

Chapter 5

When We Ignore Unethical Behavior

Since the 2008 financial collapse, fingers have pointed in many directions. Targeted guilty parties have included irresponsible banks, greedy home-buyers, speculators, the Democratic Congress (for pushing to give low-income borrowers too much credit), and the Bush administration (for poor decision making and regulatory neglect). But at least part of the problem stems from the failure of independent credit-rating agencies to appropriately rate the riskiness of the mortgage-backed securities they assessed. “The story of the credit rating agencies is a story of colossal failure,” according to Representative Henry Waxman (D-CA), chairman of the House Oversight and Government Reform Committee.¹ Waxman’s committee found strong evidence that the executives in charge of the rating agencies were “well aware that there was little basis for giving AAA ratings to thousands of increasingly complex mortgage-related securities, but the companies often vouched for them anyway.”

The purpose of credit-rating agencies is to educate outside stakeholders of the creditworthiness of issuers of debt obligations (including companies, nonprofit organizations, and federal, state, and local governments) as well as the debt instruments these financial organizations sell to the public. These agencies exist because of their presumed objectivity, yet their compensation has been tied to anything but objectivity. Former high-level employees of the rating agencies testified before Waxman’s committee that a conflict of interest exists in the U.S. credit-rating system. Specifically, the largest credit-rating agencies—including Standard & Poor’s, Moody’s, and Fitch—are paid by the companies they rate instead of by the investors who have the most to lose from inaccurate ratings. The largest ratings agencies have made enormous profits by giving top ratings to securities and debt issuers, and not necessarily by providing the most accurate assessments of these securities and issuers. In addition, agencies with the most lax standards have been, not surprisingly, best at winning business from new clients, giving the agencies financial incentives to positively assess securities. Compounding the problem further, the rating agencies have been selling consulting services to the same firms whose securities they have been rating.

It may seem obvious that if rating agencies have an incentive to please the companies they assess, an environment emerges in which independent, unbiased assessments are no longer possible. Yet not everyone believes this is an obvious conclusion. Defenders of the rating agencies have argued that the agencies’ knowledge of the importance of ensuring a firm’s integrity would protect them from issuing biased assessments. This belief, while admirable, is overly optimistic. Worse yet, it prevented society from seeing the unethical behavior of the parties involved. Just as the federal government failed to address the inherent conflict of interest in the auditing industry in the pre-Enron era, our leaders failed to make changes to the credit-rating industry that might have headed off disaster. In both cases, the unethical behavior

of others appeared opaque to many people.

There are many reasons why we do not notice the unethical behavior of others. To begin with, we are busy paying attention to other things. As we will discuss in more detail in [Chapter 6](#), we pay attention to goals for which we receive rewards and too often ignore those for which we do not. We are not usually rewarded for noticing the unethical behavior of others. What's more, human beings have a remarkable ability to overlook the obvious. In one study, psychologist Ulric Neisser asked his Cornell undergraduate students to watch a video in which two visually superimposed groups of three players were passing basketballs.² One trio wore white shirts, and the other trio wore dark shirts. The students in Neisser's study were instructed to count the number of passes made among the trio wearing white shirts. The dual video, as well as the grainy nature of the film, made the task moderately complex. Before reading on, feel free to watch the video and try to accurately count the passes among players wearing the white shirts at www.blindspots-ethics.com/neisser.

As you may have guessed, this is a trick experiment. While you were busy counting passes, you—like most people who try this task—probably failed to see a woman who clearly and unexpectedly walked through the basketball court carrying an open umbrella. (If you don't believe she was there, go look again.) Only one in five of Neisser's Cornell undergraduate participants spotted the woman with the umbrella. When we show this video in our classrooms to MBA and executive students, far fewer than one in five people notice the woman, just as we failed to notice her when we first watched the video. Because they are focusing closely on one task—in this case, counting passes—people miss very obvious information in their visual world.

Neisser's video offers evidence that our focus on one set of tasks can blind us to other readily available information in our environment. Moving beyond simple busyness and distraction as an answer, using the lens of behavioral ethics, this chapter maps the multiple reasons why we overlook the unethical behavior of others. Why do we look the other way when, objectively, it should be clear to us that someone is doing something wrong? We begin by discussing the role of motivated blindness, or the tendency for people to overlook the unethical behavior of others when it is not in their best interest to notice the infraction. Second, we explore indirect blindness, or the tendency not to notice unethical actions when people do their dirty work through the behavior of others. Third, we examine the role of a slippery slope in noticing the unethical behavior of others. Finally, we examine how the tendency to value outcomes over processes can affect people's assessments of the ethicality of others' choices.

Motivated Blindness

In the controversial 2008 fictional film *The Reader*, based on the novel of the same name by the German writer Bernhard Schlink, an illiterate former Nazi guard, Hanna Schmitz, faces charges in a war crimes trial for a terrible episode she took part in during World War II. Schmitz and five other female SS guards were leading hundreds of prisoners in a death march

in 1944. One night, the prisoners camped in a church for shelter, and the guards locked them in. The church was bombed and caught fire, but none of the guards unlocked the doors, and the three hundred prisoners burned to death inside.

