



# **No harm, no foul: The outcome bias in ethical judgments**

**Francesca Gino  
Don A. Moore  
Max H. Bazerman**

**Working Paper**

**08-080**

Copyright © 2008, 2009 by Francesca Gino, Don A. Moore, and Max H. Bazerman

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Running Head: NO FOUL

No harm, no foul:  
The outcome bias in ethical judgments

Francesca Gino<sup>1</sup>

Don A. Moore<sup>2</sup>

Max H. Bazerman<sup>3</sup>

<sup>1</sup> *Kenan-Flagler Business School, University of North Carolina at Chapel Hill*

<sup>2</sup> *Tepper Business School, Carnegie Mellon University*

<sup>3</sup> *Harvard Business School, Harvard University*

The authors thank the Center for Behavioral Decision Research at Carnegie Mellon University and the Center for Decision Research at the University of North Carolina at Chapel Hill for support.

We thank Bill Bottom, Daylian Cain, Joshua Knobe, Lamar Pierce, David Pizarro, Zachariah Sharek, and three anonymous reviewers for their thoughtful comments on an earlier draft of the paper.

### Abstract

We present six studies demonstrating that outcome information biases ethical judgments of others' ethically-questionable behaviors. In particular, we show that the same behaviors produce more ethical condemnation when they happen to produce bad rather than good outcomes, even if the outcomes are determined by chance. Our studies show that individuals judge behaviors as less ethical, more blameworthy, and punish them more harshly, when such behaviors led to undesirable consequences, even if they saw those behaviors as acceptable before they knew its consequences. Furthermore, our results demonstrate that a rational, analytic mindset can override the effects of one's intuitions in ethical judgments. Implications for both research and practice are discussed.

*Key words:* outcome bias, unethical behavior, judgment, ethical decision making

No harm, no foul:  
The outcome bias in ethical judgments

A home seller neglects to inform the buyer about the home's occasional problems with flooding in the basement: The seller intentionally omits it from the house's legally required disclosure document, and fails to reveal it in the negotiation. A few months after the closing, the basement is flooded and destroyed, and the buyer spends \$20,000 in repairs. Most people would agree that the seller's unethical behavior deserves to be punished. Now consider the same behavior on the part of a second seller, except that it is followed by a long drought, so the buyer never faces a flooded basement. Both sellers were similarly unethical, yet their behavior produced different results. In this paper, we seek to answer the question: Do people judge the ethicality of the two sellers differently, despite the fact that their behavior was the same? And if so, under what conditions are people's judgments of ethicality influenced by outcome information?

Past research has shown some of the ways that people tend to take outcome information into account in a manner that is not logically justified (Baron & Hershey, 1988; Allison, Mackie, & Messick, 1996). Baron and Hershey (1988) labeled this tendency as the *outcome bias*. Extending prior work on the effect of outcome severity on judgments (Berg-Cross, 1975; Lipshitz, 1989; Mitchell & Kalb, 1981; Stokes & Leary, 1984), their research found that people judge the wisdom and competence of decision makers based on the nature of the outcomes they obtain. For instance, in one study participants were presented with a hypothetical scenario of a surgeon deciding whether or not to perform a risky operation (Baron & Hershey, 1988). The surgeon knew the probability of success. After reading about identical decision processes, participants learned either that the patient lived or died, and were asked to rate the quality of the

surgeon's decision to operate. When the patient died, participants decided it was a mistake to have operated in the first place.

This core result has been shown to be robust across various settings, including medical decision making, salespeople's decision making, and laboratory gambles (Baron & Hershey, 1988; Marshall & Mowen, 1993). In this paper, we explore how outcome information affects judgments of ethicality, and describe a series of laboratory studies which shed light on the judgmental process that underlies ethical assessments of behaviors that produced either positive or negative outcomes. Our studies also examine the role of rational versus intuitive mindsets as potential moderators of the relationship between outcome information and ethical judgments.

The research presented here offers three main theoretical contributions. First, it examines the role of the outcome information in assessments of ethicality for a variety of behaviors ranging from the mundane to the damnable. Second, our work addresses a possible confound present in many previous studies of the outcome bias: inferences about the decision maker's knowledge based on the outcome. The issue is that people might infer from the outcome what the decision maker must have known at the time of the decision. That is, the surgeon may have had more individuating information about the patient and the patient's probability of death. The study's participants, however, only had the base rates for post-surgical morbidity. Our second and third studies address this concern and show that it cannot account for our results. Finally, this research also investigates moderating factors for the effects of outcome information on ethical judgments. In our last two studies, we demonstrate that a rational, analytic mindset can help override the moral intuitions that drive the outcome bias in ethical judgments.

#### Outcome Bias across Disciplines

While prior research on the effects of outcome information has focused on judgments about decision quality, evaluations of the individual making the decision (e.g., her competence), and attributions of blame and responsibility as the main dependent variables, our primary focus is on ethical judgments. We also consider how judgments of ethicality are used as a base for forming judgments of blame and punishment, all of which are important concepts in law, philosophy and psychology.

Philosophical discussions of ethics, dating back to Aristotle's *Nicomachean Ethics* (350 BC/1974), have generally connected ethics closely to blame and punishment. According to philosophers, blame and punishment are directly assessed based on morally harmful actions (Bok, 1979). Under the law, however, ethics and punishment are only related more weakly. Consider the case of two hit-men, who engaged in identical behavior with identical motives: each of them tried to kill someone. While Gunnar shot and killed his target, Clod shot and missed. Both criminals did the same thing with the same goal, but the outcomes were different. According to the law, Clod is not even eligible for the most serious charge of murder, while Gunnar is (Kadish, Schulhofer, & Steiker, 2007).

The law punishes for two types of reasons – retributive and utilitarian (Kadish et al., 2007). Retributive punishments are justified based on whether people deserve the punishment, and are backward looking. Sometimes, an outraged society wants someone to pay for a grievous harm, and the law can provide for this form of retributive punishment. In contrast, utilitarian views are based on the deterring, rehabilitating, and incapacitating (Kadish et al., 2007). It is easy to argue that Gunnar and Clod are equally in need of rehabilitation and incapacitation. Even if this is so, a utilitarian perspective might endorse different punishments if the behavior produced better deterrent effects among other potential miscreants.

Consistent with arguments for a distinction between ethics and punishment, Kaplow and Shavell (2002) suggest that legal policies should be exclusively assessed by their effects on individual welfare and not based on issues of fairness. Thus, in their opinion, fairness should not be treated as an independent evaluative principle.

The law, then, does not consider the “outcome bias” a mistake, because the purpose of punishment goes beyond simple fairness. Thus, while ethics, blame, and punishment are clearly related, ethics is the construct closest to the action of the protagonist, and logically least connected to the issues raised by adjunct legal concerns. Our paper focuses on the role of outcomes in making judgments about the actions and ethics of the protagonists.

#### The Rationality of Ethical Condemnation

Too often, workers are evaluated based on results and not based on the quality of the decision (Rosenzweig, 2007). Given that most consequential business decisions involve some uncertainty, organizations too often over-reward luck and under-reward decision quality. If two decision makers use the same set of information to make the same final decision, their decisions should be evaluated similarly even if one resulted in a favorable outcome and the other resulted in an unfavorable one (Hastie and Dawes, 2001; Mazzocco, Alicke, & Davis, 2004). However, a long stream of research on the effects of outcome severity on judgment has robustly shown that ordinary decision evaluations are not consistent with these rational prescriptions (Allison, Mackie, & Messick, 1996). That is, people routinely use outcome information to evaluate the quality of a decision. This violates principles of rationality since it assigns weight to information that is not relevant for understanding the individual’s state of mind, abilities, or performance at the time of the decision.

Prior research on the use of outcome information focused on its effects on decision quality and evaluations of the individual making the decision (Baron & Hershey, 1988; Hershey & Baron, 1992). Related research has shown that the severity of a decision outcome influences attributions of both blame and responsibility, independently of other event features (e.g., Lowe & Medway, 1976; McKillip & Posavac, 1975; Schroeder & Linder, 1976; Shaw & Skolnick, 1971; Ugwuegbu & Hendrick, 1974; Walster, 1966). For instance, Alicke and Davis (1989) found that people blame a homeowner who shoots a trespasser more strongly when the trespasser is an innocent victim rather than a dangerous criminal. Related research by Mazzocco et al. (2004) demonstrated that even in the case of extremely low prior negligence on the part of the actor being evaluated, outcome effects on blame and sentencing recommendations still persevere. As noted by Mazzocco et al. (2004: p. 132):

In addition to violating the expectations of rational decision theories, these findings also breach the prescriptions of moral and legal philosophers who aver moral blame and legal responsibility should be based on prior sources of culpable behavior (Fischer, 1986; Morris, 1961). As embodied in jurisprudence, these a priori culpability sources include intentional, negligent or reckless wrongdoing (American Law Institute, 1962).

We suggest that outcome information exerts an important influence on judgment in ethical contexts. Preliminary evidence supporting this prediction is provided by work examining the effects of outcomes on children's moral judgments (e.g., Berg-Cross, 1975; Leon, 1982; Stokes & Leary, 1984; Surber, 1977). For instance, Berg-Cross (1975) found that first-grade children judge acts that resulted in large consequences (blind in one eye) as more deserving of punishment than behaviors that resulted in small consequences (sore eye).

Consistent with these findings, we argue that people judge unethical actions differently based on whether those actions led to positive or negative outcomes. That is, we suggest that outcome information impacts an observer's evaluation of the ethicality of a target's behavior. We



distinguish ethical from unethical actions based on Jones's (1991) definition of unethical behaviors as acts that have harmful effects upon others and are "either illegal or morally unacceptable to the larger community" (p. 367). Examples of unethical behaviors include violations of ethical norms or standards (whether they are legal standards or not), stealing, cheating, lying, or other forms of dishonesty. We hypothesize that:

Hypothesis 1: Individuals judge ethically-questionable behaviors that produced a negative outcome as more unethical than behaviors that produced a positive outcome.

