Organizational Behavior and Human Decision Processes 136 (2016) 95–105

Contents lists available at ScienceDirect

Organizational Behavior and Human Decision Processes

journal homepage: www.elsevier.com/locate/obhdp

Bounded awareness: Implications for ethical decision making

Max H. Bazerman [a,*], Ovul Sezer [b]

[a] Harvard University, Baker Library 453, Boston, MA 02163, United States
[b] Harvard Business School, Baker Library – 444B, Boston, MA 02163, United States

ARTICLE INFO

Article history:
Received 22 June 2015
Revised 2 November 2015
Accepted 8 November 2015

Keywords:
Ethics
Decision making
Bounded awareness

ABSTRACT

In many of the business scandals of the new millennium, the perpetrators were surrounded by people who could have recognized the misbehavior, yet failed to notice it. To explain such inaction, management scholars have been developing the area of behavioral ethics and the more specific topic of bounded ethicality—the systematic and predictable ways in which even good people engage in unethical conduct without their own awareness. In this paper, we review research on both bounded ethicality and bounded awareness, and connect the two areas to highlight the challenges of encouraging managers and leaders to notice and act to stop unethical conduct. We close with directions for future research and suggest that noticing unethical behavior should be considered a critical leadership skill.

Imagine that you are an investment advisor. You have a client who has a long-term perspective and a moderate tolerance for risk. You are considering one of four feeder funds to recommend to your client: the Tobacco Trade Investments Fund, the Alpha Investments Fund, the Fortitude Investments Fund, or the Power Trade Investments Fund. Fig. 1 shows the returns for each of the funds over the last four years, as well as the return for the S&P 500. Which fund do you recommend?

Most study participants, including MBA students and executives, opt for Fortitude (Zhang, Fletcher, Gino, & Bazerman, 2015). Yet, when people are asked, instead, whether they see a problem with any of the funds, most participants identify that Fortitude's return is impossible. Quite simply, no fund can outperform the market consistently with so little volatility. Asking participants to focus on whether a problem exists brings this to light immediately, but without that request, they instead concentrate on Fortitude's high returns and low volatility. Essentially, asking participants which fund to recommend leads to ethical fading, or the tendency of the ethical elements of a decision to fade from our attention when it is diverted elsewhere (Tenbrunsel & Messick, 1999).

In fact, if you picked a fund with the same return pattern as the hypothetical Fortitude fund, your feeder fund went bankrupt, as shown in Fig. 2. As it turns out, that fund was invested with Bernie Madoff.

Bernie Madoff's Ponzi scheme led to enormous losses—as much as $65 billion—for his investors. Madoff sold most of his investments through feeder funds—that is, funds that sold investments to their retail customers, and invested all or most of these funds with Madoff. The feeder funds typically earned a small percentage (1–2 percent) of the money invested plus 20 percent of any investment profits earned (common in the hedge fund world). Madoff reported impressive profits, and the feeder funds performed quite well in the process. While Madoff's reports could not have accurately portrayed the fund's performance, many people with vast financial expertise didn't notice the impossibility of Madoff's returns.

We argue that those individuals who recommended Madoff to their investors, even though they had the knowledge and intelligence to know that dramatically outperforming the market over nine years with very little volatility was impossible, ignored their moral obligation. It is true many of these individuals did not consciously recognize what they knew to be the case: that Madoff's reports were impossible. Thus, they did not recognize that their failure to notice, and the similar failure of many of their peers, could be the keystone in the perpetuation of a fraud that would cheat many people of their life savings. But however powerful the psychological forces underpinning their choice to ignore what they knew, we argue that it was a choice. When an individual has the knowledge and the intelligence to notice unethical conduct and the power to act on it, they have a moral obligation to do so. The consequences of the failure of such actors in the case of Madoff's fraud as well as scandals such as those involving Volkswagen, Enron, and Worldcom documents this obligation and shows how they bear some of the responsibility; Madoff could not have

* Corresponding author.
E-mail addresses: mbazerman@hbs.edu (M.H. Bazerman), osezer@hbs.edu (O. Sezer).
URLs: http://www.people.hbs.edu/mbazerman (M.H. Bazerman), http://www.hbs.edu/faculty/Pages/profile.aspx?facId=508167 (O. Sezer).

http://dx.doi.org/10.1016/j.obhdp.2015.11.004
0749-5978/

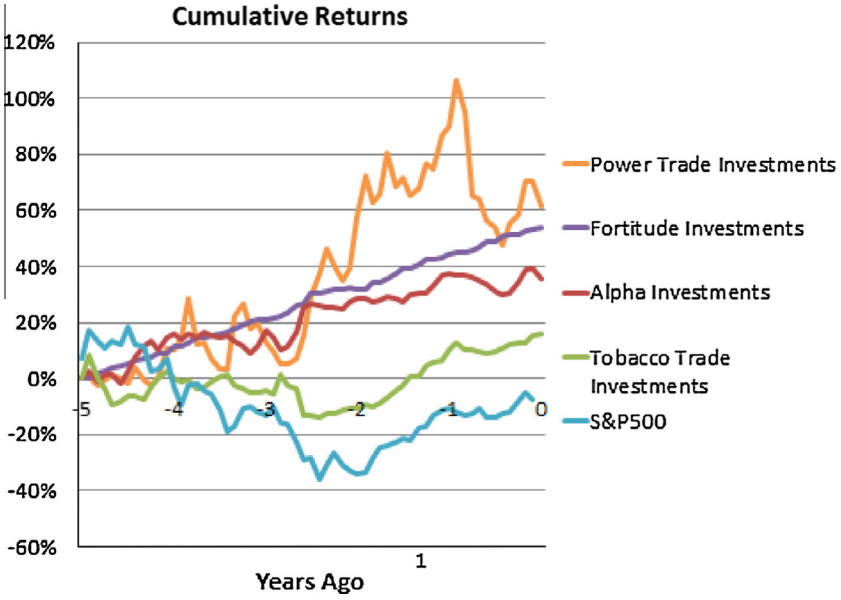

Fig. 1. The cumulative returns of four feeder funds.

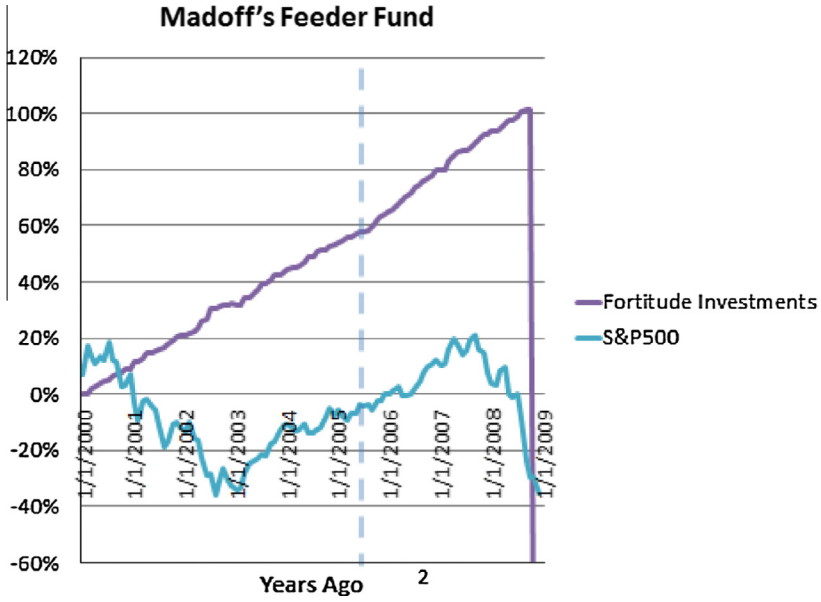

Fig. 2. Bernard Madoff's feeder fund.

perpetuated his fraud without the contributions of financial experts who knew better.

Two psychological processes underlie experts' failure to notice the implausibility of Madoff's returns. First, their awareness of the developing problem was bounded. Second, they engaged in unethical action without intending to do anything wrong and without knowing they were doing so. We will review both the literature on bounded awareness and bounded ethicality to argue that when bounded awareness prevents actors from noticing the unethical actions of others, this failure to notice becomes a form of unethical action – one that affects watchdogs and leaders on a regular basis. Thus, while bounded awareness focuses on a cognitive failure, and bounded ethicality focuses on an ethical failure, we argue that the combination, the failure to notice and act on unethical action around us, constitutes a further ethical failure. We will provide overviews of the psychology of bounded awareness and the literature on behavioral ethics. We will then attempt to integrate the two and conclude with our views about where this research should head in the future.

1. Bounded awareness

One of the core concepts of organizational behavior is bounded rationality (March & Simon, 1958; Simon, 1957). Bounded rationality argues that "human rationality is very limited, very much bounded by the situation and by human computational powers" (Simon, page 34, 1983). This broad concept has come to highlight the difference between traditional economists' and psychologists' views of human decision making. Bounded rationality most commonly manifests in systematic and predictable mistakes, which humans make on a regular basis (Kahneman & Tversky, 1973, Kahneman & Tversky, 1979; Bazerman & Moore, 2013).

