear the evidence and make the judgment. They are then supposed to disregard the evidence if they determine that the confession was coerced. This raises two questions: First, can people make accurate judgments about whether a confession was coerced? This is a fairly standard question about people’s metacognitions, although one about the circumstances surrounding someone else’s behavior rather than about their own learning or memory. Second, can people use their metacognitions to appropriately control their memory, and ultimately, their verdict?

An experiment addressing these issues was conducted by Kassin and Sukel (1997), using mock jurors who were given a confession that was elicited under low or high pressure. The mock jurors correctly realized that the high-pressure confession was coerced, and they remembered the (mock) judge’s admonition to disregard it. But even though they said they would disregard it in their judgments, their conviction rates indicated otherwise. People who were on the mock juries in which no confession was given turned in a conviction rate of 19%. The conviction rate was 63% in the low-pressure (admissible confession) condition. In the high-pressure inadmissible confession condition, the conviction rate was lower than it was with the admissible confession, but it was still 44%—much higher than the rate given when no confession had been heard. Thus, although people knew that the testimony was inadmissible and they thought they could disregard it and, indeed, they thought that they *had* disregarded it, they were unable to do so.

Malpractice, Liability, and Hindsight Bias

By now you can probably infer the implications of hindsight bias for medical malpractice. Once a person suffers a negative medical outcome, anyone assessing the situation will think that the doctors involved should have known that there was a high likelihood of it occurring, so if only they had taken the proper steps, it could have been avoided. Malpractice. In medical circles this bias is known as the retrospectroscope! Many think that hindsight bias is largely responsible for many malpractice claims and for the runaway costs of physicians’ medical-liability insurance.

Retrospective blame can also be an issue for mental health professionals. Consider the following case. In the late 1960s, Tatiana Tarasoff and Prosenjit Poddar dated several times. When Tarasoff began to date other

men, a distraught Poddar sought counseling. He divulged his intention to kill Tarasoff to his therapist, who had him detained by police. Poddar, however, was released after he agreed not to pursue Tarasoff. Later, Poddar tracked down Tarasoff and stabbed her to death. No one had been warned that Poddar posed a threat to her life! Subsequently, the 1974 Tarasoff decision established that therapists have a duty to warn potential victims that a patient may be dangerous. If therapists do not use reasonable care in such situations, they may be liable for negligence. When a patient becomes violent, the perceived probability of violence shifts from the time of the assessment by the therapist to the time of the assessment by the judge hearing the case. The law—not taking hindsight bias into account—says the therapist should have known. LaBine and LaBine (1996) have documented many instances of this phenomenon. Similarly, people blame their financial woes on auditors: They should have known, and they may face liability suits for not having said so. The evaluative judgments of auditors' performance, however, may be due to hindsight bias (Anderson, Lowe, & Reckers, 1993).

Other Examples of Hindsight Bias

We have focused, in this chapter, on the implications of hindsight bias for decisions made in the courtroom. We would be remiss, however, to close without noting that the effect is pervasive. Once people know the answer to anything, they think that it was obvious. After (but not before) solving an insight problem, people think that it was easy. They may even think that a student who is still struggling to find the answer—as they did—is deficient in some way.

And woe to the brilliant individual who makes a new discovery. In science, such a true discovery is too often belittled by naysayers who claim that it's obvious. Obvious in retrospect! To thwart such naysayers, one can make them state *their* predictions before revealing one's discoveries. Unfortunately, however, hindsight is likely to distort even their memory of what they predicted. Similarly, in politics, once the outcome of an election is known, those pundits who (by accident, perhaps) said the “right” thing only look normal and rational. They don't look like prophets, because everybody knew it all along. But those who, with the same evidence, made a different assessment have egg on their faces.

Several researchers (Ofir & Mazursky, 1997; Sanna, Schwarz, & Small, 2002) have begun an effort to look at situations in which people claim not “I knew it all along,” but rather “I *never* would have known that.” We

applaud this new line of research. But we are nevertheless willing to bet that once the answer about “we *never* would have known *that*” is finally known, most of us—the colleagues of these brave explorers—will say that we knew it all along.