Not only did Hanna fail to save the three hundred prisoners, but she also testified that, during the war, she had followed orders and chosen ten prisoners to send to the gas chambers of Auschwitz each month. When asked during the trial about her failure to unlock the doors of the church, Hanna (played by actress Kate Winslet) gives the judge a confused look. "Obviously," she says matter-of-factly, "for the obvious reason: we couldn't. We were guards. Our job was to guard the prisoners."³ If the guards had freed the prisoners from the burning church, Hanna explains, they would not have been able to control the crowd. In the chaos, the prisoners would have escaped, and Hanna would not have performed her job properly. Pressed further to explain why she failed to free the prisoners, Hanna shouts, "We were responsible for them!" Bewildered, she asks the judge, "What would you have done?"

We have no interest in defending the actions of this fictional character. But the portrayal of Hanna in *The Reader*, to the consternation of some Jewish groups that criticized the story, suggests that the character did terrible things without recognizing the ethical implications of her actions. She was uneducated, grew up following the orders of superiors, took a position with the SS for practical reasons, and simply did not see the option of freeing the prisoners trapped in the burning church. In *The Reader*, Hanna accepts her fate (prison), but throughout much of her life, she fails to view her own behavior as unethical.

Hanna's behavior—and her denial that she did anything wrong—is an extreme case, and it is a fictional case. However, we argue that Hanna's lack of recognition parallels what many do wrong in the interest of their group, organization, or country. This behavior is consistent with emerging evidence that significant numbers of people are capable of engaging in massive harm without realizing they are doing so. In a 2009 study of 2,800 employees, 49 percent reported they had observed some type of wrongdoing on the job in the previous year, despite the considerable efforts that organizations are taking to improve their employees' ethical behavior. Unfortunately, wrongdoing isn't a new fad: The ethical scandals at Arthur Andersen, Enron, Health-South, Tyco, and WorldCom were preceded by earlier ethical scandals at General Electric, Investors Overseas Services, Lincoln Savings & Loan, Sears, and Shoney's.

Throughout this book, we have noted a core finding of behavioral ethics: that people who have a vested self-interest in a situation have difficulty approaching the situation without bias, even when they view themselves as honest. Here we argue that this bias extends to the observation of others: that is, if you are motivated to turn a blind eye to someone's unethical behavior, you won't see it. The term *motivated blindness* describes the common failure of people to notice others' unethical behavior when seeing that behavior would harm the observer. When party A has an incentive to see party B in a favorable light, party A will have difficulty accurately assessing the ethicality of party B's behavior. Across most major scandals of the last decade, many people—members of boards of directors, auditing firms, rating agencies, and so on—had access to the appropriate data and should have noticed and acted on the unethical behavior of others. Yet they did not do so, at least in part because of the

psychological tendency not to notice bad data that we would prefer not to see.

One striking aspect of the story of the credit-rating agencies is how closely it resembles the story of auditing firms that emerged about seven years earlier. The most prominent scandal in the early part of the new millennium was the fall of Enron, the most famous business collapse of our time. How did Arthur Andersen, Enron's auditor, vouch for the firm's financial health during the time that Enron was concealing billions of dollars in debt from its shareholders? Quite simply, Arthur Andersen had ample reason to be afflicted by motivated blindness. In 2001, Andersen earned millions from Enron, then its second-largest client: \$25 million in auditing fees and \$27 million in consulting fees. Andersen had a strong motivation to retain and build on these lucrative contracts. Obviously, finding problems with your auditing client's books is no way to keep it as an ongoing client. In addition, it is likely that many Andersen auditors hoped to be hired by Enron, as a number of their colleagues had been.

Enron's collapse was not unique. Soon after the company's fall, major financial scandals unfolded at other major corporations, including World-Com, Global Crossing, Tyco International, and Parmalat. In each case, auditors were implicated for failing to bring wrongdoing to light. These scandals may not have occurred if members of these firms had taken note of the unethical behavior of their colleagues and clients rather than overlooking it. These cases shed light on a weakness of the U.S. auditing system: it allows motivated blindness to thrive.

Max and his colleagues tested the strength of such conflicts of interest by giving study participants information about the potential sale of a fictional company. The participants' task was to estimate the company's value.⁴ Participants were assigned to one of four roles: buyer, seller, buyer's auditor, or seller's auditor. All participants read the same information, including information that could help them estimate the worth of the firm. Those acting as auditors provided estimated valuations of the company's worth to their clients. As the literature on self-serving biases discussed earlier in the book would suggest, sellers submitted higher estimates of the company's worth than did prospective buyers.⁵ More relevant to this chapter, the auditors, who were advising either the buyer or the seller, were strongly biased toward the interests of their clients: sellers' auditors publicly concluded that the firm was worth far more than did buyers' auditors.