The literature on these predictable biases has had a powerful impact on behavioral economics, including the development of "nudge units" that help governments to apply these insights to policymaking around the globe (Thaler & Sunstein, 2009). From tax collection to education to charity giving, several governments across the world have started to implement interventions that address these predictable biases to design environments that lead individuals to more optimal and societally beneficial decisions.

Despite the power of the decision-bias perspective, bounded rationality research has been limited to exploring how people misuse and mis-integrate the information that is part of their cognitive set. Chugh and Bazerman (2007), in contrast, focus on a different aspect of the bounds to our rationality. Specifically, they argue that people have "bounded awareness" that prevents them from focusing on easily observable and relevant data. Bounded awareness leads people to not "see" important, accessible, and perceivable information during the decision-making process. Chugh and Bazerman (2007) describe a "focusing failure" that occurs when the information needed and available to make a good decision does not align with the information a decision maker considers.

Neisser (1979) identified a form of bounded awareness when he asked a sample of undergraduates to observe a video of two visually superimposed groups of players passing basketballs. One group of players wore white shirts and the other group wore dark shirts. Researchers asked participants to count the number of passes between members of one of the two groups. The visual overlap between the players and the grainy quality of the video makes the task moderately difficult. They were so focused on the task that, afterwards, only 21% of participants reported seeing a woman who walked through the basketball court carrying an open umbrella (see https://www.youtube.com/watch?v=wcjnJ1B7N0E). On their second viewing, without being distracted by the task of counting passes, every participant could clearly see the woman with the umbrella. By focusing on a particular task, they missed obvious information in their visual world. This demonstration has been well replicated (Simons, 2000; Simons & Chabris, 1999; Simons & Levin, 2003). For instance, Simons and Chabris (1999) created a more contemporary video in which a person in a gorilla costume walks through a basketball game, thumping her chest, and is clearly and comically visible for more than five seconds. Again, most people fail to notice her.

Neisser described this effect as "inattentional blindness." While we are inspired by the work on inattentional blindness, we follow Chugh and Bazerman (2007) in using the term bounded awareness to capture the argument that much of the information that we miss is not visual. That is, people fail to notice the crucial information they need to make a good decision (or fail to notice when it's missing and needs to be obtained), instead relying on the data in front of them. In addition, people fail to notice that there are other options to consider, and they fail to notice the critical roles of other parties involved in the decision process (Bazerman, 2014).

Gilbert, Wilson, and colleagues (Gilbert & Wilson, 2000; Wilson, Wheatley, Meyers, Gilbert, & Axsom, 2000) discuss a related concept that they call "focalism." Discuss a related concept that they call "focalism." Focalism is the tendency to focus too much on a particular event (the "focal event") and too little on other important events. For example, people tend to overestimate the impact of a good or bad event (such as winning the lottery or being paralyzed) on their overall happiness (Wilson et al., 2000). They focus on a particular event and not on the hundreds of other elements of life that affect their happiness.

Schkade and Kahneman (1998) identify the "focusing illusion" as the human tendency to pay too much attention to a subset of available information while neglecting unattended information. They surveyed college students in the Midwest and in Southern California about their life satisfaction and their perceptions of others' life satisfaction. Both groups reported similar levels of life satisfaction. However, both groups rated Californians as having greater life satisfaction than Midwesterners. Schkade and Kahneman suggest that by focusing participants on location, the question caused the weather to be a salient factor in their responses, leaving many other factors that determine life satisfaction out of focus.

Similarly, Fox and Tversky (1998) find that participants who are asked to assess the likely success of a specific sports team focus on the team itself and ignore the many other teams that could affect the success of the focal team. The researchers recruited basketball fans as their study participants at a point when the 1995 National Basketball Association championship series was down to eight teams. They asked different participants to assess the probability that each team (from Chicago, Indiana, Orlando, New York, Los Angeles, Phoenix, San Antonio, and Houston) would win the championship and the probability that the winning team would come from either the Eastern or Western conference. If all of the participants were unbiased, the sum of the probabilities for the eight teams and the sum of the probabilities for the two conferences should each have added up to 100%. Indeed, the joint probability for the two conferences was 102%, which is a quite reasonable result. By contrast, the sum of the probabilities for the eight teams was 218%. Fox and Tversky suggest that as participants focus on each team, they develop a logic for that team winning, while failing to think about why other teams might win instead.

Another type of focusing failure is the failure of competitive actors to consider critical information needed for negotiation effectiveness—namely, the decisions of other parties and the rules of the game they are playing, find Idson et al. (2004) and Tor and Bazerman (2003). Their research uses variations of three well-studied decision problems: the Monty Hall game (Friedman, 1998; Nalebuff, 1987), the Acquiring a Company problem (Ball, Bazerman, & Carroll, 1991; Samuelson & Bazerman, 1985), and multi-party ultimatums (Messick, Moore, & Bazerman, 1997). Across these three games, most individuals make the wrong decision, although the games require no complex analytical reasoning. Rather, people make the wrong choice because they focus on their own thoughts and actions, ignoring key pieces of information, such as the rules of the game and the decisions of the opposing party.

Across these very different research arenas, bounded awareness poses a critical challenge, one that becomes clearer in the context of our bounded ethicality.

2. Behavioral ethics and bounded ethicality

2.1. Behavioral ethics

Social scientists have been developing knowledge about the actual behavioral determinants of ethical behavior for a long time. Scholars such as Brief and Motowidlo (1986), Treviño (1986) and Messick and Tenbrunsel (1996) pioneered the behavioral study of business ethics years before it became a popular topic. Among the many different definitions of behavioral ethics, Treviño, Weaver, and Reynolds (2006) define it as "individual behavior that is subject to or judged according to generally accepted moral norms of behavior" (page 952).

An early effort to chart out the area of behavioral ethics was an edited volume by Messick and Tenbrunsel (1996). This excellent collection provided an initial organization of what we knew about the psychology of ethics. In 1996, Messick and Bazerman organized early knowledge about behavioral ethics in terms of three types of theories that people use when making decisions: theories about the world, theories about other people, and theories about themselves. This early work departed from the tradition of studying ethics by focusing on the few "bad apples" whose unethical

behavior could be understood according to a rational analysis of its costs and benefits (Becker, 1978) to the study of the psychological tendencies that lead even good people to process information and make decisions in ways that lead to unethical behavior that they would not have anticipated.

Before the collapse of Enron in 2001, ethics scholarship and teaching in professional schools was largely in the hands of philosophers and self-defined business ethicists. The study and teaching of ethics in professional schools was largely normative; it lacked a basis for offering prescriptions or a descriptive understanding of how people actually respond to ethical challenges. For example, most approaches to ethics was focused on the role of individual traits in explaining intentional unethical behavior (Rest, 1986). Across several models (Jordan, 2009; Kohlberg, 1981; Rest, 1986; Reynolds, 2006), morality is treated as a rather stable personality attribute that individuals develop through moral development. These approaches neglect the power of situational factors influencing unethical behavior. In the last 15 years, the descriptive study of ethical behavior has emerged as a core area of empirical study and pedagogical development in the field of organizational behavior. This new, descriptive approach seeks to explain the ethical and unethical behavior of people based on situational and social factors.

Bazerman and Gino (2012) highlight the distinction between traditional approaches to ethics and behavioral ethics with their summary of the well-known trolley dilemma (Foot, 1978; Thomson, 1976; for detailed review, see Greene, 2013). The trolley dilemma is often described in the following manner (see Fig. 3 for visual depiction):

> A runaway trolley is headed for five railway workmen who will be killed if it proceeds on its present course. The only way to save these people is to hit a switch that will turn the trolley onto a side-track where it will run over and kill one workman instead of five. Ignoring legal concerns, is it ethically okay to turn the trolley in order to save five people at the expense of one?

This problem highlights the difference between the utilitarian and deontological approach in philosophical tradition. Utilitarianism is captured by the idea of doing "the greatest good for the greatest number of people." For utilitarians, there is a clear answer to the trolley dilemma: you hit the switch, acting to kill one person instead of five. In contrast, deontologists, basing their approach on Kant (1785/1959), call for judging the morality of an action based on its adherence to rules or duties. This approach judges the ethicality of behavior by the motives of the person who acts, not on the consequences of actions. Deontologists would not hit the switch.

Behavioral ethicists study the factors that influence people's answers to ethical dilemmas such as the trolley dilemma. Most people say that it is okay to hit the switch, using utilitarian reasoning to justify their choice: five people dying is worse than one person dying (Greene, 2013).

