


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Organizational Behavior and Human Decision Processes

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

Justified ethicality: Observing desired counterfactuals modifies ethical perceptions and behavior

Shaul Shalvi^{a,*}, Jason Dana^b, Michel J.J. Handgraaf^a, Carsten K.W. De Dreu^a

^a Department of Psychology, University of Amsterdam, The Netherlands

^b Department of Psychology, University of Pennsylvania, United States

ARTICLE INFO

Article history:

Received 23 September 2010

Accepted 10 February 2011

Available online 8 March 2011

Accepted by Dr. Julie Irwin

Keywords:

Lies

Deception

Ethical decision making

Behavioral economics

Behavioral ethics

Morality

Ethics

Counterfactual thinking

Norm Theory

Mutability

ABSTRACT

Employing a *die-under-cup* paradigm, we study the extent to which people lie when it is transparently clear they cannot be caught. We asked participants to report the outcome of a private die roll and gain money according to their reports. Results suggest that the degree of lying depends on the extent to which self-justifications are available. Specifically, when people are allowed to roll the die three times to ensure its legitimacy, but only the first roll is supposed to “count,” we find evidence that the highest outcome of the three rolls is reported. Eliminating the ability to observe more than one roll reduces lying. Additional results suggest that observing desired counterfactuals, in the form of additional rolls not meant to determine pay, attenuates the degree to which people perceive lies as unethical. People seem to derive value from self-justifications allowing them to lie for money while feeling honest.

© 2011 Elsevier Inc. All rights reserved.

Morality, like art, means drawing a line somewhere. Oscar Wilde (1854–1900)

Introduction

Both within and outside organizations, daily life provides individuals ample opportunities to gain financially by bending the rules and lying. To some extent, the decision whether to lie is an economic calculation driven by whether the potential gain from lying is greater than the likelihood of getting caught times the magnitude of subsequent punishment (Alingham & Sandmo, 1972; Becker, 1968). To some extent, however, lying or not may depend on purely intrinsic moral standards and ethical considerations. For example, the more aversive the consequences of one's dishonesty on others' outcomes, the less likely it is that one will lie (Gneezy, 2005). Furthermore, people avoid major lies even when the chances of getting caught are essentially zero (Fischbacher & Heusi, 2008; Shalvi, Handgraaf, & De Dreu, 2011). Mazar, Amir, and Ariely (2008) explain the tendency to lie a little bit but not as much as one possibly could by proposing that people lie to some degree

to increase their profit, but not so much as to threaten their positive self-concept as honest individuals. To paraphrase Oscar Wilde, it seems that when considering lying or not, people draw the line somewhere. But the exact location of the line remains elusive, and it is yet unknown what leads people to decide that for a certain amount they would lie, but for a larger amount they would not.

We study lying by employing a *die-under-cup* paradigm (Shalvi et al., 2011; based on Fischbacher & Heusi, 2008) that allows participants to report the outcome of a die roll only they can see and gain money according to their reports. Consistent with previous studies (Fischbacher & Heusi, 2008; Gino, Ayal, & Ariely, 2009; Gino, Norton, & Ariely, 2010; Lundquist, Ellingsen, & Johannesson, 2009; Mazar et al., 2008; Shalvi et al., 2011) we find evidence of incomplete dishonesty. Reports differ significantly from the distribution of an honest die roll, but participants still lie far less than they could have. Novel to the current work, we find that the extent to which people allow themselves to lie depends critically on the availability of self-justifications that no one else knows about. Participants who were allowed to roll the die additional times to verify that it was legitimate lied to a greater extent and appear to report the largest number they saw on any roll. Participants reported the largest number they saw even though they knew that additional rolls were not supposed to “count” for determining pay. Observing a desired higher number on a roll that was not meant to “count” seems to justify lying using this one specific high number.

* Corresponding author. Address: Department of Psychology, University of Amsterdam, Roetersstraat 15, 1018 WB, The Netherlands. Fax: +31 206390531.

E-mail address: s.shalvi@uva.nl (S. Shalvi).