Were the auditors' judgments intentionally biased, or was bounded ethicality at play? To answer this question, the auditors were asked to estimate the company's true value, as assessed by impartial experts, and were told they would be rewarded for the accuracy of their private judgments. Auditors for the sellers reached estimates of the company's value that, on average, were 30 percent higher than the estimates of auditors who served buyers. This evidence shows that, rather than making a conscious decision to favor their clients, the participants assimilated information about the target company in a biased way. Being in the role of the auditor biased their estimates and limited their ability to notice the bias in their clients' behavior. Thus, even a purely hypothetical relationship between an auditor and a client distorted the judgments of those playing the role of auditor. Furthermore, we replicated this study with actual auditors from one of the "Final Four" large auditing firms as our participants and received similar

results. Undoubtedly, a long-standing relationship involving millions of dollars in ongoing revenues would have an even stronger effect.

When a client behaves unethically, its auditor doesn't see this unethical behavior for the same reason the client doesn't see its own unethical behavior. Bias in the direction of those who pay their bills (their clients) prevents auditors from distancing themselves from their clients. From the perspective of behavioral ethics, auditors become more like their clients than they would be if no such motivation existed; as a result, they are unlikely to see the unethical actions and biases in their clients' behavior. The client's bounded ethicality transfers to the auditor.

Motivated blindness appears to be responsible for the failure to notice others' unethical behavior in many domains. Consider the widespread use of steroids in baseball. In 2007, Barry Bonds of the San Francisco Giants surpassed Hank Aaron to become the all-time leader in career home runs, perhaps the most valued record in Major League Baseball. Law enforcement agencies, the baseball commissioner, and fans now question whether Bonds's performance truly surpassed that of Aaron. Many believe that Bonds used steroids or other drugs to improve his performance, especially given that his longtime trainer was indicted for supplying steroids to athletes. Similar suspicions have swirled around other MLB superstars, including Sammy Sosa, Roger Clemens, David Ortiz, Manny Ramirez, and others. In July 2009 it surfaced that MLB had known of at least 100 players who had tested positive for using performance-enhancing drugs.

In light of the steroid scandal, baseball fans tend to direct their wrath at the players who cheated (and got caught) for tainting the sport. Yet the nature of competition in Major League Baseball, the related financial rewards, and lax enforcement of drug rules were all contributing factors that gave players a strong incentive to use steroids. In fact, many players may have felt they would have been at an unfair disadvantage if they *didn't* use steroids. Fingers should also be pointed at the MLB commissioner, the San Francisco Giants team, and the players' union. None of these groups investigated the rapid changes in Bonds's and other players' physical appearances, their enhanced strength, and their increased power at the plate as these changes occurred. Given that sports journalists and many fans understood that a massive steroid problem existed throughout MLB, why didn't the commissioner, individual teams, or the players' union address the problem? The answer, we believe, lies in the fact that these groups benefited financially, at least in the short term, from the steroid use of players such as Bonds. Steroid use led to home runs, home runs increased attendance, and increased attendance generated more profit for the league, the teams, and the players. These benefits prevented MLB management from noticing problems it preferred not to see.

Was steroid use that easy to notice? Take a look for yourself. In [figure 6](#) we plot the number of home runs hit by the players with the first-, second-, and third-most home runs each year from 1990 to 2009. The peaks between 1998 and 2001, typically recognized as the height of the steroid era in baseball, should have provided reasonably good evidence for the MLB to act (along with the other evidence available). To rule out the possibility that a few stellar players during this era skewed the results, we averaged the number of home runs hit by the home run

leader from 1991 to 1994. This average was forty-four. We then counted the number of players in each year of the 1998–2001 steroid era who hit that number of home runs or more. Ten players in 1998, eight players in 1999, six players in 2000, and nine players in 2001 matched or beat the average number of home runs hit by the home run leaders between 1991 and 1994. This simple arithmetic suggests that an extraordinary number of players were hitting balls out of the park during the steroid era—and that noticing these unusual statistics shouldn't have been difficult.