Contrast this result to common responses to the footbridge dilemma (see Fig. 4), which is as follows (Foot, 1978; Greene, 2013):

> A runaway trolley threatens to kill five people. You are standing on a footbridge spanning the tracks next to a railway worker wearing a large backpack, in between the oncoming trolley and the five people. The only way to save the five people is to push this man off the bridge and onto the tracks below. The man will die as a result, but his body will stop the trolley from reaching the others. (You can't jump yourself because you don't have enough weight on you to stop the trolley, and there's no time to put the backpack on.) Ignoring legal concerns, is it okay to save the five people by pushing this stranger to his death?

Most people oppose the idea of pushing the person (Barak-Corren, Tsay, Cushman, & Bazerman, in preparation; Greene, 2013). Now, instead of citing "the greatest good for the greatest number," respondents note that pushing the man would be "murder," that "people have rights," and that "the ends don't justify the means." (Greene, 2013).

While it was philosophers who first noted the apparent inconsistency in people's responses to the two problems (as a thought problem rather than an experiment, Foot, 1978; Thomson, 1976), behavioral ethicists view this observation alone as being of little help in improving ethicality in professional organizations. Psychologists, organizational behaviorists, and an emerging group of experimental philosophers (e.g., Greene, 2013; Knobe, 2003) run experiments to gather information about what people actually do and say when faced with moral dilemmas, explore how to improve ethicality, and examine how to most effectively help people make more deliberative ethical decisions.

Greene (2013) notes that it is critical that the original trolley dilemma kills the one bystander as result of an action to save lives—that is, the death is the byproduct of a seemingly wise action. In contrast, people are far less comfortable with the idea of using the person as a means of saving the other five, as would occur in the footbridge problem. Greene (2013) summarizes research that shows that these two problems trigger activity in very different brain regions, providing evidence paralleling the intuition of many that these are two very different problems (more on this later in article).

While philosophers have argued that philosophical thinking is central to moral education, will make us better citizens, and provide the courage to stand up for justice (summarized by Schwitzgebel, 2009), others argue that there is no empirical evidence to support these claims (e.g., Posner, 1997). As experimental philosopher Schwitzgebel (2009) notes, in contrast to most people's expectations, ethicists do not behave more ethically than others, as evidenced by the fact that across 31 leading academic

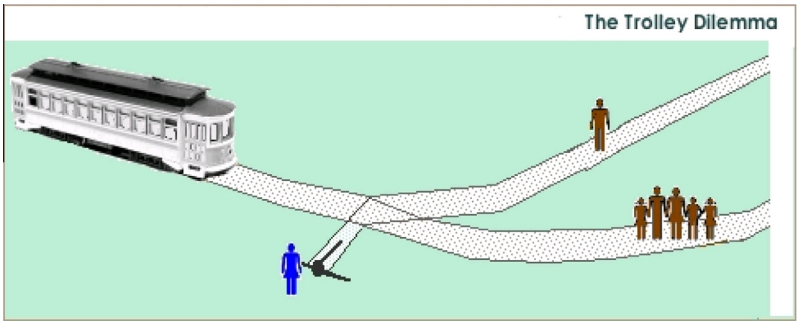

Fig. 3. The trolley dilemma.

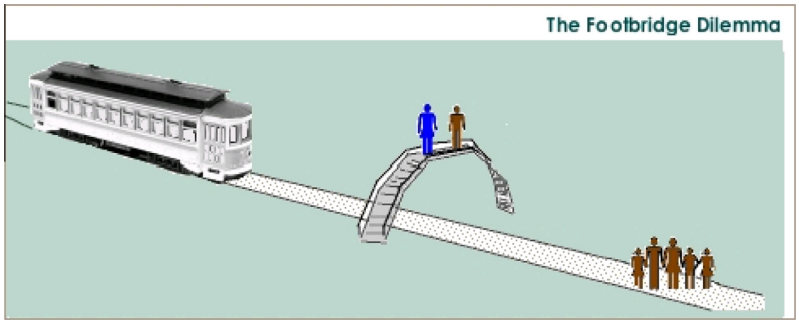

Fig. 4. The footbridge dilemma.

libraries in the United States and in the United Kingdom, books on ethics had been stolen 50–100% more often than non-ethics philosophy books that were comparable in age and popularity. More striking, the ethics texts most likely to be missing were obscure ones that only advanced students and professors of ethics are likely to seek out. These findings suggest that the traditional normative approach to ethics focuses on how people should behave in moral contexts but does not describe how unethical behavior occurs. Behavioral ethicists, on the other hand, aim to describe how people actually behave when they face an ethical dilemma.

2.2. Bounded ethicality

Within the broad topic of behavioral ethics is the much more specific topic of *bounded ethicality* (Chugh, Banaji, & Bazerman, 2005). Chugh et al. (2005) define bounded ethicality as the psychological processes that lead people to engage in ethically questionable behaviors that are inconsistent with their own preferred ethics. That is, if they were more reflective about their choices, they would make a different decision. This definition runs parallel to the concepts of bounded rationality (March & Simon, 1958) and bounded awareness (Chugh & Bazerman, 2007). In all three cases, a cognitive shortcoming keeps the actor from taking the action that she would choose with greater awareness. Importantly, if people overcame these boundaries, they would make decisions that are more in line with their ethical standards. Note that behavioral ethicists do not ask decision makers to follow particular values or rules, but rather try to help decision makers adhere more closely to their own personal values with greater reflection.

The literature on bounded ethicality has been reviewed elsewhere (Bazerman & Gino, 2012; Sezer, Gino, & Bazerman, 2015; Zhang, Gino, & Bazerman, 2014), and we do not seek to repeat such a review here. However, it may be useful to highlight some of the most common effects in this literature: Implicit attitudes and discrimination (Banaji & Greenwald, 2013), in-group/out-group biases (Bazerman & Gino, 2012) discounting the future (Wade-Benzoni, 1999), and overclaiming credit (Caruso, Epley, & Bazerman, 2006; Epley, Caruso, & Bazerman, 2006).

Bounded ethicality applies to all of us, yet few understand the degree to which our actual behavior can depart from the ethicality to which we aspire. To take one example, in 2004, Justice Antonin Scalia refused a request from the Sierra Club to recuse himself from a Supreme Court case, *Cheney v. U.S. District Court for D.C.* The case dealt with the question of whether Cheney should be required to provide information about an energy task force he led during the Bush administration's formulation of its energy policy. At the time, Scalia was a hunting buddy of Vice President Dick Cheney. The Sierra Club argued that the friendship was a threat to Scalia's objectivity in the Cheney case. "If it is reasonable to think that a Supreme Court justice can be bought so cheap, the nation is in deeper trouble than I had imagined," Scalia said in response (Janofsky, 2004). Scalia's comments indicate that he was unaware of the strong evidence that conflicts of interest have a strong psychological basis. Substantial evidence shows that people's fairness assessments are biased by the way in which they want to see the world (Bazerman, Loewenstein, & Moore, 2002; Moore, Tetlock, Tanlu, & Bazerman, 2006). Put differently, people may engage in ethically questionable acts without intending to do anything wrong. However, the rules governing Scalia's behavior on the Supreme Court, like most rules and laws concerning corrupt behavior, were only designed to guard against intentional corruption (Banaji, Bazerman, & Chugh, 2003).

This is one of the examples of the ways in which even well-intentioned people may engage in a wide variety of unethical actions without recognizing that they act unethically. The systematic patterns of behaviors in which people act unethically without their own awareness (Bazerman & Gino, 2012; Chugh et al., 2005) explain a number of troubling behaviors. As Messick (1994) argues, mortgage loan discrimination against minorities is likely rooted in lenders' unconscious favoritism toward in-groups rather than explicit hostility toward out-groups. Most people view themselves as environmentally or fiscally responsible, but behave contrary to these values, often without awareness that they are acting in ways that may harm future generations (Wade-Benzoni, 1999). In addition, Kern and Chugh (2009) show that the framing of information sways our ethical judgments: namely, that individuals are more likely to cheat to avoid losses than to obtain gains, simply because the same objective information is framed differently. Furthermore, misguided goals can cause people to focus too much on their objectives and blind them to ethical aspects that seem unrelated to the goal (Moore & Gino, 2013; Ordóñez, Schweitzer, Galinsky, & Bazerman, 2009; Staw & Boettger, 1990; Tenbrunsel, Wade-Benzoni, Messick, & Bazerman, 2000; Wolfe, 1988; Yaeger, 1986).

Overall, research shows that factors that fall outside of our awareness lead us to be ethically bounded. To highlight the bounded ethicality perspective, let's return to "Trolleyland." Barak-Corren et al. (in preparation) sought to determine whether different answers to the trolley and footbridge dilemmas reflect different preferences or whether judgment is bounded in one of these problems. They ran a three-condition design. The first condition was a modified version of the trolley dilemma, with only three people on the default set of tracks, such that the decision was whether to switch to save three lives at the cost of one life. The second condition was similar to the footbridge problem: five people were on the track, but instead of pushing one person off of the bridge, the choice was whether or not to flip a switch that would open a trap door that the man on the bridge was standing on and drop him onto the track, leading him to suffer an instant and painless death and save the five people on the track. Not surprisingly, for deontological reasons, significantly more switching occurs to save three lives in the first condition than to save five lives in the second condition.