Importantly, our paradigm insures that these justifications are completely private; no one else knows what numbers the participant has seen. Additional studies corroborate this pattern of behavioral results by showing that of all possible profit boosting possibilities, participants report that they are most likely to use one of the additional rolls as means to increase their pay. Addressing the psychological driver for the behavioral pattern we found, we obtained evidence that participants judge dishonest reports to be less dishonest when they are equal to one of the additional rolls. The idea that people need self-justifications for lying is grounded in the notion that “people are likely to arrive to conclusions that they want to arrive at, but their ability to do so is constrained by their ability to construct seemingly reasonable justifications for these conclusions” (Kunda, 1990, p. 480).

Self-justifying unethical behavior

Initial evidence that people derive value from having justifications to dishonestly benefit themselves comes from work by Batson and colleagues (Batson, Kobrynowicz, Dinnerstein, Kampf, & Wilson, 1997; Batson, Thompson, Seufferling, Whitney, & Strongman, 1999). Participants in their studies had to determine whether they or another participant would have to perform an undesirable task and were provided with a coin they could toss to assist them in making this decision, so that they could either make the decision themselves, or use a coin toss to make the decision for them. Supporting their moral hypocrisy prediction, people who claimed to use the coin to make the decision “won” significantly more often (between 80% and 90%) than predicted by an honest toss. Participants were thus able “to appear fair by flipping the coin, yet still serve self-interest by ignoring the coin” (Batson & Thompson, 2001, p. 55). This line of research supports the idea that people find value in appearing fair and moral. They seek a ‘fair’ procedure (coin toss) to justify their self benefiting outcome to others.

A second line of evidence supporting the notion that people seek to appear fair while serving their self-interest comes from work on the ultimatum bargaining game, in which a proposer offers a division of a commodity (e.g., chips to be converted to money), and a responder can accept or reject the proposed division. If the responder accepts, the commodity is divided as proposed; if the responder rejects, neither party receives anything (Güth, Schmittberger & Schwarze, 1982). Pillutla and Murnighan (1995, 2003; see also Kagel, Kim, & Moser, 1996) used a modified version of the game providing only proposers with the value of the chips for themselves as well as for responders. By manipulating the value of the chips for the responders to be lower vs. equal to the value of the proposer’s chips, Pillutla and Murnighan were able to disentangle between proposer’s desire to act in a fair way (i.e., propose offers that are fair in monetary terms) or merely appear fair (i.e., propose offers that seem fair in terms of chips offered but are actually self-serving in monetary terms). Results indeed indicated that proposers made offers that seemed fair while they actually were not.

Further empirical evidence supporting the notion that justifications lead people to act on their self-interest while appearing moral comes from early work by Snyder, Kleck, Strenta, and Mentzer (1979). In this study, people had to choose in which of two rooms they would watch a movie and fill out a questionnaire – one with a handicapped person or with a non-handicapped person. When the same movie clip was presented in both rooms, people were more likely to sit with the handicapped individual. But when different movies were presented, providing a justification for selecting one room over the other, a majority avoided the handicapped person. This finding further suggests that when people have an apparent justification for not sitting with the handicapped person by way of their preference for one of the two clips, their preferred clip

matches the room which allows more morally questionable behavior (i.e., avoiding a handicapped person).

These lines of work clearly indicate that people find value in that they can appear fair and moral in the eyes of others and that if questioned about their decision they will be able to justify it by the ‘honest’ procedure they employed (i.e., a coin toss), the ‘fair’ offer they place on the table, or their ‘preferred’ means of entertainment (i.e., a specific video clip). However, in many daily situations we act in solitude without having to justify our (un)ethical behavior to anyone but ourselves. Mazar et al. (2008) indeed demonstrated that people value maintaining their honest self-concept; that is, they value feeling and not just appearing honest. Even though participants in their studies could (and did) lie anonymously, these participants refrained from lying to the fullest extent that they could. Mazar and colleagues explained this tendency to refrain from lying to large extents by suggesting that people have a desire to hold a positive self view which includes seeing themselves as ethical and honest people. Lying modestly thus allows one to serve two desires simultaneously – to benefit financially as well as to maintain an honest self-concept. We propose that justifications – here, in the form of observed desired counterfactuals – have an economic value in this equation. Justifications are valuable as they allow one to lie for money while maintaining an honest self-concept. Most importantly, our approach does not require a restriction of the magnitude of the lie as means to allow maintaining an honest self-concept. We propose that people would be able to maintain feeling honest when lying to different degrees (even a lot) conditional on having a justification for their unethical acts.