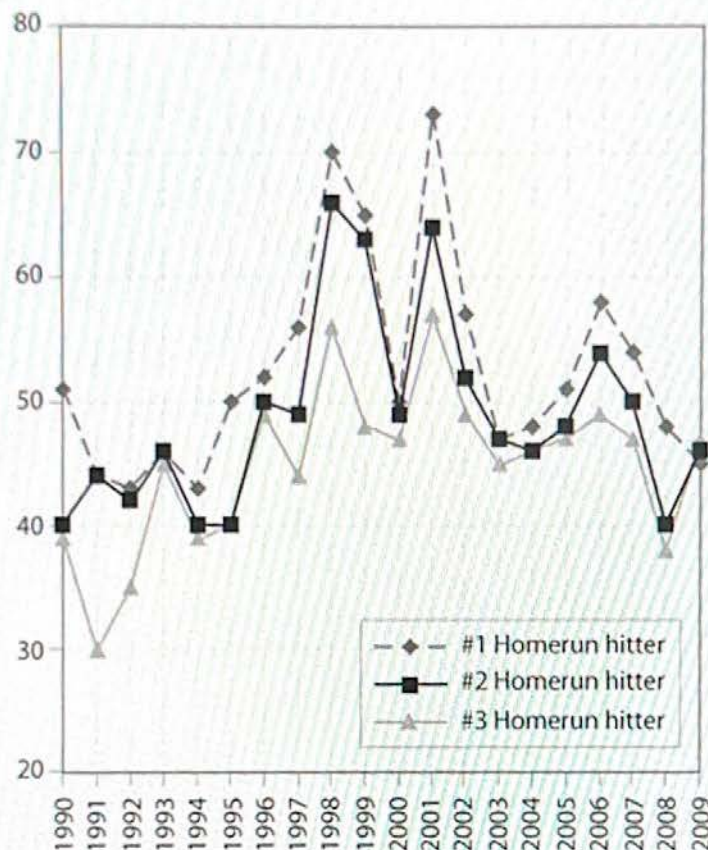


Figure 6. Home runs hit by the top three MLB players, 1990–2009

Motivated blindness can cause people at the highest levels of society to engage in behaviors that they would never condone with greater awareness. Consider the child sexual abuse scandals that have rocked the Catholic Church. How did the abuse run rampant for decades without being stopped by the church's hierarchy? To take one striking example, Cardinal Bernard F. Law, the archbishop of Boston, failed to act on the enormous amount of child abuse that occurred under his jurisdiction. He admitted in court papers that he knew about accusations against John J. Geoghan, later convicted as a child molester, yet Law returned the priest to parish work. Law also admitted that he kept James Foley active in his ministry despite learning in 1993 that the priest had fathered two children with a woman in his parish and, in 1973, had fled the scene when she took a lethal dose of pills in an apparent suicide attempt. Law kept many other criminals and church rule-breakers active in the priesthood.⁶

Making the question more complex, Cardinal Law, a former civil rights activist, had dedicated his life to helping others. All the evidence suggests that Law was an ethical person who made some very highly unethical and probably illegal decisions in his executive role. Why did he tolerate illegal, abusive behavior? Law testified that, in retrospect, he relied on outdated medical and psychiatric advice regarding the ability of the abusers to curtail their behavior when deciding whether to keep them in the church. It is quite possible that Cardinal Law believed that priests such as Geoghan would be able to control their behavior. It is also possible that Cardinal Law's desire for abusers to be reformed blinded him to obvious evidence that the immoral and criminal behavior was likely to be repeated.

More recently, Cardinal Joseph Ratzinger, the current pope, has been accused of cover-ups of other sex abuse scandals within the Catholic Church, including last-minute transfers of accused priests to other parishes and emphasizing loyalty to the church over truly ethical responsible behavior. Without excusing any behavior that led to the abuse of children, we believe that it is possible that the pope's loyalty to his organization may have blinded him to the seriousness of his actions. Rather than a defense of unethical behavior, motivated blindness offers a psychological explanation of how unethical behavior may come about.

As shown in these examples and many others, we are not only blind to our own unethical actions but also to the unethicality of those around us. The motivation to remain blind to the unethical behavior of others comes at us in many forms, including fear, incentives, organizational loyalty, and organizational culture. To behave more ethically, we need to remove our blinders and examine the effects of these forces on our judgment.

Indirect Blindness

Imagine that your company produces a slow-selling item. It has few customers, but the ones who like the item would pay much more for it than you currently charge. Imagine that these customers are hostage to your pricing increases because you have a monopoly on the product and because they need it to stay healthy. You are aware that any significant increase in the product's price would attract negative publicity that would cost you more than you would gain from the price increase. What would you do to solve this puzzle?

In August 2005, Merck, a major pharmaceutical firm, found an answer. Merck sold off a slow-selling but effective cancer drug named Mustargen to Ovation, a smaller pharmaceutical firm, along with a second cancer drug called Cosmegen.⁷ A chemotherapy drug used to treat lymphoma, Mustargen was used by fewer than 5,000 patients and generated annual sales of only about \$1 million for Merck at the time it was sold.

At first glance, it looks as if Merck had found an effective means of moving a slow-moving drug out its busy manufacturing system. But it turns out that manufacturing Mustargen was not the issue for Merck. After selling the rights to Mustargen and Cosmegen to Ovation, Merck continued to manufacture the drugs for Ovation on a contract basis.

If making a small amount of a product was inefficient, why would Merck continue to

produce Mustargen? Consider what happened after Merck completed its deal with Ovation: Ovation increased the price of Mustargen by approximately tenfold and raised the price of Cosmegen by even more. It turns out that Ovation often buys small-market drugs from visible pharmaceutical firms that have public-relations problems associated with dramatically increasing the price of drugs needed by their consumers. In a different transaction, Ovation purchased the drug Panhematin from Abbott Laboratories and increased the price nearly tenfold; Abbott continued to manufacture the drug. Merck's decision to sell Mustargen and Cosmegen to Ovation suggests that its leaders hoped to see headlines such as "Merck Sells Two Drugs to Ovation" rather than headlines such as "Merck Gouges Cancer Patients, Increases Cancer Drug Prices by 1,000 percent."