### Blame and Punishment

Judgments about the ethicality of an actor's behavior are central to the case for blame and punishment. Judgments of how harshly to punish a behavior and how much to blame a certain behavior are affected by how unethical the behavior was in the first place (as suggested by work in the philosophy literature, see e.g., North, 1987, as well as by psychology research, see e.g., Goldberg, Lerner, & Tetlock, 1999). Research in the marketing literature has shown that ethical considerations are important in determining consumers' evaluations of a company and subsequent purchasing behavior. For instance, consumers' perceptions of a company's ethicality affect their willingness to buy its products (e.g., Creyer & Ross, 1997). Consistent with this research, Jones (1991) suggested a moral intensity model to address the issue-contingent nature of ethical decision making. According to this model, decision makers systematically evaluate the "intensity" of the essential features of the moral issue at hand when facing decisions in ethical domains. The perceived overall intensity of a moral issue then determines the decision maker's judgment, intent and subsequent behavior. In Jones' model, moral intensity varies depending on factors such as the magnitude of the consequence of the moral act at stake, the degree of social

agreement that the moral act is unethical, the likelihood of the act being carried out, as well as the likelihood of the act actually bringing about the expected consequences (Jones, 1991).

Building on this previous research, we suggest that an observer's perception of the ethicality of a target's behavior will be the base to determine whether the target deserves blame for the outcome resulting from such behavior, and whether the target deserves to be punished for their behavior. Thus, we hypothesize that:

Hypothesis 2a: Judgments of ethicality mediate the relationship between the nature of outcome information and judgments of punishment for the behavior under consideration.

Hypothesis 2b: Judgments of ethicality mediate the relationship between the nature of outcome information and judgments of blame for the behavior under consideration.

#### Are Intuitive Evaluations of Ethicality More Susceptible to the Outcome Bias?

While the outcome bias has been shown to be robust across a variety of settings, we suggest that its effect on judgments of ethicality is moderated by the approach decision makers use to process information. Various theories in psychology and cognitive sciences suggest that individuals can operate in either of two different modes of thought (Chaiken, 1980; Devine, 1989; Epstein, 1994; Epstein, Lipson, Holstein, & Huh, 1992; Wilson, Lindsey, & Schooler, 2000). Information is processed differently by these two systems. The first one (labeled as System 1, see Stanovich, 1999; Stanovich & West, 2000) is intuitive, quick and heuristic. System 1 includes instinctive behaviors that are innately programmed and it is intimately associated with affect. This mode of processing information is based on "gut-feelings" and emotions. The second one (System 2) is analytic, rational and systematic. System 2 permits abstract hypothetical thinking that cannot be achieved by System 1 and it operates at the conscious level.

As demonstrated by a variety of studies, the distinction between these two modes of thinking is useful for understanding social behavior and individual judgment across several contexts. For example, the use of System 1 leads individuals to utilize cues such as the representativeness and availability heuristics (see Kahneman & Tversky, 1982; Tversky & Kahneman, 1974), producing errors in probability judgments. However, when participants are first motivated to process information in a rational mode, reliance upon heuristic cues is reduced and more effortful and accurate judgments occur (for a review, see Epstein & Pacini, 1999). Thus, as these studies suggest, if individuals commonly form judgments intuitively, asking them to think rationally changes their evaluations of given behaviors or situations.

Research in psychology has suggested that individuals possess an “intuitive ethics,” defined as “an innate preparedness to feel flashes of approval or disapproval toward certain patterns of events involving other human beings” (Haidt & Joseph, 2004: 56; see also Haidt, 2007). Haidt (2001) proposed that individuals make moral judgments quickly and intuitively, based on their gut feelings of right and wrong. Similarly, ethical judgments may be based on one’s own intuitions, and these intuitions might conflict with the decisional outcome of a rational approach to judgment.

Because of our intuitive ethics, the differential impact of outcome information on ethical judgments might occur only in one mode of thought, the intuitive one, but not in the rational mode of thought. This reasoning leads us to hypothesize that:

Hypothesis 3: Outcome information biases ethical judgments more when the observer is using an intuitive mindset than when the observer is using a rational, deliberative mindset.

*The Present Research*

We tested our predictions in six laboratory studies in which participants were asked to consider various ethically-questionable behaviors. Participants were also given information about the outcome of such behaviors and were asked to rate the ethicality of the described actions with or without the outcome information. The studies manipulated whether the described behaviors were followed by a negative or positive consequence, both between subjects (Studies 1 and 4) and within subjects (Studies 2 and 3). In addition, our studies manipulated whether participants approached their judgments with a rational or intuitive mindset (Studies 5 and 6). The results illustrate how the outcome bias affects ethical judgments and demonstrate the psychological forces at work in these effects.

### Study 1

In Study 1 we sought to provide initial evidence that individuals would judge behaviors that resulted in negative consequences as more unethical than behaviors that resulted in positive consequences. We also sought to investigate the relationship between ethical judgments and judgments of punishment and blame.

### Methods

*Participants.* In return for extra course credit, 120 undergraduates (53% male) from a university in the United States participated in the study. The average age of participants was 20 ( $SD = 1.15$ ).

*Procedure.* Participants were randomly assigned to one of two conditions: positive-outcome condition or negative-outcome condition. In each condition, participants read three scenarios, which described the behavior of a pharmaceutical researcher, an auditor, and a toy company (see Appendix A for a description of the scenarios used). In creating the materials, we matched a less unethical action with the negative outcome, and a more unethical action with a

positive action, in order to test whether the outcome would overcome the ethicality of the actual decision.

In the positive-outcome condition, the descriptions included elements of unethical or questionable practices but a positive outcome. For instance, one scenario read,

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 within his firm. He believed that the product was 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.

In the negative-outcome condition, instead, the descriptions of unethical or questionable practices were less extreme from an ethical standpoint, but the outcome for each scenario was negative. We changed how ethically-questionable the described practices were (instead of keeping them constant) to make the test for an outcome bias in the ethics realm more conservative. The scenario describing the behavior of a pharmaceutical researcher in the negative-outcome condition read,

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 within his firm. As the deadline approaches, he notices that four subjects were withdrawn from the analysis due to technicalities. He believes that the data in fact is 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.

After reading each scenario, participants were asked to use a 7-point scale to indicate 1) the extent to which they found the described behavior unethical, 2) how harshly they would punish the described behavior, and 3) the extent to which they would blame the described behavior for the observed outcome.

*Pilot.* We conducted a pilot study with a separate group of undergraduates ( $N = 87$ , 53% male) to test that we were effective in writing the scenarios so that the elements of unethical or questionable practices were indicative of a less egregious behavior in the negative-outcome condition than in the positive-outcome condition of the main study. The pilot used the same scenarios as in Study 1, but this time no outcome information was given to participants. The order in which scenarios were presented to participants was counterbalanced. Participants in the highly-unethical behavior condition read the scenarios included in the positive-outcome condition of Study 1. For instance, the scenario describing the behavior of the pharmaceutical researcher read,

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 within his firm. He believed that the product was 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.

Participants in the slightly-unethical behavior condition read the scenarios included in the negative-outcome condition of Study 1. The scenario describing the behavior of the pharmaceutical researcher in this case read,

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 within his firm. As the deadline approaches, he notices that four subjects were withdrawn from the analysis due to technicalities. He believes that the data in fact is 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.

Each participant read three scenarios and for each used a 7-point scale to indicate 1) the extent to which he or she found the behavior described in the scenario unethical, 2) how harshly he or she would punish the behavior described in the scenario, and 3) what outcome he or she

thought would result from the described behavior. We used participants' ratings on each of these three questions as the dependent variable in repeated-measures ANOVAs in which experimental condition served as between-subjects factor. The results showed that behaviors were rated as more unethical in the positive-outcome ( $M = 5.88, SD = 0.70$ ) than in the negative-outcome condition ( $M = 4.83, SD = 1.23$ ),  $F(1, 85) = 23.80, p < .001, \eta^2 = .219$ . In addition, participants indicated that (1) they would punish such behaviors more harshly (5.64 vs. 4.85,  $F[1, 85] = 14.89, p < .001, \eta^2 = .15$ ) and (2) they thought such behaviors would lead to more negative outcomes (5.53 vs. 4.95,  $F[1, 85] = 9.64, p = .003, \eta^2 = .10$ ). The results of this pilot study suggest that the behaviors described in the three scenarios of the positive-outcome condition used in Study 1 can in fact be considered more extreme than the behaviors described in the three scenarios of the negative-outcome condition. Thus, the results of this pilot study indicate that the manipulation used in Study 1 was effective.

### Results

Table 1 reports the descriptive statistics and correlations for the main variables included in Study 1. We first tested whether people's judgments of an action's ethicality would vary based on its consequences. We began by comparing participants' ratings for the three variables measured in the study (unethicality, punishment, blame) across conditions. The ethicality ratings were subjected to a repeated-measures ANOVA in which experimental condition served as between-subjects factor. This analysis revealed that unethicality ratings were higher in the negative-outcome ( $M = 5.77, SD = 0.79$ ) than in the positive-outcome condition ( $M = 5.28, SD = 1.13$ ),  $F(1, 118) = 7.35, p = .008, \eta^2 = .059$ . These results support Hypothesis 1 and suggest that the outcome bias extends to ethical judgment.

Similarly, the desire to punish was stronger in the negative-outcome ( $M = 5.46$ ,  $SD = 0.94$ ) than in the positive-outcome condition ( $M = 5.01$ ,  $SD = 1.10$ ),  $F(1, 118) = 5.81$ ,  $p = .017$ ,  $\eta^2 = .047$ . Finally, the rating for blame was higher in the negative-outcome ( $M = 5.49$ ,  $SD = 0.90$ ) than in the positive-outcome condition ( $M = 4.65$ ,  $SD = 1.40$ ),  $F(1, 118) = 15.60$ ,  $p < .001$ ,  $\eta^2 = .117$ .