Consider, for example, a person using an online tax calculator to determine if she is entitled to a tax refund. If, after entering all her information, the initial results are disappointing, she may explore modified scenarios that would lead her closer to the desired refund. During this process of exploration, she may eventually discover that if one item had been different – e.g., if her partner had moved in with her on the 30th of June instead of the actual 1st of July, making them a shared household for 1 day more than 6 months instead of 1 day less than 6 months, her refund would be better. Experiencing the feeling of a desired outcome (a better tax return) resulting from this observed counterfactual information (moving in sooner) makes lying on the return in this one specific way seem more acceptable. We argue that observing desired counterfactual information justifies lying, as it modifies the extent to which a dishonest act is perceived as unethical, increasing one’s likelihood to lie. Importantly, we further maintain that this psychological mechanism occurs even when the justification is only for oneself. That is, that people find value in feeling honest even when they know that they lie.

Desired counterfactuals as justifications

The proposition that self-justifications are valuable in allowing people to feel honest while lying implies that honesty is not perceived as a sharp contrast between lying vs. being honest, but rather as a continuum that stretches between these two ends. Deontologically, a lie is a lie and is wrong. We however propose that when a desired state is observed people’s perception of what is ethical and what is unethical is modified. Accordingly, a given dishonest act can be perceived as more or less ethical depending on the availability of a self-justification for doing it, and whether such justification is used. Initial support for the idea that ethical evaluations (and subsequent behavior) are not a right vs. wrong dichotomy but rather a continuum ranging between the two ends, comes from work on the elastic justifications (Hsee, 1995, 1996). For example, Schweitzer and Hsee (2002) asked sellers of a car to provide a buyer with a mileage estimate from a range of possible values. Sellers lied more when the provided range was wide rather

than narrow, as they could justify the lie by their increased uncertainty about the true mileage. It seems that sellers processed the information about the car's mileage in a self-serving manner (Dunning, 2005; Dunning, Heath, & Suls, 2004; Ehrlinger & Dunning, 2003), allowing them to gain financially. Conversely, lying may decrease when available information provides no means to justify an unethical act.

The question remains, how did observing a desirable piece of information influence lying? Car sellers in Schweitzer and Hsee's study (2002) who received a wide range estimate may have been influenced by the (desired) high end of the range, which potentially served as a cognitive anchor to which they adjusted their mileage estimates (Kahneman, 1992; Mussweiler, 2003; Ritov, 1996; Tversky & Kahneman, 1974). Alternatively, merely observing a desired outcome that could have been true may have liberated some people's standards enough so that reporting the highest number seemed ethically acceptable. Similarly, in the tax example above, learning that reporting moving in with a partner 1 day before the move actually took place may make this one specific lie more legitimate than other lies which are based on information that was not observed.

Past work indeed reveals that considering upward counterfactuals "that are evaluatively better than actuality" (Roese, 1997, p. 134), influences people's (emotional) interpretation of factual reality (Mandel & Dhami, 2005; Zeelenberg et al., 1998) and their subsequent decisions (Markman, McMullen, & Elizaga, 2008; Morris & Moore, 2000; Reichert & Slate, 2000; Ritov & Baron, 1995; Roese, 1994; for a review see, Epstude & Roese, 2008). But how does observing a desired piece of information justify lying? And why would lying using this specific piece of observed information feel more legitimate compared to lying using other pieces of information that were not observed?