How did Merck get away with this clever strategy? Merck succeeded because human intuition does not sufficiently hold people and organizations accountable for such indirect unethical behavior. Even when data suggesting unethical intent is obvious, we still let those who behaved unethically off the hook. Notice that we are not commenting on the ethicality of increasing prices for needed cancer medication. In fact, we generally believe that high profits in the pharmaceutical sector have helped to create the vast array of amazing drugs that are available to patients today. But it is important to become aware of how difficult it is to "see" the indirect but unethical actions of others; this awareness can help us identify individuals and organizations that intentionally create opaqueness. If Merck did indeed assume that a tenfold price increase in a cancer drug would attract negative attention, we believe most people would view the decision to hide the increase through an intermediary such as Ovation as a manipulative, unethical strategy. Merck's apparent strategy often works, as do many other similar strategies, because the public and the press too often fail to notice the dirty work that individuals and organizations perform through intermediaries. Most of us fail to hold others sufficiently accountable for their indirect unethical actions.

This argument was tested more precisely by Max and his colleagues in an experimental study designed to mirror the environment of the Merck story.⁸ Participants in the study read the following passage:

A major pharmaceutical company, X, had a cancer drug that was minimally profitable. The fixed costs were high and the market was limited. But, the patients who used the drug really needed it. The pharmaceutical was making the drug for \$2.50/pill (all costs included), and was only selling it for \$3/pill.

The participants were then divided into two groups. Members of one of the groups were asked to assess the ethicality of the following action:

A. The major pharmaceutical firm raised the price of the drug from \$3 per pill to \$9 per pill.

The other group was asked to assess the ethicality of a different course of action:

B. The major pharmaceutical X sold the rights to a smaller pharmaceutical. In order to

recoup costs, company Y increased the price of the drug to \$15 per pill.

As we expected, people who read action A judged the behavior of the pharmaceutical firm more harshly than did participants who read action B, despite the fact that action A would have had a smaller financial impact on patients.

It is useful to note that these participants responded to only one of the two options, not to both (what experimental researchers call a “betweensubjects design”). We then went on to present a third group of participants with both possible actions and asked them to judge which action was more unethical. Now the preferences were reversed: When they could compare the two scenarios, participants viewed action B as being more ethically problematic than action A. This finding is consistent with substantial research showing that this type of “side by side” or “joint” evaluation leads to more reflective and rational assessments than “separate” (one at a time) evaluations. Yet it is important to recognize that most real-world, morally questionable actions come to us one action at a time.

We replicated this result in domains other than drugs, such as contaminated land and pollution controls. We consistently found that when study participants were judging one option, they significantly discounted the unethicity of the focal firm acting through an intermediary. Yet when they were asked to compare an indirect action to a direct action, they saw through the indirectness and made their assessments based on the magnitude of the harm created by the action.⁹ Further, we improved the transparency of the intent of the pharmaceutical firm in the indirect condition by making it clear that the firm understood the implications of selling off the drug and would profit by doing so. Even with extraordinary transparency, participants viewed indirect action, under separate evaluation, to be less unethical than direct action.

Finally, an economist, Luke Coffman, turned our question into an experimental game in which the goal was to find out how much other actors would punish a party for acting unethically either directly versus indirectly.¹⁰ Luke created what he calls a “four-player dictator game.” In the more common two-person dictator game, player A is given a fixed amount of money and faces a choice between giving none, some, or all of this money to player C. Player C is a passive recipient of player A’s decision. In Luke’s game, player A, who had \$24, was given the option of playing the dictator game, as the dictator, with player C, or he could sell the rights to the game to player B at a negotiated price. If player B bought the game from player A, player B assumed the role of the dictator and had \$24 to allocate in a dictator game played with player C. The final step in the game is that player D had the opportunity to punish player A (but not player B) for his actions by reducing player A’s final payoff. As expected, when player A stayed in the game (i.e., did not sell the rights to the game to player B), player D typically punished player A for giving smaller amounts to player C, with the size of the punishment directly related to the amount of money player A kept for himself. More interestingly, and consistent with our studies, when player A did sell the rights to the game to player B, rather than choosing to be a greedy dictator, player D decreased the amount of punishment dramatically. That is, participants punished those who engaged in direct unethical behavior more than they punished those who engaged in indirect unethical behavior. This

difference held up even when the net harm to player C was the same, and in later versions, when player A could fully predict how player B's decisions would affect player C.

This type of behavioral ethics research demonstrates that by engaging in indirect action under predictable circumstances, decision makers trigger *indirect blindness* in the eyes of observers and thus are let off the hook for the harms they cause. Members of organizations routinely delegate unethical behavior to others in their organizations. Managers tell their subordinates to "do whatever it takes" to achieve production or sales goals, for example, leaving open the possibility of aggressive or even unethical tactics. U.S. companies outsource production to offshore subcontractors that are inexpensive because they are less constrained by costly labor and environmental ethical standards. Partners at accounting firms remind junior auditors about the importance of retaining a client that has inappropriate accounting practices. Across many other situations, people overlook the problematic ethical implications of others' behavior when the actions occur indirectly.