Hypothesis 2a predicted that ethical judgments would mediate the relationship between the outcome and judgments of punishment for the behavior under consideration. We tested this hypothesis using the criteria prescribed by Baron and Kenny (1986). In our first regression, we used outcome information as the independent variable (1 = positive, 0 = negative) and the ratings for punishment as the dependent variable. As expected, this relationship was significant and negative ( $\beta = -.45$ ,  $p = .017$ ). In the second regression, we tested the relationship between outcome information and ratings for unethicity. This relationship was also significant and negative ( $\beta = -.48$ ,  $p = .008$ ), indicating that those in the positive outcome condition reported lower ratings for unethicity than did those in the negative outcome condition. In the final step, we included outcome information and ratings for unethicity as independent variables and judgments of punishment as the dependent variable. Supporting Hypothesis 2a (Sobel test,  $Z = -2.59$ ,  $p = .009$ ), the path between outcome information and judgments of punishment became insignificant ( $\beta = -.09$ ,  $p = .52$ ) when the direct influence of judgments of ethicality was included in the regression ( $\beta = .75$ ,  $p < .001$ ). These results suggest that judgments of ethicality fully mediate the relationship between outcome information and judgments of punishment. We depict the mediation results in Figure 1.

Hypothesis 2b predicted that ethical judgments would mediate the relationship between the nature of outcome information and judgments of blame. We conducted a mediation analysis



this time using blame judgments as the dependent variable. Supporting Hypothesis 2b (Sobel test,  $Z = -2.42, p = .02$ ), the path between outcome information and blame judgments became less significant ( $\beta = -.57, p = .004$ ) when the direct influence of judgments of ethicality was included in the regression ( $\beta = .57, p < .001$ ). These results suggest that judgments of ethicality partially mediate the relationship between outcome information and judgments of blame. Figure 2 depicts these results.

### Discussion

The results of our first study suggest that participants were more critical in the scenarios presented in the negative-outcome condition than they were in the scenarios presented in the positive-outcome condition. Yet, as demonstrated by the pilot study, the behavior described in each scenario was more unethical in the positive-outcome condition than in the negative-outcome condition. In addition, the results of Study 1 showed that ethical judgments mediated the relationship between outcome information and judgments of both blame and punishment.

### Study 2

The first study showed that learning about outcomes influences people's assessment of ethicality. Our second study makes two new contributions. First, we start addressing an alternative explanation for our findings: the "inside knowledge" explanation. This explanation argues that the decision-makers whose behavior is being evaluated must have had better information about the likely consequences of their actions than did the judges who are evaluating them. This is a viable explanation for most prior studies investigating the outcome bias or related phenomena (e.g., Lipshitz, 1989; Mitchell & Kalb, 1981).

Consistent with this explanation, it seems reasonable to assume that the decision-makers had inside knowledge of the chances of a negative outcome. Some of this uncertainty on the part

of participants in the role of judges is resolved by learning the outcome. In other words, learning that the toy company sold a poisonous product that killed six children (as in the case of one of the scenarios used in Study 1) implies that this possibility may have been in the mind of the company's managers at the outset, and they knew that this might be a consequence of outsourcing. After all, if one has no way to assess the quality of a decision, then all one is left with is assessing its outcome. Study 2 starts addressing this alternative explanation by employing scenarios in which the decision makers do not have any inside knowledge about the chances of a positive or negative outcome—it is a chance outcome over which the decision maker has no control.

The second contribution of our second study is the use of a within-subjects design, in which participants first evaluate the quality of the decision without knowing the outcome, then learn the outcome and evaluate the decision again using the same criteria. This within-subjects design allows us to test our contention that the outcome bias does not depend on incorrectly imagining how one would have evaluated the decision in the absence of outcome knowledge, and also allows us to show that the outcome bias observed in Study 1 truly results from the valence of outcome information.

### Methods

*Participants.* Fifty-eight individuals (69% male; 91% students) participated in the study and were paid \$7. Most participants (91% of them) were students from local universities. The average age of participants was 26 ( $SD = 8.58$ ).

*Procedure.* The study included three phases. Participants were randomly assigned to one of two conditions: positive-outcome or negative-outcome conditions. In the first phase of the study, participants were asked to read three scenarios, which were the same across conditions

(see Appendix B for a description of the scenarios used). Each scenario included elements of unethical or questionable practices, but no information about outcomes. For instance, one of the scenarios read,

A government agency in a developing country finds itself dealing with a natural disaster in which several thousand poor peasants have been made homeless during the winter. The agency must decide what sort of short-term housing it will provide for the refugees. The inexpensive option is tents, which will probably be fine, given the mildness of the local winters—overnight temperatures only fall below freezing once every four years or so, on average. The more expensive option is to put up temporary shacks that would provide more shelter against the cold. But the shacks would be more expensive and would force the agency to cut funding to other (less urgent) programs. In the end, the agency's commissioner decides to provide only tents for the refugees.

Participants in both conditions read this same scenario. After reading each of three scenarios, participants rated the ethicality of the described behavior using a 7-point scale (from 1 = very ethical to 7 = very unethical).

In the second phase of the study, participants completed a 7-minute filler task, included in order to reduce the pressure for participants to respond consistently in the first and third phases.

In the third phase, participants were asked to read a new set of scenarios. The scenarios presented to participants at this stage included the same information as in Phase 1, but this time the scenarios also included information about the outcome. We varied whether the outcome was negative (e.g., “The winter is substantially colder than expected, and fifty children among the refugees die of exposure to the cold”) or positive (e.g., “The winter is quite mild, and the tents provide sufficient shelter”). Finally, participants rated the ethicality of the behavior described in each scenario using a 7-point scale.

## Results

As one might expect, given that the manipulation occurred in the third phase of the study, there were no significant differences across conditions in the rating for the ethicality of the

described behavior in Phase 1 ( $M_{negative-outcome} = 4.65$  vs.  $M_{positive-outcome} = 4.83$ ,  $F [1, 56] < 1$ ,  $p = .54$ ). Yet, in Phase 3 (after the manipulation occurred) participants rated the behaviors described in the negative-outcome condition as more unethical ( $M = 6.26$ ,  $SD = 1.04$ ) than those in the positive-outcome condition ( $M = 3.83$ ,  $SD = 1.68$ ),  $F (1, 56) = 75.24$ ,  $p < .001$ ,  $\eta^2 = .57$ .

In order to verify that ratings of unethicality increased from Phase 1 to Phase 3 for participants in the negative-outcome condition, we conducted a repeated-measures ANOVA with type of outcome (positive vs. negative) as between-subjects factor and change in rating as within-subjects factor (repeated measure on scenario). We used the ratings of unethicality provided in both Phase 1 and 3 as the dependent measure. This analysis revealed a significant main effect for type of outcome,  $F (1, 56) = 27.03$ ,  $p < .001$ ,  $\eta^2 = .33$ . The main effect of change in rating did not reach significance ( $p = .11$ ), but the change in rating X type of outcome interaction did ( $F [1, 56] = 49.91$ ,  $p < .001$ ,  $\eta^2 = .47$ ). The interaction is depicted in Figure 3.

While in the positive-outcome condition, the ratings of unethicality significantly decreased from Phase 1 to Phase 3 ( $F [1, 29] = 11.97$ ,  $p = .002$ ,  $\eta^2 = .29$ ), they significantly increased in the negative-outcome condition ( $F [1, 27] = 51.67$ ,  $p < .001$ ,  $\eta^2 = .66$ ). These results suggest that, when participants learned the outcomes of the described behaviors were positive, they rated the behaviors as more ethical than they did when they were not given information about the outcome. Instead, when participants learned the outcomes of the described behaviors were negative, they rated the behaviors as more unethical than they did when they were not given information about the outcome. These results provide further support for Hypothesis 1, which predicted that in ethical domains individuals would judge behaviors that produced a negative outcome as more unethical than behaviors that produced a positive outcome.

## Discussion

The results of our second study provide further support for the hypothesis that ethical judgments are influenced by the nature of the outcome produced by ethically-questionable behaviors. Using a within-subjects design, Study 2 demonstrates that even when participants saw and rated the ethicality of a behavior before, their opinions changed when they learned the outcome of the described behavior. In particular, participants judged the decisions that led to negative outcomes as unethical, even if they did not think so before, and they have not learned anything about the decision-maker when receiving information about the decision outcome. These results demonstrate that the effect of outcome information on ethical judgments observed in Study 1 results from the valence of outcome information provided to participants.

### Study 3

Our second study attempted to rule out the inside knowledge explanation for the outcome bias. Yet, one might argue that learning about the nature of the outcome can lead decision makers to make different assumptions about information the actor might have had at the time of the decision (e.g., the riskiness of the decision) and that even in the scenarios used in Study 2 participants could have made implicit assumptions about information known to the decision maker once they learned about the outcome. Study 3 used a very simple scenario in which such assumptions cannot be made.

### Methods

*Participants.* Seventy-one college students (39 male) participated in the study and were paid \$5. The average age of participants was 21 ( $SD = 1.37$ ).

*Procedure.* The study used the same procedure as in Study 2 but varied the type of scenario participants read. Participants were randomly assigned to one of two conditions:

positive-outcome or negative-outcome conditions. In the first phase of the study, participants were asked to read the following scenario,

Imagine you were to judge the behavior of participants in a study we conducted in the past. The study examines the behavior of participants in a given role (Player A) who are asked to make allocation decisions of resources between themselves and participants in another role in the study (Player B). Participants in the study were randomly assigned to one of two roles (Player A or Player B) and randomly paired up so that we had one Player A and one Player B in each dyad. Player A is asked to consider the following two options: Option 1) Player A receives \$5, and Player B receives \$5. Option 2) Player A receives \$6, and the experimenter will toss a fair coin. If heads, then Player B will receive \$5. If tails, then Player B will receive nothing. Imagine Player A chose the second option, namely the one in which Player A receives \$6 and Player B's payoff is determined by a coin toss.