To explore whether these answers reflect people's values, Barak-Corren et al. (in preparation) designed a third condition in which they used a procedure developed by Bazerman, Loewenstein, and White (1992) of comparing how people make decisions about one concrete option in comparison to how they make choices between options. Research shows that people are more deliberative when choosing among multiple options (Bazerman, Moore, Tenbrunsel, Wade-Benzoni, & Blount, 1999). Thus, in the third condition, Barack-Corren et al. asked participants to imagine there were two trains coming down two tracks, and the decision maker only had time to flip the switch to save three lives (at the expense of one) in the story of the Condition 1 train, flip the switch to save five lives (at the expense of one) in the story of the Condition 2 train, or do nothing. This time, flipping the switch to have the man fall from the bridge to save five lives was more popular than flipping the switch to switch the train to a different track to save three lives.

This study suggests that people often act more emotionally and deontologically when examining one option and that comparative decision making leads to more utilitarian approaches. Similarly, Bohnet, van Geen, and Bazerman (2015) show that while people show massive discrimination against females when choosing someone to hire to complete a math task, this discrimination largely disappears if the decision maker chooses between two or more candidates. These examples suggest that if we can identify contexts where people seem to be acting against their own more deliberative ethical preferences, we can devise environments to help them move more toward their more reflective ethical preferences.

More recently, we have been involved in research that documents the surprising degree to which people ignore the unethical actions of others. For example, Moore et al. (2006) find evidence that while we recognize others' conflicts of interest, we fail to recognize conflicts of interest that we ourselves face that lead to corrupt behavior. Thus, when we have a desire not to see the unethical actions of others, we do not, and we commit this failure often without our own awareness. As we describe in the next section, this tendency often leads individuals to ignore misconduct they witness in their own organizations.

3. Bounded awareness and ethical decision making

From Bernie Madoff's Ponzi scheme to corporate scandals in firms such as Enron, WorldCom, and Tyco International, instances of unethical behavior have proliferated in this century. The public reckoning typically results in a list of specific individuals who are deemed responsible for the disastrous outcomes. However, targeting the guilty seems to have no effect on the prevalence of unethical behavior. Incidents continue to make headlines, like the manipulation of the Libor (the London Interbank Offered Rate), a widely used benchmark interest-rate index, by five large global banks. More recently, high-level officials at FIFA, the international governing body of soccer, were accused of engaging in money laundering, bribery, wire fraud, and abuse of their positions to steal millions of dollars. Apart from these officials, there likely were many bystanders who failed to notice information that should have alerted them to unethical conduct. How could regulators, policymakers, and other decision makers have failed to see the evidence that was readily available? How do people ignore obvious evidence of wrongdoing until disaster strikes?

Bounded awareness is an important factor when bystanders fail to act to prevent others' ethical transgressions, as researchers have documented (Bazerman & Tenbrunsel, 2011). In significant cases of high-level organizational fraud, key decision makers—such as boards of directors, rating agencies, and auditors—typically are found to have had sufficient information to detect the unethical behavior, yet they failed to notice it, with direct consequences for society. If these people could have noticed and condemned others' ethical breaches at the time, these scandals could have been prevented.

Several psychological processes of bounded awareness lead individuals to ignore others' unethical behavior: the outcome bias, the role of the slippery slope, motivated blindness, and indirect blindness. These processes are largely about our cognitive failure to notice and act on others' unethical actions. Yet, many of the effects that we discuss have a motivational root. That is, our motivations affect our cognitions. The motivation to overlook others' unethical behavior affects us in many forms, including these cognitive biases.

3.1. Outcome bias

Before the financial collapse of 2008, investment advisors who consistently bought risky mortgage products were financially rewarded, while advisors who refrained from taking these risks produced lower returns and therefore were denied these rewards (Sorkin, 2010). Those who tried to make sound investments were penalized for their good intentions, while those who overlooked the signs of a forthcoming crisis were rewarded. In other words, the desirability of end results greatly influenced individuals' evaluations. And, before a bad outcome happened, few were concerned with a poor decision process that led to overinvestment in inappropriately risky assets.

Recent research in behavioral ethics suggests that individuals are far more likely to condemn others' questionable behavior when it leads to a bad outcome than when it does not. In a recent study, Sezer, Zhang, Gino, and Bazerman (in preparation) found that when evaluating others, lab participants, as well as managers from real companies with hiring and promotion authority, neglected information about people's intentions and relied too much on the favorability of outcomes when judging their actions. The study used a game that generates accidentally generous or accidentally selfish outcomes, developed by Cushman, Dreber, Wang, and Costa (2009). Specifically, researchers told study participants that their partners were asked to choose between the following two options (Cushman et al., 2009; Sezer et al., in preparation):

> Option A: You roll a six-sided die. If it comes up 1–4, you get $10, and the other party gets $0. If it comes up a 5, you get $5, and the other party gets $5. If it comes up a 6, you get $0, and the other party gets $10.
> Option B: You roll a six-sided die. If it comes up a 1, you get $10, and the other party gets $0. If it comes up 2–5, you get $5, and the other party gets $5. If it comes up a 6, you get $0, and the other party gets $10.

A partner who selects Option A has signaled selfish intentions, because the distribution of the outcomes is favorable to the die roller. On the other hand, selecting Option B signals fair intentions, because the distribution of the outcomes favors both parties equally. Researchers informed participants of their partners' die selection, and then the results of the roll and their winnings from the game. Then they gave participants the opportunity to punish or reward the partner who selected between the two options, without incurring any financial costs to themselves. Duplicating the results of Cushman et al. (2009) and Sezer et al. (in preparation) found that when deciding whether to punish or reward, participants paid more attention to the outcome of the rolled die—a random outcome—than to their partners' intentions, as signaled by their option selection. This evidence of the neglect of intentions is consistent with the long stream of research that shows that people overweigh outcome information in their evaluations of decision quality (Allison, Mackie, & Messick, 1996; Baron

& Hershey, 1988). Baron and Hershey (1988) refer to this tendency to overvalue outcome information as the outcome bias.

The outcome bias leads people to blame others too harshly for making sensible decisions that lead to accidentally unfavorable outcomes. Conversely, people often wait too long to condemn unethical behavior when they view its outcome positively—and they notice ethical violations only after a bad outcome emerges. Findings from behavioral ethics research has provided evidence that the outcome bias can alter ethical judgments by leading individuals to perceive the same behavior as more or less unethical depending on the outcome information. For instance, Gino, Moore, and Bazerman (2012) presented participants with various ethically questionable behaviors and provided outcome information to some participants and not others. The results indicated that the two groups of participants judged the same actions differently. Neglecting other critical information with a bias toward outcomes limits individuals' ability to notice unethical conduct. As this suggests, rather than judging the ethicality of others' decisions, we judge the outcome; consequently, we are often too slow to respond to others' unethical behavior.

Favorable outcomes lead people to ignore egregious behavior, and unfavorable outcomes lead them to condemn behavior they had previously ignored. This phenomenon has played out time and again in the U.S. auditing industry. Its practices prevent auditors from making independent evaluations of client firms (Bazerman et al., 2002; Frankel, Johnson, & Nelson, 2002; Moore et al., 2006), yet only disastrous outcomes have drawn attention to the misalignment between the current structure of the audit industry and required ethical standards (Bazerman & Watkins, 2004; Levitt & Dwyer, 2002).

The presence of an identifiable victim influences the outcome bias. Research on this effect suggests that people are far more concerned about specific identifiable victims than they are about statistical, unknown victims (Kogut & Ritov, 2005; Small & Loewenstein, 2003; Small & Loewenstein, 2005). In addition, Gino and colleagues found that people tend to judge unethical behavior far more harshly when it harms specific, identifiable victims than when it harms an anonymous group of people (Gino, Shu, & Bazerman, 2010); even supplying the name of the victim increases sympathy for him or her (Small & Loewenstein, 2003). As a form of the outcome bias, the identifiable victim effect shows a number of ways in which factors that have no relevance to the ethicality of a decision affect bystanders' evaluations of it.