Summary

People often confidently believe in their memories of past events. Such high confidence in memory is often not misplaced. In our everyday lives, many of our memories are valid and we should be confident in them. Unfortunately, memory can too easily be altered in numerous ways: You may remember that a person was present at a crime scene because you later saw the person walking down the street. You may recall an event that was merely suggested to you during repeated interviews with police. You may come to believe that an event represents a true memory because you were forced to confabulate that memory. In many of these cases, people will be highly confident in their false memories and may stick tenaciously to them even though hard physical evidence suggests otherwise.

Perhaps worse, misplaced confidence in invalid memories—as well as hindsight bias that people have difficulties escaping from—can have dire consequences in many settings. These consequences are perhaps most evident in the judicial system in which confident witnesses are readily believed (even without physical evidence to support their testimony). Society as a whole is largely ignorant of how mere repetition can lead to false memories—that is, most people lack the metacognitive knowledge about how their minds operate and how easily they can be tricked. Besides introducing some general issues relevant to metacognition and the law, we hope this chapter will make you think twice before you swear “to tell the truth, the whole truth, and nothing but the truth,” or fully believe that the testimony you are considering—even if from a witness testifying in good faith—must be true.

DISCUSSION QUESTIONS

1. In this chapter, we discussed a variety of situations in which people were led to develop a false memory of the past: People remembering a yield sign (after one was suggested to them) when in fact they saw a stop sign, and a child being led to believe he was lost

at a mall when in fact he never was lost. Do you believe that people can be made to believe that *any* kind of event happened to them? Which kinds of event do you think would be difficult to implant, and for what kinds of people?

2. Morgan et al. (2004) reported what you might consider to be a rather surprising finding: When people were interrogated, their ability to later identify the interrogator was much worse with an interrogator who evoked a lot of stress than with an interrogator who evoked little stress. One might expect that people would remember the stressful interrogator, as if the stress would focus their attention on the person who was being abusive. Think of several explanations for why people might have difficulty remembering a stressful interrogator. How could you evaluate (using an experiment) your explanations for why stress can impair memory?

MINDS-ON ACTIVITY: IMPLANTING FALSE MEMORIES

Using a simple method (Roediger & McDermott, 1995), it is relatively easy to make people recall words that they did not originally hear. In many cases, the victims are surprised to learn their memories are not valid. To demonstrate this false-memory effect, you will probably only need one participant—such as a friend or family member—although the demonstration can be attempted with multiple people at once. Here's what you need to do. For each list below, read one word at a time (for about 3 seconds a word), and tell your participants to try to remember them for an upcoming test. After you've finished reading each list, have them count backwards from a 3-digit number (e.g., 475 or 899) by threes for 30 seconds or so. This rehearsal-prevention task will ensure that they just can't repeat back the final few words on the list. After 30 seconds of counting backwards, have them (a) write down every word they can remember, and (b) make a retrospective confidence judgment for each word that they recalled (e.g., from 0 to 100, where 0 means they have no confidence that the word was on the list and 100 means they are 100% confident it was there). Now, repeat for the other lists.

When you're finished, check to see if they recalled the seductive lures: sleep (for List 1), needle (List 2), and sweet (List 3). Were they just as confident in their memory of the words that were on the list as they were for the seductive lures? How might you help people to reduce the likelihood of recalling the seductive lures? If you have ideas, try your experiment again (but with new participants) to find out if your technique actually reduces false memories.

<i>List 1</i>	<i>List 2</i>	<i>List 3</i>
bed	thread	sour
rest	pin	candy
awake	eye	sugar

tired	sewing	bitter
dream	sharp	good
wake	point	taste
snooze	prick	tooth
blanket	thimble	nice
doze	haystack	honey
slumber	thorn	soda
snore	hurt	chocolate
nap	injection	heart
peace	syringe	cake
yawn	cloth	tart
drowsy	knitting	pie

SOURCE: Lists are a subset from Roediger and McDermott (1995).

CONCEPT REVIEW

For the following questions and exercises, we recommend you write down the answers on a separate sheet in as much detail as possible, and then check them against the relevant material in the chapter.

1. How has DNA evidence been informative with respect to showing that witnesses' extreme confidence in their memories can be entirely misplaced?
2. What is hindsight bias, and what are the most prominent explanations for it?
3. How accurately can people detect when other people are lying? Why might lie detection be so difficult?
4. Describe evidence that indicates that in fact jurors are influenced by the confidence of eyewitnesses' testimony.