Participants in both conditions read this same scenario. After reading each of three scenarios, participants indicated the extent to which the described behavior was unfair, unethical and wrong using a 7-point scale (from 1 = not unethical at all / not unfair at all / not wrong at all to 7 = very unethical / unfair / wrong).

In the second phase of the study, participants completed a 5-minute filler task, included in order to reduce the pressure for participants to respond consistently in the first and third phases.

In the third phase, participants were asked to read the same scenario but this time we included information about the outcome. We varied whether the outcome was negative (e.g., "Player A chose Option 2. The experimenter tossed the coin, and the result of the coin toss was tails. Player A thus received \$6, and Player B received no payment for the study") or positive (e.g., "Player A chose Option 1. The experimenter tossed the coin, and the result of the coin toss was heads. Player A thus received \$6, and Player B received \$5 for the study"). Finally, participants rated the unethicality, unfairness and wrongness of the behavior described in each scenario using a 7-point scale as they did in the first phase.

## Results

The three measures used to capture judgments of ethicality (unfair, unethical, and wrong) were highly correlated with one another and we thus averaged them to create a single measure (both Cronbach's alphas  $> .80$ ). As expected, there were no significant differences across conditions in the rating for the ethicality of the described behavior in the first phase ( $M_{negative-outcome} = 4.34$  vs.  $M_{positive-outcome} = 4.63$ ,  $F [1, 69] < 1$ ,  $p = .45$ ,  $\eta^2 = .01$ ). Yet, in the third phase, when participants received outcome information, participants rated Player A's behavior as more unethical in the negative-outcome condition ( $M = 5.86$ ,  $SD = 0.98$ ) than in the positive-outcome condition ( $M = 3.75$ ,  $SD = 1.58$ ),  $F (1, 69) = 45.09$ ,  $p < .001$ ,  $\eta^2 = .40$ .

Next, we conducted an ANOVA with type of outcome (positive vs. negative) as between-subjects factor and change in rating as within-subjects factor. We used the ratings of unethicality provided in both Phase 1 and 3 as the dependent measure. This analysis revealed a significant main effect for type of outcome,  $F (1, 69) = 10.39$ ,  $p = .001$ ,  $\eta^2 = .13$ . The main effect of change in rating did not reach significance ( $p = .12$ ), but the change in rating X type of outcome interaction did ( $F [1, 69] = 34.53$ ,  $p < .001$ ,  $\eta^2 = .33$ ). The interaction is depicted in Figure 4.

While in the positive-outcome condition, the ratings of unethicality significantly decreased from Phase 1 to Phase 3 ( $F [1, 35] = 8.50$ ,  $p < .01$ ,  $\eta^2 = .20$ ), they significantly increased in the negative-outcome condition ( $F [1, 34] = 30.83$ ,  $p < .001$ ,  $\eta^2 = .48$ ).

### Discussion

The results of Study 3 are consistent with the findings of Study 2 and provide further evidence for the outcome bias in the ethics realm. They also rule out the inside knowledge explanation as a potential explanation for the findings observed in the first two studies. In Study 3, indeed, there was no information contained in the outcome that could have shed light on Player A's state of mind at the time of Player A's choice.

## Study 4

We have argued that ethical judgments, focused as they are on the mind and intentions of the protagonist, ought not be affected by outcome information. However, it is possible that the severity and importance of ethical outcomes might actually make them even more susceptible to the outcome bias. We designed a fourth study to investigate such possibility. Our fourth study adds a new condition (labeled as negative-financial outcome condition) in which the described outcome is not only negative but also not-ethically relevant –or at least not as extreme as the outcome used in the negative-outcome conditions of Studies 1 and 2.

### Methods

*Participants and Procedure.* One-hundred forty individuals (52% male) participated in the study in exchange of \$6. Most participants (82% of them) were students from local universities. The average age of participants was 24 ( $SD = 5.31$ ). Participants were randomly assigned to one of three conditions: positive-outcome condition, negative-financial outcome condition or negative-unethical outcome condition. In all cases, participants read a description of a company's decision to locate its manufacturing facility in India that turned out badly and lost money (see Appendix C for a description of the scenarios used). In one condition, this bad outcome was due to a leak of expensive chemicals. In another condition, the bad outcome was due to a leak of cheap chemicals that caused harm to residents living around the factory.

After reading the scenario, participants in each condition answered two questions. The first one asked them to indicate how ethical they thought the CEO was on a 7-point scale (ranging from 1 = unethical to 7 = ethical). The second asked them to indicate whether they believed that the CEO deserved to be punished again on a 7-point scale (ranging from 1 = No, to 7 = Yes).



## Results and Discussion

Outcome information influenced ethicality ratings,  $F(2, 137) = 16.64, p < .001, \eta^2 = .20$ . Ethicality ratings were higher in the positive-outcome ( $M = 4.04, SD = 1.71$ ) than in both the negative/financial-outcome condition ( $M = 2.93, SD = 1.47; t[91] = 3.35, p = .001$ ) and the negative/unethical-outcome condition ( $M = 2.28, SD = 1.30; t[92] = 5.65, p < .001$ ). In addition, ethicality ratings were higher in the negative/financial-outcome condition than in the negative/unethical-outcome condition ( $t[91] = 2.29, p = .024$ ).

Similarly, the desire to punish was affected by outcome information,  $F(2, 137) = 27.49, p < .001, \eta^2 = .29$ . It was stronger in the negative/unethical-outcome ( $M = 5.83, SD = 1.57$ ) than in either the negative/financial-outcome ( $M = 4.65, SD = 1.68; t[91] = 3.50, p = .001$ ) or positive-outcome conditions ( $M = 3.28, SD = 1.77; t[92] = 7.42, p < .001$ ). Furthermore, the desire to punish was stronger in the negative/financial-outcome condition than in the positive-outcome condition ( $t[91] = 3.85, p < .001$ ).<sup>1</sup>

Taken together these results indicate that participants evaluated the behavior as more positively and less deserving of punishment when its consequences were negative but without strong ethical implications than when they were negative and with strong ethical implications. This suggests that ethical judgments may be particularly vulnerable to the influence of the outcome bias.

### Study 5

In Study 5, we start examining the psychological process that underlies the outcome bias in individuals' judgments of ethicality. Is the outcome bias in ethical judgments the result of a judge's gut feelings of right and wrong or the result of rational deliberation? Individuals might

---

<sup>1</sup> Consistent with the results of Study 1, we also find that judgments of ethicality mediate the relationship between outcome information and judgments of punishment.

use a rational strategy and explicitly judge behaviors that resulted in a negative outcome as more unethical than behaviors that resulted in a positive outcome. Or they might use their own intuitions, which might conflict with a rational deliberation. Hypothesis 3 predicted that the differential impact of outcome information on ethical judgments would occur only in one mode of thought, the intuitive one, but not in the rational mode of thought.

To explore this hypothesis, we designed a study which used a paradigm developed by Epstein and his colleagues (Epstein et al., 1992) and employed by other researchers as well (e.g., Pizarro, Uhlmann, & Bloom, 2003). By employing this paradigm, we induced participants to respond to outcome information in a rational manner, and then we compared these responses to those of participants who were asked to respond in an intuitive manner. Following Epstein et al. (1992) as well as others (e.g., Pizarro et al., 2003), we also altered the order in which participants received intuitive versus rational instructions. We expected that when first placed in a rational mindset, participants would correct their intuitions and would not judge behaviors as more unethical when they lead to a negative outcome. However, when first placed in an intuitive mindset, participants would continue to associate negative outcomes with unethical behavior, even at the second stage, when they are asked to respond rationally. In this case, individuals commit themselves to their intuitive judgments and make efforts to rationalize their intuitions as argued by prior research (Haidt, 2001; Epstein, 1994; Epstein et al., 1992; Pizarro et al., 2003).

### Method

*Participants.* Eighty-five individuals (52 female, 33 male) participated in the study for monetary compensation. Each received \$6 for their participation. Most participants (68% of them) were students from local universities. The average age of participants was 28 ( $SD = 12.22$ ).

*Design and Procedure.* The study employed a 2 (within participants: rational instructions vs. intuitive instructions) X 2 (between participants: rational instructions first vs. intuitive instructions first) design. Participants in the study received two sets of scenarios, each containing two vignettes. Each of the two sets contained a vignette describing the behavior of a decision maker which led to negative consequences (version 1) and, on the same page, a vignette describing the behavior of a similar decision maker which led to positive consequences (version 2). The vignettes used were identical to the vignettes describing the behavior of a pharmaceutical researcher and the behavior of an auditor in the positive and negative-outcome conditions used in Study 1.

The order in which participants read both versions of the vignettes was randomized (half of the participants received the negative outcome vignette first, and the other half received the positive outcome vignette first in each of the two sets of scenarios). As an example, a participant might have received two vignettes describing the behavior of a pharmaceutical researcher (in one case such behavior led to positive consequences [positive-outcome vignette] and in another case it led to negative consequences [negative-outcome vignette]) in set 1, and then two vignettes describing the behavior of an auditor (again, in one case such behavior led to positive consequences and in another case it led to negative consequences) in set 2. In order to easily distinguish between the pharmaceutical researcher (or auditor) described in the positive-outcome vignette and the pharmaceutical researcher (or auditor) described in the negative-outcome vignette, we referred to the decision maker using a gender-neutral first name. So, for instance, the first sentence of the positive-outcome vignette for the pharmaceutical researcher read,

A pharmaceutical researcher, Jamie, defines a clear protocol for determining whether or not to include clinical patients as data points in a study.