3.2. Slippery slope

In addition to the outcome bias, the gradual erosion of others' unethical behavior has caused people to overlook warning signs of many of the biggest business scandals of the new millennium, including Enron's collapse and the UBS rogue trader scandal in 2011. Research has shown that individuals fail to see gradual changes that occur in front of their eyes (Simons, 2000). For instance, in one study, an experimenter with a basketball stopped pedestrians to ask for directions (Simons, Chabris, Schnur, & Levin, 2002). While the pedestrian was responding to the question, a group of confederates walked between the experimenter and the pedestrian. As they passed, another individual from the group replaced the experimenter. Most participants in the study failed to notice the change, but when researchers asked about the basketball the first experimenter had been holding, they recalled it. This suggests that the participants became accustomed to the change that occurred—thus they failed to notice explicitly that a visual change took place. Simons and Levin (1997) labeled this inability to detect visual changes as change blindness.

Recent behavioral ethics research has investigated the influence of gradual change on unethical behavior, or what researchers refer to as the slippery slope of unethical behavior (Tenbrunsel & Messick, 2004). Related research has found that when ethical decision making is at odds with self-interest, individuals often attempt to rationalize their unethical behavior (Pittarello, Leib, Gordon-Hecker, & Shalvi, 2015; Shalvi, Eldar, & Bereby-Meyer, 2012; Shalvi, Gino, Barkan, Ayal, & feeling moral. Current Directions in Psychological Science, in press) to maintain their view of themselves as moral (Mazar, Amir, & Ariely, 2008; Messick & Bazerman, 1996; Tenbrunsel, 1998). And the ethical numbing that occurs as a result of this slippery slope helps individuals rationalize their unethical behavior (Tenbrunsel & Messick, 2004).

When individual cheating cannot be identified and payoffs depend solely on self-reports, participants inflate their performance (Gino & Pierce, 2009; Mazar et al., 2008; Schweitzer, Ordóñez, & Douma, 2004). Welsh, Ordóñez, Snyder, and Christian (2015) investigated whether the slippery slope might exacerbate cheating levels. In this study, researchers gave their participants an opportunity to cheat in a problem-solving task. They told some participants that for each correct problem they solved, they would make $0.25 in the first round, $1.00 in the second round, and $2.50 in the final round. This gradual-change condition slowly increased incentives to cheat. In the abrupt-change condition, participants would not be compensated in the first two rounds, but would be compensated $2.50 per problem in the final round. Thus, there was an abrupt increase in their incentive to cheat. The researchers found that in the gradual-cheating condition, cheating levels increased in each successive round, but stayed lower and more constant in the abrupt-change condition. Sixty percent of participants cheated in the final round under the gradual condition, whereas only 30% of participants cheated in the final round in the abrupt-change condition. The slippery slope effect makes it easier to rationalize ethical indiscretions.

Gino and Bazerman (2009) performed a set of laboratory experiments to explore how the slippery-slope phenomenon affects the ability to notice others' unethical behavior. They assigned some participants to the role of estimators and others to the role of approvers. Estimators had to estimate the total value of pennies in glass jars, with awards for high estimates. Approvers acted as watchdogs to catch instances of overestimation. Consistent with the research on change blindness, approvers were much less likely to report the overestimates of those who gradually inflated their estimates over time than those who made the change in one abrupt shift. This research suggests that gradual deterioration of ethical behavior along a slippery slope limits people's ability to notice erosion of ethical behavior.

Related research has found that organizations are also vulnerable to the slippery-slope effect. For example, Schrand and Zechman (2012) examined firms for inappropriate accounting practices between 1996 and 2003. Researchers concluded that the majority of the fraud was rooted in a desire to adjust for a weak performance period, but that executives who were involved in minor accounting manipulation continued to violate accounting principles in upcoming periods. Their commitment to manipulation escalated when the company's performance did not improve. Interestingly, neither auditors nor the board of directors noticed these problems at the time.

Overall, these results suggest that people are much more likely to ignore others' misconduct when conditions encourage a side down a slippery ethical slope. The slippery-slope effect explains why so many ordinary people have ignored others' unethical behavior.

3.3. Motivated blindness

Research in social and cognitive psychology has documented that individuals see information that supports their view (Holyoak & Simon, 1999; Koehler, 1991), but fail to notice readily

available information that contradicts their preferences (Gilovich, 1991). People often evaluate evidence in such a way that affirms the conclusions they desire (Kunda, 1990) and validates their perceptions of reality (Shafir, Simonson, & Tversky, 1993). Similarly, individuals do not notice others' unethical behavior when noticing is not in their best interest (Bazerman & Tenbrunsel, 2011). This common failure of oversight—the tendency to overlook unethical behavior when it hurts the observer—is called motivated blindness.

One striking example of motivated blindness is the widespread use of steroids in Major League Baseball. In the 2007 season, the public started to question the performance of several superstars. In 2009, it became apparent that the use of performance-enhancing drugs was very common in the sport. Yet the commissioner, teams, players' union, and even the fans, who had financial or emotional investments in players' superb performance, were motivated to ignore signs of steroid use.

Members of boards of directors, credit rating agencies, or auditing firms are all susceptible to motivated blindness to unethical behavior because of their conflicts of interest (Bazerman et al., 2002; Moore et al., 2006). Arthur Andersen's failure to identify Enron's ethical lapses in its audits reflects such motivated blindness. Accounting firms are incentivized not to detect client errors because their clients can rehire them, fire them, or offer individual auditors employment within the firm (Bazerman et al., 2002). In a system where parties have a vested self-interest in viewing data in a certain way, they are not cognitively capable of making an unbiased and independent judgment (Moore et al., 2006). Individuals across a variety of industries with a conflict of interest (Issacharoff, 2005; Kassier, 2005) face a similar ethical dilemma between making an unbiased and independent judgment or acting in their own interest (Moore et al., 2006).

Moore, Tanlu, and Bazerman (2010) investigated whether individuals recognize their own conflicts of interest. They assigned participants to the role of a buyer, seller, buyer's auditor, or sellers' auditor of a fictional company. All participants received the same information and were asked to estimate the value of the firm. Consistent with the prior research on self-serving biases (Babcock & Loewenstein, 1997), the role to which participants were assigned led to substantial differences in their estimated valuations; sellers provided higher estimates for the company than buyers did. Further, participants who were assigned to the role of auditors were strongly biased toward their clients' interests.

These findings prompted the researchers to explore whether this decision was intentional or whether people were engaging in this behavior without their own awareness (Chugh et al., 2005). Moore et al. (2010) told the participants assigned to serve as auditors that they should give their best estimate of how a neutral expert would assess the company's true value, as they would be rewarded for the accuracy of their judgments. Sellers' auditors concluded that the firm was, on average, worth 30% more than that did buyers' auditors. Simply being assigned to a certain role in an experiment disabled participants from making an unbiased judgment. Given that the only reward promised was for accurate judgments, these results suggest that individuals engage in this behavior without even realizing they do so. In other words, at least a significant portion of this failure occurs unconsciously.

Some scholars have argued that professional auditors might be less vulnerable to these biases due to their expertise. However, Moore et al. (2010) replicated their results with professional auditors from large auditing firms (Moore et al., 2010). A large body of research suggests that experts with special training and knowledge are susceptible to the same biases that affect others (Buchman, Tetlock, & Reed, 1996; Cain, Loewenstein, & Moore, 2005; Cain, Loewenstein, & Moore, 2011; Cuccia, Hackenbrack, & Nelson, 1995; Loewenstein, Cain & Sah, 2011; Lowenstein, Sah & Cain, 2012; Moore et al., 2006). Motivated blindness affects ordinary people who view themselves as honest, as well as experts with special training and knowledge (Babcock & Loewenstein, 1997; Babcock, Loewenstein, Issacharoff, & Camerer, 1995; Pittarello et al., 2015; Shalvi et al., in press). These findings suggest that government involvement is needed to overhaul several industries, including the auditing industry, to create structures that allow organizations to provide what they promise society (Bazerman & Gino, 2012). This research also highlights that motivated blindness is an unconscious and very common failure to which even good people may be susceptible without their own awareness. Therefore, understanding the conditions under which motivated blindness influences our judgments is crucial to creating environments that enable us to make unbiased assessments.

3.4. Indirect blindness

People also change their ethical judgments when unethical acts are carried out through third parties (Gino & Bazerman, 2009; Gino & Pierce, 2009). For instance, in 2005, Merck, one of the world's leading pharmaceutical firms, owned the rights to a cancer drug called Mustargen. Mustargen had a small market share and generated sales of only about $1 million a year (Bazerman & Tenbrunsel, 2011). Merck could have increased profits by simply raising the price, but the media does not like dramatic price increases that affect cancer patients. So, Merck sold the rights for Mustargen to Ovation, a smaller pharmaceutical firm (Berenson, 2006). Merck continued to produce the drug on a contract basis after selling the rights to the drug. Ovation quickly increased the price by 1000 percent. The pharmaceutical giant sold the rights to Ovation, knowing that the smaller firm could increase the price while insulating Merck from negative publicity that might have damaged its image. If Merck had simply increased the price on its own by 500 percent, the public might have seen the increase in prices as the unethical exploitation of patients. But, through Ovation, the public largely ignored the dramatic price increase.