Here's another example. Max lives in Boston, is a longtime football fan, and doesn't like cheaters. Thus, Max was quite disappointed by behavior that occurred during the 2007 National Football League's season. The New England Patriots in that year were arguably one of the greatest football teams of all time. Unfortunately, Bill Belichick, the team's highly visible head coach, threatened the team's reputation by blatantly cheating. When the Patriots were playing the New York Jets (a weak team) early in the 2007 season, Belichick directed an assistant to film the Jets' private defensive signals—a clear violation of the rules, as Belichick well knew.¹¹ NFL commissioner Roger Goodell fined Belichick \$500,000, fined the Patriots \$250,000, and penalized the Patriots in the form of taking away one of their future high-value draft choices.

Clearly, Belichick was guilty. But what about the Kraft family, which owns the Patriots? They hired Belichick, encouraged him to win, and offered no criticism of the coach after the incident. The ethics of the Kraft family were largely unquestioned by the media, and Patriots fans did not seem overly concerned about the reputable family's behavior. The Kraft family's notable silence on the issue was indirectly unethical and, as a result, went unnoticed.

When people stand by the unethical actions of their subordinates, they own that unethical action. Their silence suggests that their only problem with the unethical action is that it was detected. We should hold executives accountable for the actions of their employees when all evidence suggests that the organization tolerated unethical behavior. Unfortunately, behavioral ethics research has provided abundant evidence that outsiders overlook the unethical actions of actors who work through indirect parties.

Unethical Behavior on a Slippery Slope

According to an interesting folk tale, if you place a frog in a pot of hot water, the frog will jump out. However, if you put the frog in a pot of warm water and raise the temperature gradually, the frog will not react to the gradual change in temperature, and it will cook to death.

While the story happens to be untrue, it is a fine analogy for the failure of most people to notice the gradual erosion of ethical standards. As we suggested in the previous chapter, our unethical behavior often occurs on a slippery slope. We excuse ourselves for committing one tiny infraction and then allow ourselves to commit increasingly unethical infractions as time passes.

Behavioral ethics research shows that people also commonly fail to notice the slippery slope of *others'* unethical behavior. In addition to Bernard Madoff's feeder funds, the broader professional investment community and the U.S. Securities and Exchange Commission (SEC) did not notice that his funds' performance was inconceivable. Why not? Part of the story is motivated blindness. Another part, though, is that this fraud developed slowly, over at least a fifteen-year period. When fraud occurs on a slippery slope, the impossibility of returns such as Madoff's is likely to go unnoticed.

In fact, people are capable of ignoring clear warning signals of others' unethical behavior. Beginning in 1999, independent financial fraud investigator Harry Markopolos repeatedly attempted to warn the SEC that Madoff's returns were not legally possible. Yet all indications suggest that the SEC did not take these accurate warnings seriously. As a result, Madoff's fraud involving more than \$50 billion didn't come to light until the mega-swindler himself confessed.

Now imagine that an accountant with a large auditing firm (perhaps Arthur Andersen) is in charge of the audit of a large company (perhaps Enron) with a strong reputation. For three years in a row, the client's financial statements are extremely ethical and of high quality. As a result, the auditor approves the statements and has an excellent relationship with its client. The next year, however, the company commits clear transgressions in its financial statements, stretching and even breaking the law in certain areas.

Now imagine a different scenario. This time, the auditor notices that the corporation stretches the law in a few areas the first year but does not appear to break the law. The next year, the firm is even more unethical, committing a minor violation of federal accounting standards. The third year, the violations are a bit more severe. In the fourth year, the auditing firm finds itself facing the type of severe violations described in the previous paragraph, in which the client crosses the ethical line abruptly.

How is the corporation's auditor likely to react in each of these two scenarios? In the first situation, the auditor probably would refuse to certify that the financial statements were acceptable according to government regulations. In the second scenario, it is far less likely that the auditor would notice the same severe ethical transgression. In other words, auditors would be more likely to notice and refuse to sign the statements in the first version of the story than in the second one, even if the unethical behavior was the same in the last year described in both stories.¹²

In our research with David Messick and Francesca Gino, we explored whether this "slippery slope" pattern of behavior can explain the common failure to notice the egregious behavior of others.¹³ Using laboratory studies with features similar to those described in these two stories, we found that people are less likely to perceive changes in others' unethical behavior if the changes occur slowly over time rather than abruptly.¹⁴

Visual perception research, such as the basketball-passing video that we described at the start of the chapter, demonstrates that we frequently fail to notice changes that occur right in front of our eyes.¹⁵ In one study investigating “change blindness,” an experimenter holding a basketball stopped pedestrians to ask for directions. As each pedestrian gave directions, a group of confederates (research assistants) walked between the experimenter and the pedestrian. As the confederates walked by, the experimenter handed the basketball to one of the confederates. After the pedestrian completed giving directions, the experimenter asked her if she had noticed any sort of change while she was talking. Most of the individuals in the study did not notice any change. Yet when they were asked directly if they had seen a basketball, many recalled seeing the basketball, and many could even recount specific characteristics of the ball. Thus, while the pedestrians failed to notice explicitly that a small, incremental (but obvious) change took place, it was possible that they could have done so had they been attuned to it. Similarly, many Madoff investors can now recount the evidence that they should have perceived long before he confessed to his crimes.