Kahneman and Miller's Norm Theory (1986) postulates that norms are computed after the event has occurred rather than in advance. Specifically, they proposed that "each stimulus selectively recruits its own alternatives (Garner, 1962, 1970) and is interpreted in a rich context of remembered and constructed representations of what it could have been, might have been, or should have been" (Kahneman & Miller, 1986, p. 136). A key aspect in Norm Theory is the level of mutability of the stimulus (e.g. information, object, circumstances) which is the extent to which this stimulus is modifiable. Low mutability suggests that it is very difficult to imagine the situation had been different, while high mutability means that this is easy. When it is possible to think of alternative scenarios leading to an unfavorable outcome (i.e., high mutability) people tend to focus on upward counterfactual comparisons (Markman, Gavanski, Sherman, & McMullen, 1993; Roese & Olson, 2007). The reason for this is that considering such desirable (upward) counterfactuals leads to reflective "as if" thinking: People reflect on the counterfactual "as if" it was real, thereby reducing the contrast between this piece of counterfactual information and factual reality (Markman & McMullen, 2003; Kahneman & Varey, 1990). Simply put, when reality is disappointing and desirable information becomes available, it is much easier to imagine the specific alternative scenario which would have happened if this piece of counterfactual information was true.

Here we propose that people observing a desired piece of information may feel that it is more legitimate to use this piece of information since it 'nearly' occurred, serving their self-interest as a result. When a desired piece of information is observed, it is this specific information, more than other desired alternatives that were not observed, that becomes readily available for consideration as an alternative fact. Put differently, when reality is disappointing and a desired piece of information is observed, people more easily construe reality by using the observed information that is available to them than any other information that was not

observed. We thus suggest that when an observed counterfactual is desirable compared to factual reality, dishonestly using this counterfactual to boost profit may feel less unethical than using other counterfactuals that were not observed.

Overview and predictions

We placed people in a situation in which they could lie to gain financially and manipulated whether they were able to observe desired counterfactuals. Specifically, we adapted a simple paradigm introduced by Fischbacher and Heusi (2008) in which people privately roll a die and are paid based on what outcome they report. Across their different experimental conditions including manipulations of the profit generated by the lie, the influence of the report on another person's outcomes, and the level of privacy in receiving pay for the task (i.e., by the experimenter vs. paying oneself from an envelope), a robust pattern of monotonically increasing reports emerged – higher outcomes were reported more often compared to lower outcomes. The deviation from the expected distribution of a fair roll reflected that people lied to some degree for financial gain. Based on our reasoning that people might use observed desired counterfactuals to boost profit, we focused on one aspect of Fischbacher and Heusi's paradigm that was not investigated before, namely the irrelevant rolls used to verify the die's legitimacy and not to determine pay. If reporting a desired counterfactual indeed seems like a more legitimate lie, people may report the highest outcome they observed on all rolls instead of only the one that should actually determine pay. People may have been reporting the highest of the multiple-rolls they saw – that is, the most desirable counterfactual they observed.

To address this possibility, we adapted Fischbacher and Heusi's (2008) paradigm to eliminate these extra rolls. Holding all other aspects of the paradigm constant, we introduced a condition we call *single-roll* that allows participants to verify that the die are legitimate while only rolling the die once during the task. We contrast this condition to a *multiple-rolls* condition in which participants verify the die's legitimacy by rolling it additional times after the roll that is to determine pay. Our main prediction was thus:

Hypothesis 1. Reducing people's ability to observe desired counterfactuals reduces lying.

Put differently, seeing only one die roll will reduce lying compared to seeing additional rolls because seeing only one roll reduces one's ability to justify the lie to him or herself by reporting a desired counterfactual that almost happened. Because reports using the die rolling paradigm were private, we could only statistically infer what strategies participants used. In Experiment 2, we explicitly asked participants who were explained this task to indicate what they would report when faced with various combinations of first and additional rolls. Doing so assessed if people indeed over-report the highest of these rolls more than any other profit boosting option available to them.