This example inspired Paharia, Kassam, Greene, and Bazerman (2009) to examine individuals' tendency to ignore price gouging when the main actor uses an intermediary. Since Paharia et al. (2009) reasoned that, had Merck raised the price directly, citizens would react in a negative way, they offered a scenario similar to Merck's behavior with respect to Mustargen:

> 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.

In one condition, participants assessed the ethicality of the following action:

> A: The major pharmaceutical firm raised the price of the drug from $3/pill to $9/pill.

In the other condition, participants evaluated 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/pill.

As predicted, participants judged Action A as more unethical as compared to those who judged Action B, even though Action A delivered a better price to patients.

Past research has found that evaluating two or more options at a time leads to more rational and reflective judgments than

one-at-a-time evaluations (Bazerman, Gino, Shu, & Tsay, 2011; Bazerman, Gino, Shu, & Tsay, 2013). Therefore, Paharia et al. (2009) provided a third group of participants with both scenarios and asked them to judge which action was more unethical. These participants judged Action B to be more ethically problematic than Action A.

Paharia et al. (2009) replicated these results across different domains, including contaminated land or pollution controls. Once again, participants judging one scenario at a time rated decision makers more harshly when they engaged in the unethical behavior directly than when they acted through an intermediary, while participants who compared the direct and indirect acts simultaneously reversed their opinion.

In further studies, researchers have tested whether the transparency of a firm's objectives can help mitigate indirect blindness (Paharia et al., 2009). In this study, participants were randomly assigned to one of four conditions. In the first condition, researchers gave a group of participants a scenario in which Company X raised the price of the drug directly. In the second condition, another group of participants were told that another firm, Company Y, which had bought the drug from Company X, raised the price itself. In the third condition, participants were told that Company X sold the drug to Company Y and was aware that Company Y would increase the price of the drug. Finally, in the fourth condition, researchers told participants that Company X made a contract with Company Y to increase the price and sell the drug indirectly. The greater the transparency of the firm's intent, the more harshly participants rated Company X. But they perceived direct action as more unethical than even the most transparent final condition. Taken together, these results demonstrate that people often ignore unethical action if the unethical actor has an intermediary do its dirty work.

In another study, Coffman (2011) created an experimental game to find out how much other parties would penalize an actor who engages in direct or indirect unethical behavior. In a four-player game adapted from the standard dictator game, the first mover received a fixed amount of money and determined how much to allocate to the receiver. The first mover could also sell the rights to the game to the intermediary at a negotiated price. If the intermediary bought these rights, then she played the dictator game with the same receiver. Finally, the punisher received a different fixed amount—and the opportunity to reduce the first mover's final payoff. The results showed that when the first mover played the dictator game with the receiver, the punisher penalized the first mover for giving smaller allocations. However, the punisher would reduce the first mover's payoff significantly less if he or she sold the rights to the game to the intermediary. In other words, as in the earlier experiments, participants punished direct unethical behavior more than they punished indirect unethical behavior. These results provide further support that intermediaries reduce punishment. Individuals assess the ethicality of direct unethical behavior far more harshly than indirect unethical behavior due to indirect blindness.

4. Future directions: How to improve ethicality

Part of the inspiration for the bounded ethicality agenda is the idea that if we can understand how humans deviate from their own more reflective preferences, we can design interventions that will lead them to their more ethical selves. Bazerman and Gino (2012) have discussed what such interventions might look in terms of behavioral ethics more broadly. This section focuses on opportunities for researchers to identify how to help people and organizations better notice and respond to the unethical actions of others.

We offer this agenda with full awareness that research on debiasing human judgment has a long history of limited success (Bazerman & Moore, 2013; Fischhoff, 1982). Yet, in the new millennium, we have witnessed two approaches to improving decision making that hold much more promise: (1) switching from System 1 thinking to System 2 thinking, and (2) nudges. We believe both provide powerful ideas for how to convince people and organizations more broadly to notice unethical behavior before disaster strikes.

4.1. Moving to System 2

Stanovich and West (2000) make the distinction between System 1 and System 2 cognitive functioning. We use System 1 thinking most of the time. It is our intuitive system, which is fast, automatic, effortless, implicit, and emotional. System 2 thinking refers to the slower, conscious, effortful, explicit, and logical processes we use to arrive at certain decisions (Kahneman, 2011). Decision analysis involves training people to use their System 2 processes in a rigorous manner. Chugh (2004) notes that the busier people are and the more they have on their minds, the more likely they are to rely on System 1 thinking. She further argues that the pace of managerial life provides a natural push toward System 1 thinking even in circumstances that clearly call for deliberation (Chugh, 2004). While System 2 processing cannot occur for even the majority of managerial decisions, we need to be aware that System 1 thinking is more prone to the biases so well documented in behavioral decision research (Bazerman & Moore, 2013). In addition, System 1 processing is less likely than System 2 to notice on others' unethical decisions and behavior.

Going back to the Madoff decision that introduced this article, we note that Zhang et al. (2015) observed that people are quick to invest in Fortitude (Madoff), but that the simple prompt "do you see any problem with any of the funds" quickly leads to the System 2 observation that Fortitude's returns are not possible. We believe that we need research to discover the conditions that prompt System 2 thinking so that individuals will notice unethical behavior more often.

We believe that as behavioral ethics permeates professional schools, it is important to think about the uniqueness of managerial and leadership positions in organizations. By definition, managers and leaders are responsible not just for their own behavior, but also for the behavior of others. In terms of ethics, this means that managers and leaders are responsible not only for acting ethically, but for noticing when unethical contact is occurring within their organizations. Many managers and leaders have not yet accepted this responsibility.

4.2. Nudging

Managers and leaders can also design institutions, organizations, and governments to change the way people see and react to problems. Thaler and Sunstein (2009) provide us with a structure to change institutions that takes into account the most accurate expectations of how humans behave, including their ethicality. They argue that by understanding how humans behave, we can intervene to design institutions and organizations that will "nudge" them toward more ethical decisions.

The good news is that most managers and leaders would like to notice the unethical behavior that occurs around them. The challenge for researchers is to devise nudges that can help leaders. Can we devise checklists that would make leaders more aware? Can we make changes to our annual reviews to identify problematic behaviors more quickly? Future research could seek to find the answers to these and other pressing questions.

4.3. Closing remark

We close with an anecdotal observation. Unethical behavior is often ambiguous. That is, when we encounter hints of unethical actions, we aren't sure that the action is unethical, and what to do about it is also unclear. Hints typically crop up in the blur of managerial life, amid other seemingly more critical problems that have a clear blueprint for action. As a result, we do not follow up by examining the ambiguous unethical behavior more closely. In reporting his personal failure as an expert witness to notice and act on corruption he observed at the Department of Justice, Bazerman (2014) writes that when he tells the story, managers absolve him quickly, given the ambiguity he faced. In contrast, investigative journalists typically mention that when *they* can't figure out what is going on in a given situation, that's a clue that they may have found an important story. One solution to noticing more might involve finding ways to encourage an investigative mindset, so that when leaders are puzzled about what is going on, they feel prompted to pursue answers, rather than turning their attention to other, more clearly specified problems.

References

Allison, S. T., Mackie, D. M., & Messick, D. M. (1996). Outcome biases in social perception: Implications for dispositional inference, attitude change, stereotyping, and social behavior. In M. P. Zanna (Ed.). *Advances in experimental social psychology* (Vol. 28, pp. 53–93). New York: Academic.

Babcock, L., & Loewenstein, G. (1997). Explaining bargaining impasse: The role of self-serving biases. *The Journal of Economic Perspectives*, 109–126.

Babcock, L., Loewenstein, G., Issacharoff, S., & Camerer, C. (1995). Biased judgments of fairness in bargaining. *The American Economic Review*, 1337–1343.

Ball, S. B., Bazerman, M. H., & Carroll, J. S. (1991). An evaluation of learning in the bilateral winner's curse. *Organizational Behavior & Human Decision Processes, 48* (1), 1–22.

Banaji, M. R., Bazerman, M. H., & Chugh, D. (2003). How (Un)Ethical are you? *Harvard Business Review, 81*(12), 56–64.

Banaji, M. R., & Greenwald, A. G. (2013). *Blindspot: Hidden biases of good people*. New York, NY: Random House.

Barak-Corren, N., Tsay, C. J., Cushman, F. A., & Bazerman, M. H. (2016). If you're going to do wrong, at least do it right: the counter intuitive effects of considering two moral dilemmas at the same time (in preparation).

Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology, 54*, 569–579.

Bazerman, M. H. (2014). *The power of noticing: What the best leaders see*. Simon and Schuster.

Bazerman, M. H., & Gino, F. (2012). Behavioral ethics: Toward a deeper understanding of moral judgment and dishonesty. *Annual Review of Law and Social Science, 8*(1), 85–104.