Instead, the first sentence of the positive-outcome vignette for the pharmaceutical researcher read,

A pharmaceutical researcher, Terry, defines a clear protocol for determining whether or not to include clinical patients as data points in a study.

After reading the vignettes, participants were asked to compare the actions of the person described in the pair of vignettes (either the pharmaceutical researcher or the auditor) in the case of positive-outcome information and in the case of negative-outcome information. A first question asked them to indicate the extent to which the described behavior was unethical using a 7-point scale ranging from 1 [Jamie/Chris's behavior is more unethical than Terry/Pat's behavior] to 7 [Terry/Pat's behavior is more unethical than Jamie/Chris's behavior], with the mid-point being 4 [Jamie/Chris's behavior and Terry/Pat's behavior are equally unethical].<sup>2</sup> The second question asked participants to indicate how harshly they would punish the described behavior using a similar 7-point scale ranging from 1 [Jamie/Chris's behavior deserves to be punished more harshly than Terry/Pat's behavior] to 7 [Terry/Pat's behavior deserves to be punished more harshly than Jamie/Chris's behavior], with the mid-point being 4 [Jamie/Chris's behavior and Terry/Pat's behavior deserve to be punished the same way]. The third question asked participants to indicate the extent to which they blamed the described behavior for the outcome using a 7-point scale ranging from 1 [Jamie/Chris's behavior deserves much more blame than Terry/Pat's behavior] to 7 [Terry/Pat's behavior deserves much more blame than Jamie/Chris's behavior], with the mid-point being 4 [Jamie/Chris's behavior and Terry/Pat's behavior deserve equal blame]. Finally, a fourth question asked them to indicate the extent to which the person described was a bad person using again a 7-point scale ranging from 1

---

<sup>2</sup> The names Chris and Pat were used in the vignettes describing the behavior of an auditor.

[Jamie/Chris is a worse individual than Terry/Pat] to 7 [Terry/Pat is a worse individual than Jamie/Chris], with the mid-point being 4 [Jamie/Chris and Terry/Pat are equally bad people].

Each participant was asked to answer this set of questions twice: the first time after reading the first set of vignettes (e.g., the two vignettes describing the behavior of a pharmaceutical researcher); the second time after reading the second set of vignettes (the two vignettes describing the behavior of an auditor).

Depending on the conditions they had been randomly assigned to (rational instructions first vs. intuitive instructions first), participants received different instructions on how to approach their judgments. In particular, participants were asked to make these judgments from either an *intuitive perspective* (i.e., “my intuitive, gut feeling is that...”), or a *deliberative perspective* (i.e., “my most rational, objective judgment is that...”). Half of the participants were instructed to respond in a rational manner first (rational instructions first condition), and the other half were instructed to respond intuitively first (intuitive instructions first condition), in order for us to test the effects of priming an intuitive mindset first (Epstein et al., 1992) on ethical judgments. In other words, participants in the rational instructions first condition were asked to respond rationally to the questions reported after the first set of vignettes, and to respond intuitively to the questions reported after the second set of vignettes. Instead, participants in the intuitive instructions first condition were asked to respond intuitively to the questions reported after the first set of vignettes, and to respond rationally to the questions reported after the second set of vignettes.

## Results

Hypothesis 3 predicted that mindset instructions would influence ethical judgments. To test this hypothesis, we used the unethicity ratings as the dependent variable in a 2

(instructions: rational vs. intuitive) X 2 (order of instructions: rational instructions first vs. intuitive instructions first) mixed-design ANOVA. This analysis revealed an insignificant main effect for order of instructions,  $F(1, 83) < 1, p = .49, \eta^2 = .006$ . The main effect of the type of instructions participants received (rational vs. intuitive) was significant,  $F(1, 83) = 10.62, p = .002, \eta^2 = .11$ . Furthermore, as predicted, this analysis revealed a significant interaction between order of instructions and type of instructions,  $F(1, 83) = 4.05, p < .05, \eta^2 = .05$ , depicted in Figure 5.

In order to identify whether this interaction was in the expected direction, we conducted one-sample  $t$  tests on judgments of unethicality. Judgments were compared to a value of 4, which indicated an identical judgment of unethicality between behaviors that led to a positive outcome and behaviors that led to a negative outcome (i.e., no outcome bias in ethical judgments). Values greater than 4 indicated that individuals were influenced by outcome information in making their ethical judgments, and judged the behavior of the decision maker described in the negative-outcome vignette as more unethical than the behavior of the decision maker described in the positive-outcome vignette. As predicted, when participants first received the rational instructions and were asked to respond rationally, they were not influenced by the outcome bias ( $M = 4.02, SD = 0.77$ ),  $t(42) < 1, p = .84$ . However, when subsequently asked to answer intuitively, their judgment showed the outcome bias ( $M = 4.93, SD = 1.61$ ),  $t(42) = 3.79, p < .001$ . In contrast, when participants first received the intuitive instructions and were asked to respond intuitively, their judgments were biased by outcome information ( $M = 4.52, SD = 1.27$ ),  $t(41) = 2.67, p = .011$ , and continued to be biased even when they were subsequently asked to respond rationally ( $M = 4.74, SD = 1.38$ ),  $t(41) = 3.47, p = .001$ . These results provide support for Hypothesis 3 which predicted that the outcome information would bias ethical judgments when individuals

form such judgments intuitively but not when they form them rationally. More specifically, the results suggest that individuals abandon their intuitions to judge the ethicality of others' behavior when asked to respond rationally, as long as the rational instructions were received first. Indeed, as our results show, this is the only condition in which participants' ethical judgments were not influenced by the outcome bias.

Note that the nature and significance of the results do not change if any of the other variables measured in the study is used as the dependent variable in the analyses presented above. Similarly, if the dependent variables assessing the degree to which individuals differentially judged behaviors which led to positive versus negative consequences are combined, then again the nature and significance of the results do not change.

### Discussion

The results of Study 5 provide further evidence for the robustness of the outcome bias in ethical domains: individuals judge behaviors that produced a negative outcome as more unethical than behaviors that produced a positive outcome. Furthermore, the results demonstrate that a rational, analytic mindset can override the effects of one's intuitions in ethical judgments, as long as a rational deliberation is required first.

### Study 6

While Study 5 provides initial evidence for the moderating effect of one's mode of thoughts, it does so by asking participants to make comparative judgments. Given that most real-world judgments are non-comparative, we conducted a sixth study in which participants were asked to make judgments of ethicality using only a specific mode of thought and without comparing behaviors leading to different outcomes.

### Method

*Participants and Procedure.* One-hundred forty two individuals (54% male) participated in the study in exchange of \$2. Participants were randomly assigned to one of four conditions: positive-outcome / rational mindset condition, negative-outcome / rational mindset condition, positive-outcome / intuitive mindset condition or negative-outcome / intuitive mindset condition. Thus, the study employed a 2 (outcome information: positive vs. negative) X 2 (mindset: rational vs. intuitive) between-subjects design. In each condition, participants read three scenarios. We used the same scenarios as in Study 2.

In the rational mindset conditions, participants were asked to rate the ethicality of the described behavior using their rational judgment:

My most rational, objective judgment is that, based on the information available at the time of the decision, the manager's [commissioner's, mayor's] decision not to build the back-up system [to provide only tents, against investing in water conservation] was...

In each scenario, participants indicated their judgments on a 7-point scale, ranging from 1 = unethical, to 7 = ethical. In the intuitive mindset conditions, participants were asked to rate the ethicality of the described behavior using their intuitive judgment:

My intuitive, gut feeling is that, based on the information available at the time of the decision, the manager's [commissioner's, mayor's] decision not to build the back-up system [to provide only tents, against investing in water conservation] was...

Also in this case, participants indicated their judgments on a 7-point scale, ranging from 1 = unethical, to 7 = ethical.

## Results

We used the ethicality ratings participants provided in a 2 (outcome information) X 2 (mindset) between-subjects repeated measure ANOVA (repeated measure on scenarios). Outcome information influenced ethicality ratings,  $F(1, 137) = 38.28, p < .001, \eta^2 = .22$ . Ethicality ratings were higher in the positive-outcome ( $M = 3.74, SD = 1.20$ ) than in the



negative-outcome condition ( $M = 2.49, SD = 1.19$ ). Replicating our previous findings, these results indicate that participants evaluated the same described behavior differently depending on the outcome information we provided. Specifically, they evaluated the behavior as more ethical when its consequences were positive than when they were negative.

The main effect of mindset was insignificant,  $F(1, 137) = 1.51, p = .22, \eta^2 = .01$ . This analysis also revealed a significant outcome information X mindset interaction,  $F(1, 137) = 4.31, p = .04, \eta^2 = .03$ , depicted in Figure 6.

When the outcome information was negative, a rational mindset led people to evaluate described decisions as less unethical (i.e., more ethical) than an intuitive mindset (2.83 vs. 2.17,  $F[1, 70] = 5.89, p = .018, \eta^2 = .08$ ). When the outcome information was positive, there was no significant difference in the ethicality ratings participants provided when approaching their judgments with a rational or intuitive mindset (3.66 vs. 3.81,  $F[1, 67] < 1, p = .57, \eta^2 = .005$ ). Furthermore, differences in the effect of outcome information were significant both in the rational mindset condition ( $M_{positive\ outcome} = 3.66$  vs.  $M_{negative\ outcome} = 2.83, F[1, 66] = 8.74, p = .004, \eta^2 = .12$ ) and in the intuitive mindset condition ( $M_{positive\ outcome} = 3.81$  vs.  $M_{negative\ outcome} = 2.17, F[1, 71] = 33.63, p < .001, \eta^2 = .32$ ).