Experiments 3 and 4 were designed to test the proposed psychological mechanism accounting for this effect. Specifically, in these experiments we tested the prediction that the reason people report observed desired counterfactuals in order to profit financially is because reporting these justified outcomes seems less unethical than dishonestly reporting a desired outcome that was never observed. Simply put, we tested if seeing a justification makes lying for profit feel more honest compared to a similar lie without justification. The influence of having a justification on lying should thus be mediated by the modification in ethical perceptions. Therefore, a given lie feels less unethical with a justifica-

tion and thus is more likely to be used. Specifically, we tested the following prediction:

Hypothesis 2. Dishonestly reporting a desired counterfactual that was observed is considered less unethical than other desired dishonest reports that were not observed.

Experiment 1

Participants rolled a die and earned money according to what they reported rolling (Fischbacher & Heusi, 2008; Shalvi et al., 2011). We placed the die under a paper cup with small hole in the top and had participants shake the cup to roll the die and look through the hole to see the result. This procedure assured participants that only they could see the die (see Fig. 1). The first roll was to determine pay equivalent in US Dollars to the number they reported. After rolling for pay, participants were also instructed to roll at least two more times to ensure that the die was legitimate. This procedure also insured that their first roll could not be observed by anyone, even after they left the experiment. In another condition, holding all other situational parameters constant, we eliminated participants' ability to observe more than one roll before reporting their outcomes. We tested if eliminating people's ability to observe a desired counterfactual in the form of a higher number on one of the subsequent rolls (meant only to verify the die's legitimacy and not to determine pay), reduces lying.

Participants and procedure

One hundred and twenty night students at an eastern US university were randomly assigned to one of two die roll conditions (*multiple-rolls* vs. *single-roll*). Participants arrived at the lab in groups of 8–19, were seated in individual cubicles and performed the *die-under-cup* task, by rolling a six-faced die under a cup, reporting the outcome, and earning in dollars the number they reported. Participants in the *multiple-rolls* condition were instructed to shake the cup once, check the outcome that determined their pay, and then roll at least two more times to make sure that the die was legitimate. They were then asked to write down the first roll outcome. In the *single-roll* condition, participants first saw a box containing all the dice that were used in their experimental session and were allowed to test that the dice were legitimate in any way they liked, including rolling them. After testing the dice, participants were then instructed on and received a *die-under-cup*. Participants were asked to roll and check the outcome, but not instructed to roll again. To assure participants that their outcomes were private in this condition, the experimenters passed around a box into which each participant swept his cup and die. This procedure was demonstrated by the experimenter before participants began the actual task. In summary, both the *multiple-rolls* and the *single-roll* conditions allowed the participants to privately roll a die, verify that it was legitimate, and hide the result from being viewed by anyone in the future, but the two condi-

tions differed critically in that the *multiple-roll* condition allowed participants to roll the die multiple times after they knew what the task was.

Results

If people use observed desired counterfactuals to unethically increase their profit, we should find less lying in the *single-roll* compared to the *multiple-rolls* condition, and the distribution of *multiple-rolls* reports should be consistent with reporting the best out of three outcomes. Indeed the *multiple-rolls* distribution differed from the uniform distribution expected from a fair die (Kolmogorov–Smirnov $Z = 2.03$, $p < .01$). This was due to underreporting of 1's, 2's, and 3's (95%CI's < chance [16.67%] for 1's and 2's and 90%CI < chance for 3's) and over-reporting 6's (95%CI > chance). In the *single-roll* condition, lying was reduced and the distribution did not significantly differ from chance (Kolmogorov–Smirnov $Z = 1.16$, *ns*). Importantly, as predicted by Hypothesis 1, people lied significantly less in the *single-roll* condition ($M = \$3.97$, $SD = 1.56$) than in the *multiple-rolls* condition ($M = \$4.45$, $SD = 1.59$), Mann–Whitney $Z = 1.93$, $p = .05$, see Fig. 2.