Bazerman, M. H., Gino, F., Shu, L. L., & Tsay, C. J. (2011). Joint evaluation as a real-world tool for managing emotional assessments of morality. *Emotion Review, 3* (3), 290–292.

Bazerman, M. H., Gino, F., Shu, L. L., & Tsay, C. J. (2013). The power of the cognition/emotion distinction for morality. *Emotion, 6*(1), 87–88.

Bazerman, M. H., Loewenstein, G., & Moore, D. A. (2002). Why good accountants do bad audits. *Harvard Business Review, 80*(11), 96–103.

Bazerman, M. H., Loewenstein, G. F., & White, S. B. (1992). Reversals of preference in allocation decisions: Judging an alternative versus choosing among alternatives. *Administrative Science Quarterly, 37*(2), 220–240.

Bazerman, M. H., & Moore, D. A. (2013). *Judgment in managerial decision making* (8th ed.). John Wiley and Sons Inc. Wiley.

Bazerman, M. H., Moore, D. A., Tenbrunsel, A. E., Wade-Benzoni, K. A., & Blount, S. (1999). Explaining how preferences change across joint versus separate evaluation. *Journal of Economic Behavior & Organization, 39*(1), 41–58.

Bazerman, M. H., & Tenbrunsel, A. E. (2011). *Blind spots: Why we fail to do what's right and what to do about it*. Princeton: Princeton University Press.

Bazerman, M. H., & Watkins, M. D. (2004). *The disasters you should have seen coming and how to prevent them*. Harvard Business School Publishing.

Becker, G. S. (1978). *The economic approach to human behavior*. Chicago: University of Chicago Press.

Berenson, A. (2006, March 12). A cancer drug's big price rise is cause for concern. The New York Times. Retrieved from <http://www.nytimes.com/2006/03/12/business/12price.html?n=Top%2FReference%2FTimes%20Topics%2FSubjects%2FF%2FFood>.

Bohnet, I., van Geen, A., & Bazerman, M. H. (2015). When performance trumps gender bias: Joint versus separate evaluation. *Management Science*. http://dx.doi.org/10.1287/mnsc.2015.2186 (in press).

Brief, A. P., & Motowidlo, S. J. (1986). Prosocial organizational behaviors. *Academy of Management Review, 11*, 710–725.

Buchman, T., Tetlock, P. E., & Reed, R. O. (1996). Accountability and auditors' judgments about contingent events. *Journal of Business Finance and Accounting, 23*, 379–398.

Cain, D. M., Loewenstein, G., & Moore, D. A. (2005). The dirt on coming clean: Perverse effects of disclosing conflicts of interest. *The Journal of Legal Studies, 34* (1), 1–25.

Cain, D. M., Loewenstein, G., & Moore, D. A. (2011). When sunlight fails to disinfect: Understanding the perverse effects of disclosing conflicts of interest. *Journal of Consumer Research, 37*(5), 836–857.

Caruso, E., Epley, N., & Bazerman, M. H. (2006). The costs and benefits of undoing egocentric responsibility assessments in groups. *Journal of Personality and Social Psychology, 91*(5), 857.

Chugh, D. (2004). Societal and managerial implications of implicit social cognition: Why milliseconds matter. *Social Justice Research, 17*(2), 203–222.

Chugh, D., Banaji, M., & Bazerman, M. H. (2005). Bounded ethicality as a psychological barrier to recognizing conflicts of interest. In D. Moore, D. Cain, G. Loewenstein, & M. Bazerman (Eds.). *Conflicts of interest: Challenges and solutions in business, law, medicine, and public policy*. New York: Cambridge University Press.

Chugh, D., & Bazerman, M. H. (2007). Bounded awareness: What you fail to see can hurt you. *Mind & Society, 6*(1), 1–18.

Coffman, L. C. (2011). Intermediation reduces punishment (and reward). *American Economic Journal: Microeconomics, 3*, 77–106.

Cuccia, A. D., Hackenbrack, K., & Nelson, M. (1995). The ability of professional standards to mitigate aggressive reporting. *The Accounting Review, 70*(July), 227–249.

Cushman, F., Dreber, A., Wang, Y., & Costa, J. (2009). Accidental outcomes guide punishment in a "trembling hand" game. *PLoS ONE, 5*, 1–7.

Epley, N., Caruso, E., & Bazerman, M. H. (2006). When perspective taking increases taking: reactive egoism in social interaction. *Journal of personality and social psychology, 91*(5), 872.

Fischhoff, B. Debiasing (1982). In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 422–432). Cambridge: Cambridge University Press.

Foot, P. (1978). *Virtues and vices*. Oxford: Blackwell.

Fox, C. R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science, 44*(7), 879–895.

Frankel, R. M., Johnson, M. F., & Nelson, K. K. (2002). The relation between auditors' fees for nonaudit services and earnings management. *The Accounting Review, 77* (Supplement), 71–105.

Friedman, D. (1998). Monty Hall's three doors: Construction and deconstruction of a choice anomaly. *The American Economic Review, 4*(88), 933–946.

Gilbert, D. T., & Wilson, T. D. (2000). Miswanting: Some problems in the forecasting of future affective states. In J. P. Forgas (Ed.), *Feeling and thinking: The role of affect in social cognition. Series: Studies in emotion and social interaction*. NY: Cambridge University Press.

Gilovich, T. (1991). *How we know what isn't so: The fallibility of human reason in everyday life*. New York: The Free Press.

Gino, F., & Bazerman, M. H. (2009). When misconduct goes unnoticed: The acceptability of gradual erosion in others' unethical behavior. *Journal of Experimental Social Psychology, 45*(4), 708–719.

Gino, F., Moore, D. A., & Bazerman, M. H. (2012). *No harm, no foul: The outcome bias in ethical judgments*. Working paper.

Gino, F., & Pierce, L. (2009). The abundance effect: Unethical behavior in the presence of wealth. *Organizational Behavior and Human Decision Processes, 109* (2), 142–155.

Gino, F., Shu, L. L., & Bazerman, M. H. (2010). Nameless+ harmless= blameless: When seemingly irrelevant factors influence judgment of (un) ethical behavior. *Organizational Behavior and Human Decision Processes, 111*(2), 93–101.

Greene, J. (2013). *Moral tribes: emotion, reason, and the gap between us and them*. New York: Penguin Press.

Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology General, 128*, 3–31.

Idson, L. C., Chugh, D., Bereby-Meyer, Y., Moran, S., Grosskopf, B., & Bazerman, M. H. (2004). Overcoming focusing failures in competitive environments. *Journal of Behavioral Decision Making, 17*, 159–172.

Issacharoff, S. (2005). Legal responses to conflicts of interest. *Conflicts of Interest: Challenges and Solutions in Business, Law, Medicine, and Public Policy*, 189–201.

Janofsky, M. (2004, February 6). Scalia's Trip With Cheney Raises Questions of Impartiality. *New York Times*. Retrieved from <http://www.nytimes.com/2004/02/06/politics/06SCAL.html>.

Jordan, J. (2009). A social cognition framework for examining moral awareness in managers and academics. *Journal of Business Ethics, 84*(2), 237–258.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Kahneman, D., & Tverksy, A. (1973). On the psychology of prediction. *Psychological Review, 80*, 237–251.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263–291.

Kant, I. (1959). *Foundation of the metaphysics of morals*. Indianapolis: Bobbs Merrill.

Kassier, J. P. (2005). *On the take: How America's complicity with big business can endanger your health*. New York: Oxford University Press.

Kern, M., & Chugh, D. (2009). Bounded ethicality: The perils of loss framing. *Psychological Science, 20*(3), 378–384.

Knobe, J. (2003). Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology, 16*, 309–324.

Koehler, D. J. (1991). Explanation, imagination, and confidence in judgment. *Psychological Bulletin, 110*, 499–519.

Kogut, T., & Ritov, I. (2005). The identified victim" effect: An identified group, or just a single individual? *Journal of Behavioral Decision Making, 18*, 157–167.

Kohlberg, L. (1981). *Essays on moral development, Vol. I: The philosophy of moral development*. Harper & Row.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin, 108*(3), 480–498.

Levitt, A., & Dwyer, P. (2002). *Take on the street*. New York, NY: Pantheon.

Loewenstein, G., Cain, D. M., & Sah, S. (2011). The limits of transparency: Pitfalls and potential of disclosing conflicts of interest. *American Economic Review, 101*(3), 423–428.

Loewenstein, G., Sah, S., & Cain, D. M. (2012). The unintended consequences of conflict of interest disclosure. *Journal of the American Medical Association, 307*(7), 669–670.

March, J. G., & Simon, H. A. (1958). *Organizations*. Oxford: Wiley.

Mazar, N., Amir, O., & Ariely, D. (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of Marketing Research, 45*, 633–644.