## Discussion

The results of Study 6 are consistent with the findings of Study 5, and provide further support for Hypothesis 3, which predicted that outcome information would bias ethical judgments significantly more when the observer is using an intuitive mindset than when the observer is using a rational, deliberative mindset.

## General Discussion

The present studies provide strong evidence of the existence of outcome effects in ethically-relevant contexts, when people are asked to judge the ethicality of others' behavior. It is worth noting that what we show is not the same as the curse of knowledge or the hindsight bias. The curse of knowledge describes people's inability to recover an uninformed state of mind (Camerer, Loewenstein, & Weber, 1989). Likewise, the hindsight bias leads people to misremember what they believed before they knew an event's outcome (e.g., Fischhoff, 1975; Fischhoff & Beyth, 1975). By contrast, we show that that outcomes of decisions lead people to see the decisions themselves in a different light, and that this effect does not depend on misremembering their prior state of mind. In other words, people will see it as entirely appropriate to allow a decision's outcome to determine their assessment of the decision's quality.

Furthermore, the results of our studies also speak to the judgmental process behind the outcome bias in ethical judgments. The results of our last two studies are consistent with the notion that the outcome bias in ethical domains is driven by an intuitive impulse, and not by a rational or logical deliberation.

#### *Contributions to Theory and Research*

These findings extend prior research in psychology on outcome effects and outcome severity which represent a source of irrationality in social judgment (Mazzocco et al., 2004). Our third study showed that participants revised their perception of the ethicality of the actor's behavior once they learned about the outcome such behavior led to, even when they learned nothing about the decision maker or what he knew. This tendency seems to contradict the rational prescriptions proposed by decision analysts.

Our findings also extend prior research in the ethics literature. Several models have been proposed to explain the decision-making process people use in situations involving a perceived

ethical problem (e.g., Ferrell et al., 1989, Hunt & Vitelli, 1986; Jones, 1991; Knobe, 2006; Trevino, 1986). A common feature across these models is the inclusion of moral philosophy frameworks and theories which are used by individuals to judge the ethicality of an actor's actions. According to some of these theories, known as teleological theories, the goodness or badness of behavior is judged based on its consequences. Utilitarianism, for instance, holds that an action is right if it produces (or tends to produce) the greatest amount of good for the greatest number of people affected by the action (DeGeorge, 1999). Actions are neither good nor bad: their "nature" is judged based on their consequences.

Our research presents a challenge for the utilitarian reasoning. If an individual's choice produces a positive outcome due purely to chance, should the actor therefore be praised? Is it reasonable to encourage or reward behavior that resulted in favorable outcomes, not because the actor willed that outcome but thanks to good fortune? This is the problem that philosophers have labeled as "moral luck" (Nagel, 1979; Williams, 1981). Actions which produced negative outcomes are perceived as more unethical than similar actions which produced positive outcomes, even in cases in which fortune was the primary cause behind those outcomes.

Our findings also present a challenge for the administration of legal punishments. While there are legitimate motives for punishment that are driven by outcomes, our results suggest that motives to punish may be magnified by ethical evaluations that are, in turn, contaminated by outcomes. In other words, those meting out punishment run the risk of punishing those responsible for bad outcomes too harshly because of illogical inferences about what the outcomes reflect about the ethicality of their choices.

*Practical Implications*

The research presented here also has practical implications. The tendency demonstrated in our studies might lead people to blame others too harshly for making sensible decisions that have unlucky outcomes. While we recognize the difference between simple scenarios like the ones used in the present studies and complex real-world decisions, the effects observed in our studies could help explain the slow reactions people tend to have when they observe others' unethical behavior. Too often, we let ethically-questionable decisions slide for a long time until they result in negative outcomes, even in cases in which such outcomes are easily predictable (Bazerman & Watkins, 2004).

The case of auditor independence provides a good example. For many years, auditing firms worked for their clients not only as auditors but also as consultants, and engaged in other activities that compromised the independence of their audits (Frankel, Johnson, & Nelson, 2002; Moore, Tetlock, Tanlu, & Bazerman, 2006). A large body of evidence has pointed to the failure of auditor independence (e.g., Levitt & Dwyer, 2002), despite the fact that independence is core to auditing (Berardino, 2000; Burger, 1984). Yet, the U.S. government did not take any action to bolster auditor independence until auditor conflicts of interest led to important failures of big firms such as Enron, WorldCom, and Tyco (Moore et al., 2006). Long before negative outcomes such as bankruptcy and loss of jobs for many employees, ample evidence was available that the existing structure compromised the ethics of the auditing profession (Bazerman, Morgan, & Loewenstein, 1997). In retrospect, only negative outcomes seemed to have motivated US legislative representatives to address the failure of auditor independence.

*What to Do About It*

Our work has implications also for organizations and public institutions. As our studies show, decision-makers should anticipate being judged not so much on the ethics of their actions but on the consequences of those actions. However ethical the decision of a manager or company is, judges (such as customers, citizens or employees) might punish the manager or company if things go wrong.

Our research could inform potential policies changes within organizations. Evaluation of others' ethical behavior is increasingly becoming an important part of performance appraisals of employees within organizations, as recommended by prior work (Buckley, 2001; Weaver, 2001; Weaver & Trevino, 1999). If ethical judgments are biased by systematic errors, then performance evaluations might wrongly reward those employees who achieved good work outcomes but did so through ethically questionable practices. These implications extend to domains outside organizational settings, such as performance evaluation of athletes or decisions to fund companies or projects which have been successful in the past.

These findings suggest that managers within organizations – or more generally people in the position of judges – would benefit from asking questions about the decision makers' behavior before outcomes are realized. Such information might be difficult to gather but it is fundamental in assuring a fair reward/punishment system within organizations and society more broadly. Our findings also speak to the role of questioning cultures which support praises to the winners without precisely evaluating the means which led to successful ends. Too often, workers are evaluated based on results and not based on the quality of the decision.

These examples highlight the importance of the study of outcome effects in ethically-relevant contexts. Future research in this area is warranted. Future studies could investigate

further the psychological mechanisms explaining the outcome bias in the ethics realm, as well as its boundary conditions. We demonstrated that adopting a rational mindset can override the intuitive, gut feelings of right and wrong that drive ethical judgments. Further work is needed to uncover factors hiding behind these intuitions. For instance, observers might feel strongly motivated to engage in a post-hoc sense making process to justify a negative outcome. These are questions that warrant further investigation.

For now, based on our findings we suggest that it is worth trying to understand a decision maker's state of mind. Sometimes bad things happen when good people are unlucky, and sometimes scoundrels get away clean. Judging decisions based on their outcomes will wind up condemning too many unlucky people and acquitting too many scoundrels.

## Appendix A: Scenario used in Study 1

<i>Negative-outcome condition</i>	<i>Positive-outcome condition</i>
<p><u>Pharmaceutical researcher</u></p> <p>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 within his firm. As the deadline approaches, he notices that four subjects were withdrawn from the analysis due to technicalities. He believes that the data in fact is 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.</p>	<p>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 within his firm. He believed that the product was 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.</p>
<p><u>Auditor</u></p> <p>An auditor is examining the books of an important client, a client that is not only valuable for their auditing fees, but also buys lucrative advisory services from the auditor's firm as well. The auditor notices some accounting practices that are probably illegal, but it would take multiple court cases to be sure about whether the action was legal or not. The auditor brings up the issue with the client, who insists that there is nothing wrong with their accounting. The client also threatens to withdraw their business if the auditor withholds their approval. The auditor agrees to let it go by for one year, and encourages the client to change their accounting practices over the next year. Six months later, it is found that the client was committing fraud, their corporation goes bankrupt, the bankruptcy is connected to the issue that the auditor noticed, and 1,400 people lose their jobs and their life's savings.</p>	<p>An auditor is examining the books of an important client, a client that is not only valuable for their auditing fees, but also buys lucrative advisory services from the auditor's firm as well. The auditor notices clearly fraudulent practices by their client. The auditor brings up the issue with the client, who insists that there is nothing wrong with their accounting. The client also threatens to withdraw their business if the auditor withholds their approval. The auditor agrees to let it go by for one year, and encourages the client to change their accounting practices over the next year. No problems result from the auditor's decision.</p>
<p><u>Toy company</u></p> <p>A toy company finds out that the products that they were selling, manufactured by another firm in another country, contains lead, which can be extremely hazardous to children. The toy company had failed to test for lead in the product, since testing is expensive and is not required by U.S. law. The lead paint eventually kills 6 children, and sends dozens more to emergency room for painful treatment for lead poisoning.</p>	<p>A toy company sells products made by another firm, manufactured in another country. The toy company knows that the toys contain lead, which can be extremely hazardous to children. The toy company successfully sells this product, makes a significant product, and no children are injured by the lead paint.</p>