The scientific study of change blindness focuses on visual perception. People also fail to notice other types of changes in their environment that lead to significant decision-making errors with ethically relevant consequences. As a result, people are less likely to notice others’ unethical behavior when it occurs in small increments—on a slippery slope—than when it occurs suddenly, a phenomenon that should put us on alert to slowly degrading ethical behavior.

Valuing Outcomes over Processes

Consider story A:

A pharmaceutical researcher defines a clear protocol for determining whether or not to include clinical patients as data points in a study. He is running short of time to collect sufficient data points for his study within an important budgetary cycle in his firm. As the deadline approaches, he notices that four subjects were withdrawn from the analysis because of technicalities. He believes that the data derived from those four subjects in fact are appropriate to use, and when he adds those data points, the results move from not quite statistically significant to significant. He adds these data points, and soon the drug goes to market. This drug is later withdrawn from the market after it kills six patients and injures hundreds of others. How unethical do you view the researcher to be?

Now consider a somewhat different story, B:

A pharmaceutical researcher defines a clear protocol for determining whether or not to include clinical patients as data points in a study. He is running short of time to collect sufficient data points for his study within an important budgetary cycle in his firm. He believes that the product is safe and effective. As the deadline approaches, he notices that

if he had four more data points for how subjects are likely to behave, the analysis would be significant. He makes up these data points, and soon the drug goes to market. This drug is a profitable and effective drug, and years later shows no significant side effects.

How unethical do you view this researcher in this second story to be? Which story do you find to be more egregious?

While we have shown you both stories, in a study based on them, we presented one story to one group of participants and the other story to another group of participants, such that each group only saw one story.¹⁶ Those who read story A were more critical of the researcher in the story than were those who read story B. Those who read story A also reported that the behavior in the first story should be punished more harshly. Yet, as you probably noticed, the researcher's behavior was more unethical in story B than in story A.

Why would people view the behavior in story A as more egregious than the behavior in story B? The outcome bias provides an answer.¹⁷ The outcome bias describes the tendency to take results into account, in a manner that is not logically justified, when evaluating the quality of the decision process that a decision maker used. Decision researchers Jon Baron and Jack Hershey were the first to find that in contexts ranging from simple laboratory gambles to medical decision-making, people judge the wisdom of decision makers based on the outcomes they obtain.

Our own research in behavioral ethics finds that people too often judge the ethicality of actions based on whether harm follows, rather than on the ethicality of the choice itself.¹⁸ As in the research on direct versus indirect effects described earlier, people are affected by this bias when they confront one story or instance at a time. Clearly, the ability to see two versions of a story that have transparent differences allows us to avoid the outcome bias and to pay attention to and compare the actions of the two researchers. When this is done experimentally, people rate story B as more egregious than story A. But, as noted earlier, most of the time, the world presents us with one situation to assess at a time. Philosophers have long debated whether we should judge ethical actions based on the rules used to decide which action should be taken or based on the outcome itself. We expect this age-old debate will continue to be fiercely argued. However, when we judge an action based on its outcome and don't consider alternative options or scenarios (as is often the case), this judgment does not reflect the actor's underlying intentions.

The outcome bias is solidly integrated into our laws. Consider the story, told by psychologist Fiery Cushman and his colleagues, of two brothers, Jon and Mark, both of whom lack a criminal record or good marksmanship but possess a quick temper.¹⁹ Imagine that a man confronts the two brothers and insults their family. Vowing to kill the guy, Jon pulls out a gun, but he misses his shot, and his target remains unharmed. By contrast, Matt decides he only wants to scare the man. He pulls out a gun, accidentally shoots the guy in the heart, and kills him. Cushman and colleagues note that in most U.S. states, Matt can expect a far longer prison sentence than Jon. In other words, the law pays more attention to outcomes than to intentions.

Cushman and colleagues have offered a brilliant experiment related to this hypothetical legal

story and to the outcome bias. Simplifying the essence of their experiment, imagine that you face a choice between the following two options. You will be playing the game you choose with an unknown other person, also a participant in the experiment.

Option A: You roll a six-sided die. If it comes up a one, two, three, or four, you get \$10 and the other party gets \$0. If it comes up a five, you get \$5 and the other party gets \$5. If it comes up a six, you get \$0, and the other party gets \$10.

Option B: You roll a six-sided die. If it comes up a one, you get \$10 and the other party gets \$0. If it comes up a two, three, four, or five, you get \$5 and the other party gets \$5. If it comes up a six, you get \$0 and the other party gets \$10.

Notice that option A is the greedy choice, as it offers you more opportunities (four out of six, to be exact) to claim \$10 for yourself. By contrast, option B is the fair choice, at least most of the time, as it offers four opportunities for the \$10 to be split evenly between you and the other party. Regardless of which choice you make, any of the three outcomes described is possible; it's just their probabilities that differ.