Messick, D. M. (1994, March 1). Mortgage-bias complexities. *Chicago Tribune*. Retrieved from <http://articles.chicagotribune.com/1994-03-01/news/9403010062_1_minority-applications-credit-worthiness-discrimination>.

Messick, D. M., & Bazerman, M. H. (1996). Ethical leadership and the psychology of decision making. *Sloan Management Review*, 9–22.

Messick, D. M., Moore, D. A., & Bazerman, M. H. (1997). Ultimatum bargaining with a group: Underestimating the importance of the decision rule. *Organizational Behavior & Human Decision Processes, 69*(2), 87–101.

Messick, D. M., & Tenbrunsel, A. E. (Eds.). (1996). *Codes of conduct: Behavioral research into business ethics*. New York: Russell Sage Foundation.

Moore, C., & Gino, F. (2013). Ethically adrift: How others pull our moral compass from true North, and how we can fix it. *Research in Organizational Behavior, 33*, 53–77.

Moore, D. A., Tanlu, L., & Bazerman, M. H. (2010). Conflict of interest and the intrusion of bias. *Judgment and Decision Making, 5*(1), 37–53.

Moore, D. A., Tetlock, P. E., Tanlu, L., & Bazerman, M. H. (2006). Conflicts of interest and the case of auditor independence: Moral seduction and strategic issue cycling. *Academy of Management Review, 31*, 10–29.

Nalebuff, B. (1987). Puzzles: Choose a curtain, duel-ity, two point conversions, and more. *Economic Perspectives, 1*(1), 157–163.

Neisser, U. (1979). The concept of intelligence. *Intelligence, 3*(3), 217–227.

Ordóñez, L. D., Schweitzer, M. E., Galinsky, A. D., & Bazerman, M. H. (2009). Goals gone wild: The systematic side effects of overprescribing goal setting. *Academy of Management Perspectives, 23*, 6–16.

Paharia, N., Kassam, K. S., Greene, J. D., & Bazerman, M. H. (2009). Dirty work, clean hands: The moral psychology of indirect agency. *Organizational Behavior and Human Decision Processes, 109*, 134–141.

Pittarello, A., Leib, M., Gordon-Hecker, T., & Shalvi, S. (2015). Justifications shape ethical blind spots. *Psychological Science, 26*, 794–804.

Posner, R. A. (1997). Against ethical criticism. *Philosophy and Literature, 21*(1), 1–27.

Rest, J. R. (1986). *Moral development: Advances in research and theory*. New York: Praeger.

Reynolds, S. J. (2006). Moral awareness and ethical predispositions: Investigating the role of individual differences in the recognition of moral issues. *Journal of Applied Psychology, 91*(1), 233–243.

Samuelson, W. F., & Bazerman, M. H. (1985). The winner's curse in bilateral negotiations. In V. Smith (Ed.). *Research in experimental economics* (Vol. 3). Greenwich, CT: JAI Press.

Schkade, D. A., & Kahneman, D. (1998). Does living in California make people happy? A focusing illusion in judgments of life satisfaction. *Psychological Science, 9*(5), 340–346.

Schrand, C. M., & Zechman, S. L. (2012). Executive overconfidence and the slippery slope to financial misreporting. *Journal of Accounting and Economics, 53*(1), 311–329.

Schweitzer, M. E., Ordóñez, L., & Douma, B. (2004). Goal setting as a motivator of unethical behavior. *Academy of Management Journal, 47*(3), 422–432.

Schwitzgebel, E. (2009). Do ethicists steal more books? *Philosophical Psychology, 22*, 711–725.

Sezer, O., Zhang, T., Gino, F., Bazerman, M. H. (2016) Overcoming the outcome bias: Making intentions matter (in preparation).

Sezer, O., Gino, F., & Bazerman, M. H. (2015). Ethical blind spots: Explaining unintentional unethical behavior. *Current Opinion in Psychology, 6*, 77–81.

Shafir, E., Simonson, I., & Tversky, A. (1993). Reason-based choice. *Cognition, 49*(1), 11–36.

Shalvi, S., Eldar, O., & Bereby-Meyer, Y. (2012). Honesty requires time (and lack of justifications). *Psychological Science, 23*(10), 1264–1270.

Shalvi, S., Gino, F., Barkan, R., Ayal, S. (2015). Self-serving justifications: Doing wrong and feeling moral. Current Directions in Psychological Science (in press).

Simon, H. A. (1957). *Models of man; social and rational*. New York: Wiley.

Simon, H. A. (1983). *Reason in human affairs*. Stanford: Stanford University Press.

Simons, D. J. (2000). Current approaches to change blindness. *Visual Cognition, 7*(1–3), 1–15.

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception, 28*(9), 1059–1074.

Simons, D. J., Chabris, C. F., Schnur, T., & Levin, D. T. (2002). Evidence for preserved representations in change blindness. *Consciousness and Cognition, 11*(1), 78–97.

Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences, 1*(7), 261–267.

Simons, D. J., & Levin, D. T. (2003). What makes change blindness interesting? *Psychology of Learning and Motivation, 42*, 295–322.

Small, D. A., & Loewenstein, G. (2003). Helping the victim or helping a victim: Altruism and identifiability. *Journal of Risk and Uncertainty, 26*(1), 5–16.

Small, D. A., & Loewenstein, G. (2005). The devil you know: The effect of identifiability on punitiveness. *Journal of Behavioral Decision Making, 18*(5), 311–318.

Sorkin, A. R. (2010). *Too big to fail: The inside story of how Wall Street and Washington fought to save the financial system—and themselves*. New York, NY: Penguin Books.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate. *Behavioral & Brain Sciences, 23*, 645–665.

Staw, B. M., & Boettger, R. D. (1990). Task revision: A neglected form of work performance. *Academy of Management Journal, 33*, 534–559.

Tenbrunsel, A. E. (1998). Misrepresentation and expectations of misrepresentation in an ethical dilemma: The role of incentives and temptation. *Academy of Management Journal, 41*(3), 330–339.

Tenbrunsel, A. E., & Messick, D. M. (1999). Sanctioning systems, decision frames, and cooperation. *Administrative Science Quarterly, 44*(4), 684–707.

Tenbrunsel, A. E., & Messick, D. M. (2004). Ethical fading: The role of self-deception in unethical behavior. *Social Justice Research, 17*(2), 223–236.

Tenbrunsel, A. E., Wade-Benzoni, K. A., Messick, D. M., & Bazerman, M. H. (2000). Understanding the influence of environmental standards on judgments and choices. *Academy of Management Journal, 43*, 854–866.

Thaler, R. H., & Sunstein, C. R. (2009). *Nudge: Improving decisions about health, wealth, and happiness*. New Haven, CT: Yale University Press.

Thomson, J. J. (1976). Killing, letting die, and the trolley problem. *The Monist, 59*, 204–217.

Tor, A., & Bazerman, M. H. (2003). Focusing failures in competitive environments: Explaining decision errors in the Monty Hall game, the acquiring a company problem, and multiparty ultimatums. *Journal of Behavioral Decision Making, 16*(5), 353–374.

Treviño, L. K. (1986). Ethical decision making in organizations: A person-situation interactionist model. *Academy of Management Review, 11*(3), 601–617.

Treviño, L. K., Weaver, G. R., & Reynolds, S. (2006). Behavioral ethics in organizations: A review. *Journal of Management, 32*(6), 951–990.

Wade-Benzoni, K. A. (1999). Thinking about the future: An intergenerational perspective on the conflict and compatibility between economic and environmental interests. *American Behavioral Scientist, 42*, 1393–1405.

Welsh, D. T., Ordóñez, L. D., Snyder, D. G., & Christian, M. S. (2015). The slippery slope: How small ethical transgressions pave the way for larger future transgressions. *Journal of Applied Psychology, 100*(1), 114–127.

Wilson, T. D., Wheatley, T., Meyers, J. M., Gilbert, D. T., & Axsom, D. (2000). Focalism: A source of durability bias in affective forecasting. *Journal of Personality & Social Psychology, 78*(5), 821–836.

Wolfe, D. (1988). Is there integrity in the bottomline? Managing obstacles to executive integrity. In S. Srivastva (Ed.), *Executive integrity: The search for high human values in organization life* (pp. 140–171). San Francisco: Jossey-Bass.

Yaeger, P. C. (1986). Analyzing corporate offenses: Progress and prospects. In J. Post (Ed.). *Research in corporate social performance and policy* (Vol. 8, pp. 93–120). Greenwich, CT: JAI Press.

Zhang, T., Fletcher, P. O., Gino, F., & Bazerman, M. H. (2015). Reducing bounded ethicality: How tohelp individuals notice and avoid unethical behavior. *Organizational Dynamics*.

Zhang, T., Gino, F., & Bazerman, M. H. (2014). Morality rebooted: Exploring simple fixes to our moral bugs. *Research in Organizational Behavior, 34*, 63–79.