## Appendix B: Scenario used in Study 2

<i>Negative-outcome condition</i>	<i>Positive-outcome condition</i>
<p><u>Sewage treatment plant</u></p> <p>A sewage treatment plant is undergoing remodeling and updating. There is a critical phase to the project during which all the treatment systems are shut off and incoming sewage is diverted to a holding tank until the new systems are activated. This critical phase of the project will last 48 hours. If there should be substantial rainfall during this 48-hour period, there is a high probability that the holding tanks will overflow into local waterways, with serious negative environmental and health consequences for the wildlife and people who live in the area. The company that runs the sewage treatment plant could invest in back-up systems that would eliminate the risk of overflow even in the case of heavy rain, but this would be expensive. Historically, the chances of rain during the planned 48 hour period are about 5%. After extensive consideration, the manager in charge of the remodeling decides against instituting the back-up plan. It winds up raining a great deal. Twenty people fall ill and wind up in the hospital, many fish die, the water is unsafe for swimming for a week, and fishers are discouraged from eating fish caught downstream.</p>	<p>A sewage treatment plant is undergoing remodeling and updating. There is a critical phase to the project during which all the treatment systems are shut off and incoming sewage is diverted to a holding tank until the new systems are activated. This critical phase of the project will last 48 hours. If there should be substantial rainfall during this 48-hour period, there is a high probability that the holding tanks will overflow into local waterways, with serious negative environmental and health consequences for the wildlife and people who live in the area. The company that runs the sewage treatment plant could invest in back-up systems that would eliminate the risk of overflow even in the case of heavy rain, but this would be expensive. Historically, the chances of rain during the planned 48 hour period are about 10%. After extensive consideration, the manager in charge of the remodeling decides against instituting the back-up plan. In fact, there is no rain during the critical 48-hour period and the plant remodeling is a complete success.</p>
<p><u>Natural disaster</u></p> <p>A government agency in a developing country finds itself dealing with a natural disaster in which several thousand poor peasants have been made homeless during the winter. The agency must decide what sort of short-term housing it will provide for the refugees. The inexpensive option is tents, which will probably be fine, given the mildness of the local winters—overnight temperatures only fall below freezing once every four years or so, on average. The more expensive option is to put up temporary shacks that would provide more shelter against the cold. But they shacks would be more expensive and would force the agency to cut funding to other (less urgent) programs. In the end, the agency's commissioner decides to provide only tents for the refugees. The winter is substantially colder than expected, and fifty children among the refugees die of exposure to the cold.</p>	<p>A government agency in a developing country finds itself dealing with a natural disaster in which several thousand poor peasants have been made homeless during the winter. The agency must decide what sort of short-term housing it will provide for the refugees. The inexpensive option is tents, which will probably be fine, given the mildness of the local winters—overnight temperatures only fall below freezing once every two years or so, on average. The more expensive option is to put up temporary shacks that would provide more shelter against the cold. But they shacks would be more expensive and would force the agency to cut funding to other (less urgent) programs. In the end, the agency's commissioner decides to provide only tents for the refugees. The winter is quite mild, and the tents provide sufficient shelter.</p>
<p><u>Water supply</u></p> <p>There is a river that runs through dry regions in Mexico. Ninety percent of the time, there has been plenty of water to supply the communities that depend on water from the river. So when the mayor of a prosperous town that drew its water from the river was asked to invest in water-conservation measures, she assumed they were unnecessary. It so happened that this town was upstream from most of the other communities that depended on the river. Upstream communities naturally have the advantage because they can take what they need first, and downstream communities left without water have little recourse. One year, rainfall is far below expectations and 46 small farms are driven out of business when the lack of water in the river leaves them unable to irrigate their land.</p>	<p>There is a river that runs through dry regions in Mexico. Eighty percent of the time, there has been plenty of water to supply the communities that depend on water from the river. So when the mayor of a prosperous town that drew its water from the river was asked to invest in water-conservation measures, she assumed they were unnecessary. It so happened that this town was upstream from most of the other communities that depended on the river. Upstream communities naturally have the advantage because they can take what they need first, and downstream communities left without water have little recourse. In the end, there was plenty of rain and more than enough water for the communities along the river.</p>



## Appendix C: Scenario used in Study 4

ChemTech is a multinational corporation that manufactures various synthetic chemicals. The firm's products include fertilizers, insecticides, and alternative fuels. The inputs to these products include many chemicals that are expensive to produce. Some are harmless and non-toxic, like hydrogen. Others are highly toxic and dangerous, such as methyl isocyanate.

ChemTech's CEO decided to build a new manufacturing facility in India. The biggest reason favoring the CEO's decision to locate in India was that the regulations surrounding construction specifications and building codes were considerably more flexible in India than in more developed economies. This allowed ChemTech to build its factory without the need for triple- and quadruple-backup systems to prevent leaks and accidental discharges of chemicals from storage tanks. This made sense because the biggest threat to ChemTech's storage systems was fluctuations in temperature faced at facilities in Europe and North America. But the location of the Indian factory featured a moderate climate where the temperature hardly ever fluctuated much.

<i>Ethically-Relevant Outcome</i>	<i>Financially-Relevant Outcome</i>
<p>ChemTech's facility in India had been operating for 8 months when disaster struck. On the night of January 17th, 2003, the temperature dropped substantially below the lowest temperature that had ever been recorded in the region during the 70 years for which weather records had been kept. One of the storage tanks for gaseous chemicals sprang a leak. Over a period of a few hours, tons and tons of hydrogen escaped. The chemical leak was prodigiously expensive for ChemTech. The lost hydrogen wound up costing more than ChemTech had saved by building the factory in India.</p>	<p>ChemTech's facility in India had been operating for 8 months when disaster struck. On the night of January 17th, 2003, the temperature dropped substantially below the lowest temperature that had ever been recorded in the region during the 70 years for which weather records had been kept. One of the storage tanks for gaseous chemicals sprang a leak. Over a period of a few hours, tons and tons of methyl isocyanate escaped. The chemical leak was prodigiously expensive for ChemTech. The restitution paid to families of those killed by the methyl isocyanate wound up costing more than ChemTech had saved by building the factory in India.</p>

## References

- Alicke, M. D., & Davis, T. L. (1989). The role of a posteriori victim information in judgments of blame and sanction. *Journal of Experimental Social Psychology, 25*, 362-377.
- 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.
- American Law Institute. (1962). *Model penal code: Official draft and explanatory notes*. Philadelphia: Author.
- Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology, 54*, 569-579.
- Baron, R. M., & Kenny, D. A. (1986). The moderator-mediator variable distinction in social psychological research: Conceptual, strategic and statistical considerations. *Journal of Personality and Social Psychology, 51*, 1173-1182.
- Bazerman, M. H., Morgan, K. P., & Loewenstein, G. (1997). The impossibility of auditor independence. *Sloan Management Review, 38*(4), 89-94.
- Bazerman, M. H., & Watkins, M. D. (2004). *The Disasters You Should Have Seen Coming and How to Prevent Them*. Harvard Business School Publishing.
- Berardino, J. F. (2000). Hearing on Auditor Independence, from <http://www.sec.gov/rules/proposed/s71300/testimony/berardi2.htm>
- Berg-Cross, L. G. (1975). Intentionality, degree of damage, and moral judgments. *Child Development, 46* (4), 970-974.
- Bok, S. (1979). *Lying: Moral choice in public and private life*. Vintage Books.

- Buckley, R. M. (2001). Ethical issues in human resource systems. *Human Resources Management Review*, 11, 1/2.
- Burger, W. United States v. Arthur Young & Co., 82-687 (Supreme Court of the United States 1984).
- Camerer, C., Loewenstein, G., & Weber, M. (1989). The curse of knowledge in economic settings: An experimental analysis. *The Journal of Political Economy*, 97(5), 1232-1254.
- Chaiken, S. (1980). Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology*, 39, 752-766.
- Creyer, E. H., & Ross Jr., W.T. (1997). The influence of firm behavior on purchase intention: do consumers really care about business ethics? *Journal of Consumer Marketing*, 14, 421-432.
- DeGeorge, R. T. (1999). *Business ethics*. Prentice-Hall, Upper Saddle River, NJ.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, 56, 5-18.
- Epstein, S. (1994). Integration of the cognitive and the psychodynamic unconscious. *American Psychologist*, 4, 709-724.
- Epstein, S., Lipson, A., Holstein, C., & Huh, E. (1992). Irrational reactions to negative outcomes: Evidence for two conceptual systems. *Journal of Personality and Social Psychology*, 62, 328-339.
- Epstein, S., & Pacini, R. (1999). Some basic issues regarding dual-process theories from the perspective of cognitive-experiential self theory. In S. Chaiken and Y. Trope (Eds.) *Dual-process theories in social psychology*. New York: The Guildford Press.

- Ferrell, O. C., Gresham, L. G., & Fraedrich, J. (1989). A synthesis of ethical decision models for marketing. *Journal of Macromarketing*, 9 (Fall), 55-64.
- Fischhoff, B. (1975). Hindsight≠foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1, 288-299.
- Fischhoff, B., & Beyth, R. (1975). "I knew it would happen." Remembered probabilities of once-future things. *Organizational Behavior and Human Performance*, 13, 1-16.
- Fischer, J. M. (1986). *Moral responsibility*. Ithaca, NY: Cornell University Press.
- 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.
- Goldberg, J. H., Lerner, J. S., & Tetlock, P. E. (1999). Rage and reason: The psychology of the intuitive prosecutor. *European Journal of Social Psychology*, 29, 781-795.
- Haidt, J. (2007). The new synthesis in moral psychology. *Science*, 316(5827), 998-1002.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108, 814–834.
- Haidt, J., & Joseph, C. (2004). Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus*, Fall, 55-66.
- Hastie, R., & Dawes, R. M. (2001). *Rational choice in an uncertain world: The psychology of judgment and decision making*. Thousand Oaks, Calif.: Sage.
- Hershey, J., & Baron, J. (1992). Judgment by outcome: When is it justified? *Organizational Behavior and Human Decision Processes*, 53, 89-93.

- Hunt, S. D., & Vitell, S. (1986). A general theory of marketing ethics. *Journal of Macromarketing*, 6 (Spring), 5-16.
- Jones, T. M. (1991). Ethical decision making by individuals in organizations: An issue contingent model. *Academy of Management Review*, 16 (April), 366-395.
- Kadish, S. H., Schulhofer S. J., Steiker C. (2007). *Criminal law and its processes: Cases and materials*. Wolters Kluwer Law & Business.
- Kahneman, D., & Tversky, A. (1982). The simulation heuristic, In D. Kahneman, P. Slovic and A. Tversky (eds.) *Judgement under uncertainty: Heuristics and biases* (pp. 201-208). New York: Cambridge University Press.
- Kaplow, L., & Shavell, S. (2002). *Fairness versus welfare*. Harvard University Press.
- Knobe, J. (2006). The concept of intentional action: A case study in the uses of folk psychology. *Philosophical Studies*, 130, 2, 203-231.
- Leon, M. (1982). Rules in children's moral judgments: Integration of intent, damage, and rationale information. *Developmental Psychology*, 18, 833-840.
- Levitt, A., & Dwyer, P. (2002). *Take on the Street*. Pantheon, New York, NY.
- Lipshitz, R. (1989). Either a medal or a corporal: The effects of success and failure on the evaluation of decision making and decision makers. *Organizational Behavior and Human decision Processes*, 44, 380-395.
- Lowe, C. A., & Medway, F. J. (1976). Effects of valence, severity, and relevance on responsibility and dispositional attribution. *Journal of Personality*, 44, 518 - 539.
- Marshall, G. W, & Mowen, J. C (1993). An experimental investigation of the outcome bias in salesperson performance evaluations. *Journal of Personal Selling and Sales Management*, 8, 31-47.