After you choose which game to play, the die is rolled and the money is paid. Cushman and colleagues then allow the other party to punish you, the chooser, by reducing your payment without incurring any cost herself. The fascinating result is that when allocating punishment, the other party typically pays more attention to the equality of the result of the rolled die—a random outcome—than to the chooser's sense of fairness (as demonstrated by which option she chose). For example, if you chose to be fair and play option B, and then rolled a one, the other party is more likely to punish you than she would if you had greedily chosen option A and rolled a five.

These results clarify our unfortunate tendency to blame people too harshly for making sensible decisions that have unlucky outcomes. Compounding the problem, judging decisions based on their outcomes means that we often wait too long to condemn unethical behavior—until after a bad outcome has occurred. Many people now question the ethics of the Bush administration's decision to invade Iraq in 2003, including its misrepresentation of the "facts" that prompted the war. But criticism of the invasion was limited in much of the United States when it seemed as if the war was going well. When the war began to drag on, many more people began to question the Bush administration's prewar tactics, such as unfounded claims of evidence of weapons of mass destruction in Iraq. The outcome bias may partially explain why so many reserved judgment on the decision to go to war until they knew what the outcome would be. We often fail to take notice of unethical behavior—and condemn it only after a harmful outcome occurs.

We now return to the case of auditors at another level of analysis. For decades, U.S. auditing firms provided both auditing and consulting services to their clients. As we noted earlier, this situation logically and psychologically compromised the independence of their audits.²⁰ Long before Enron's collapse, we had ample evidence that the existing structure compromised the ethics of the auditing profession.²¹ Despite added evidence of the failure of auditor independence and the widespread belief that independence was essential for reliable audits, it

took the glaringly obvious failures of Enron, WorldCom, Tyco, and other firms to persuade the U.S. government to address the underlying conflicts of interest that compromised auditors.²² Only these very bad outcomes motivated our legislative representatives to address the problem. But, for reasons that we will explore in [Chapter 7](#), even these changes were insufficient and poorly crafted to solve the core problem.

The outcome bias is related to research on identifiable victims.²³ The “identifiable victim effect” refers to the finding that people tend to be far more concerned with and show more sympathy for identifiable victims than statistical victims. Identifiable victims are specific people, while statistical victims are unknown, unspecified people. People tend to feel more concern for specific victims, even when no useful personalizing information about the victim is available (e.g., only a name is provided).²⁴ Now consider that the same unethical action could harm an identifiable victim, an unidentifiable victim, or no victim at all. Just as we often fail to notice unethical behavior when no victims have yet been affected by it, we are less likely to see the presence of unethical behavior when statistical victims are affected than when the victims are identifiable. Once again, differences in judgments of ethical behavior depend on the outcome of the unethical action, including our perceptions of who was affected, even though the perpetrator’s actions remain the same.

The story of Noreen Harrington, a Goldman Sachs veteran who was the whistleblower in the mutual fund late-trading scandal, illustrates how depersonalizing the victims of our unethical behavior allows such behavior to be perpetuated.²⁵ The scandals involved two questionable practices: late trading, or the illegal practice of buying and selling funds after the 4:00 p.m. market close but still receiving the 4:00 p.m. price; and market timing, which involves exploiting prices via time zone differences in international funds, a practice that is legal but can be in violation of fund rules, as it often profits “market timers” at the expense of long-term shareholders. Harrington has said that prior to blowing the whistle on these practices, she viewed them as part of “a nameless, faceless business . . . in this business this is how you look at it. You don’t look at it with a face.”²⁶ That view changed, she said, when her older sister asked her for advice on her 401(k) account. Her sister, whom Harrington characterized as one of the hardest workers she knew, was worried that the losses she saw in her retirement account would prevent her from retiring. Suddenly, Harrington “thought about this from a different vantage point,” she explains. “I saw one face—my sister’s face—and then I saw the faces of everyone whose only asset was a 401(k). At that point I felt the need to try and make the regulators look into [these] abuses.”²⁷

Our own industry—higher education—is not immune from this bias. In our discussion of in-group favoritism in [Chapter 3](#), we discussed the widespread policy of universities admitting the underqualified children of alumni. To our surprise, few commentators have publicly objected to the policy of admitting such underperforming “legacies.” The lack of outrage over this ethically questionable practice is likely due in part to the difficulty of identifying the victims of such practices—that is, those who are denied admission. Because the victims of legacy admissions policies are statistical rather than identifiable, people fail to perceive that these practices cause harm, and the behavior of those responsible goes unchecked. Even when

we do recognize the negative outcome of such policies in theory, we are often dulled by their lack of vividness when we do not know who was actually harmed.

Behavioral ethics research supports the argument that most people want to act ethically. Yet we still find ourselves engaging in unethical behavior because of biases that influence our decisions—biases of which we may not be fully aware. As we have noted in this chapter, these biases affect not only our own behavior, but also our ability to see the unethical behavior of others. Having completed our overview of the systematic mistakes the human mind makes in ethical domains, in the next three chapters, we will use this knowledge to explore implications for organizations and society, as well as opportunities to change these dysfunctional patterns of behavior.