- Mazzocco, P. J., Alicke, M. D., & Davis, T. L. (2004). On the robustness of outcome bias: No constraint by prior culpability. *Basic and Applied Social Psychology*, 26, 131-146.
- McKillip, J., & Posavac, E. J. (1975). Judgments of responsibility for an accident. *Journal of Personality*, 43, 248-265.
- Mitchell, T. R., & Kalb, L. S. (1981). Effects of outcome knowledge and outcome valence on supervisors' evaluations. *Journal of Applied Psychology*, 66, 604-612.
- Moore, D., Tetlock, P. E., Tanlu, L., & Bazerman, M. (2006). Conflicts of interest and the case of auditor independence: Moral Seduction and Strategic Issue Cycling. *Academy of Management Review*, 31, 10-29.
- Morris, H. (1961). *Freedom and responsibility*. Stanford, CA: Stanford University Press.
- Nagel, T. (1979). *Mortal questions*. New York: Cambridge University press.
- North, J. (1987). Wrongdoing and forgiveness. *Philosophy*, 62 499-508.
- Pizarro, D. A., Uhlmann, E., & Bloom, P. (2003). Causal deviance and the attribution of moral responsibility. *Journal of Experimental Social Psychology*, 39, 653-660.
- Rosenzweig, P. (2007). *The halo effect*. New York: Free Press.
- Schroeder, D. A., & Linder, D. E. (1976). Effects of actor's causal role, outcome severity, and knowledge of prior accidents upon attributions of responsibility. *Journal of Experimental Social Psychology*, 12, 340-356.
- Shaw, J. I., & Skolnick, P. (1971). Attributions for a happy accident. *Journal of Personality and Social Psychology*, 18, 380-383.
- Stanovich, K. E. (1999). Who is rational? Studies of individual differences in reasoning, Elrbaum.

- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate. *Behavioral and Brain Sciences*, 23 (5), 645–726.
- Stokes, J. M., & Leary, M. R. (1984). Evaluation of others' decisions by intellectually gifted and average children: Effects of decision consequences and decentering prompts. *Personality and Social Psychology Bulletin*, 10, 564-573.
- Surber, C. F. (1977). Developmental processes in social inference: Averaging of intentions and consequences in moral judgments. *Developmental Psychology*, 13, 654-665.
- Trevino, L. K. (1986). Ethical decision making in organizations: A person-situation interactionist model. *Academy of Management Review*, 11 (3), 601-617.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1131.
- Ugwuegbu, D. C., & Hendrick, C. (1974). Personal causality and attribution of responsibility. *Social Behavior and Personality*, 2 (1), 76-86.
- Walster, E. (1966). Assignment of responsibility for an accident. *Journal of Personality and Social Psychology*, 3, 13-19.
- Weaver, G. R. (2001). The role of human resources in ethics/compliance management: a fairness perspective. *Human Resource Management Review*, 11, 1/2, 113-135.
- Weaver, G. R., & Trevino, L. K. (1999). Integrated and decoupled corporate social performance: management commitments, external pressures, and corporate ethics practices. *Academy of Management Journal*, 42(5), 539-552.
- Williams, B. (1981). *Moral luck*. Cambridge: Cambridge University Press.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, 107, 101–126.

## Tables

Table 1

*Descriptive statistics and correlations of the main variables used in Study 1*

	<i>Mean</i>	<i>SD</i>	<i>1</i>	<i>2</i>
1. Unethicality	5.53	1.00		
2. Harshness in punishing	5.25	1.04	.72***	
3. Blame	5.09	1.24	.47***	.57***

\*\*\*  $p < .001$



## Figures Captions

*Figure 1.* Mediation analysis of judgments of punishment, Study 1. Note that \*  $p < .05$  \*\*  $p < .01$  \*\*\*  $p < .001$

*Figure 2.* Mediation analysis of judgments of blame, Study 1.

*Figure 3.* Mean rating for unethicity for both Phase 1 and Phase 3, by scenario (Study 2). Error bars represent standard errors.

*Figure 4.* Mean rating for unethicity for both Phase 1 and Phase 3 (Study 3). Error bars represent standard errors.

*Figure 5.* Mean rating for unethicity by condition (Study 4). Error bars represent standard errors.

*Figure 6.* Mean ratings for ethicality (1-7 scale) by condition (Study 5). Error bars represent standard errors.

Figures

Figure 1.




Figure 2.




Figure 3.

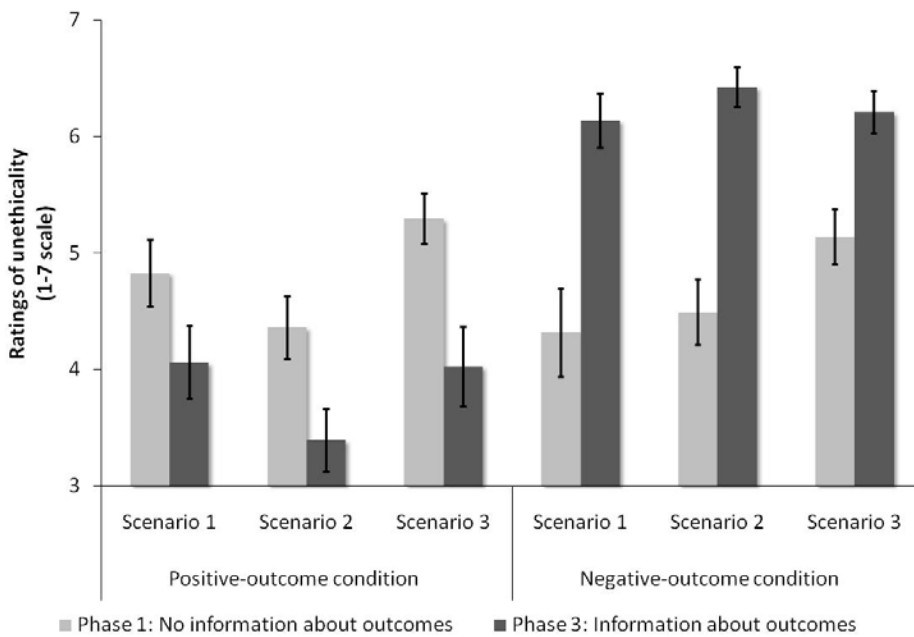


Figure 4.

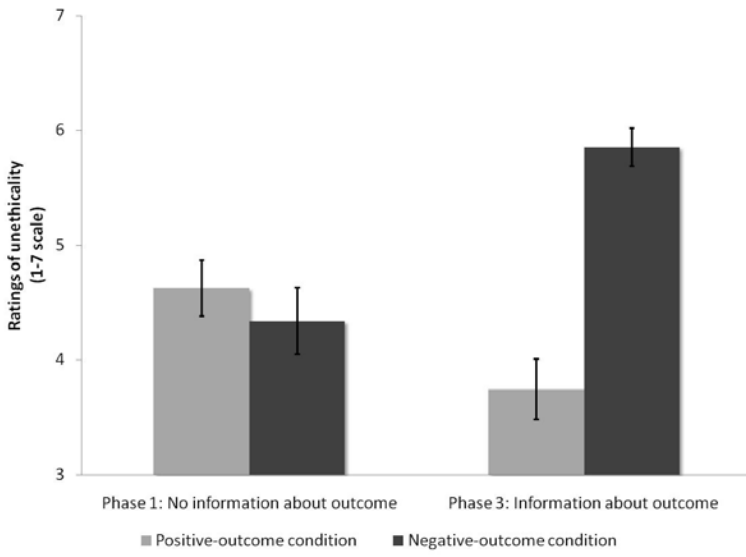


Figure 5.

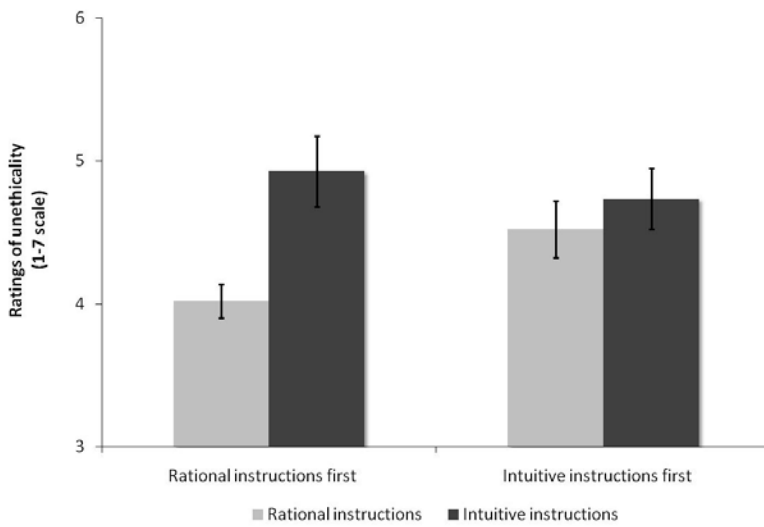


Figure 6.