We compared the distributions in each condition to the theoretical distribution of choosing the highest value of three die rolls, a mechanism for lying that we propose. In the theoretical distribution of choosing the highest of three rolls (see Fig. 2, left side), the likelihood of seeing 1 as the highest value of the three rolls is very low (it happens only if one rolls 1 three times, which is 1 of $6^3 = 216$ combinations), whereas seeing 6 as the highest value of three rolls is a more likely event (it happens in 91 of 216 three die roll combinations). If people report desired counterfactuals (the highest value they observed) the distribution of the *multiple-rolls* condition should resemble the distribution resulting from choosing the highest of several rolls. Indeed, supporting Hypothesis 1, the *multiple-rolls* distribution did not differ from the theoretical distribution of choosing the highest of three rolls (Kolmogorov–Smirnov $Z = .97$, *ns*), whereas the *single-roll* distribution did (Kolmogorov–Smirnov $Z = 2.34$, $p < .01$). Hence, we could not reject the hypothesis that people who rolled more than once chose the highest of three rolls they have observed.

The behavioral pattern collected in this study allows us to assess the potential range of honest people as well as income-maximizers that would lie in all cases (see Fischbacher & Heusi, 2008). There is no reason to assume that anyone reporting a '1' would be lying. In our sample, 8% in both single and multiple-rolls conditions reported rolling '1' suggesting that at least 8% of people would not lie whether or not they have a justification to do so. On the other extreme, 34% of participants in the *multiple-rolls* condition and 19% in the *single-roll* condition reported rolling a '6'. Because '6' would be reported an expected 1/6th of the time if participants were reporting honestly, we can calculate the proportion of people who lie to the maximum by subtracting 1/6th from the observed proportions and multiplying the outcome by 6/5. As clarified by Fischbacher and Heusi, this multiplication is required to take into account those participants that actually rolled a '6'




Fig. 1. *Die-under-cup*. Under complete anonymity conditions, participants received the option to roll a *die-under-cup* that only they could see, report what they rolled, and receive \$US as a function of the amount they reported rolling.




Fig. 2. The theoretical distribution of choosing the highest of three die rolls (left), the percentage of reported rolled outcomes in the *multiple-rolls* (middle) and *single-roll* (right) conditions, Experiment 1. The dashed line represents the honest distribution predicted by chance (16.67% per die face). The error bars represent the 95% confidence interval around the proportion.

but would have lied if they have rolled a lower number. In the *multiple-rolls* condition, the proportion of income-maximizers is thus $(34-17\%) * 6/5 = 20\%$ which is very similar to Fischbacher and Heussi's 22% estimate. However, as demonstrated here, the proportion of income-maximizers might be much lower when a justification in the form of extra rolls is removed. Calculating the proportion of liars in the single rolls condition reveals that only $(19-17\%) * 6/5 = 2.5\%$ of people may be considered as people who would lie to the maximum extent regardless of the outcome they observed.

Experiment 2

Supporting Hypothesis 1, limiting one's ability to observe desired counterfactuals reduced lying. People who saw only one roll lied less compared to those who saw multiple-rolls. Furthermore, providing initial support to Hypothesis 1, we could not reject the hypothesis that people who rolled more than once reported rolling the highest of three rolls they observed (which was the minimal number of times they were instructed to roll).

The method employed in Experiment 1 made it absolutely transparent that one's actions can never be observed. As such, it rooted out fears of detection and allowed testing internal standards. Justifications in the form of the extra rolls were in the eyes of the person rolling only; no one else could know they had these justifications. Statistical analyses further demonstrated that the distribution of reported outcomes for participants, who could roll more than one time, resembled the theoretical distribution of choosing the highest of three rolls. We could not, however, directly observe what people saw before reporting the outcome. In Experiment 2, we asked people directly what they would have reported in such ethically challenging situations. While this procedure does not motivate participants' reports financially, it does allow a more direct test of the prediction that people report the desired counterfactual they observe more than any other profit boosting options that are available to them.

Participants and procedures

Sixty-eight people (51% females, $M_{age} = 26.71$) were recruited online and paid for participation. We employed a $6 (1st\ roll: 1-6) \times 6 (2nd\ roll: 1-6)$ within subject design. Participants read

the instructions of the *multiple-rolls* condition and indicated what they would have reported in all randomly presented two die roll combinations. For example, a participant was shown rolls of 2 and 5 and asked what she would report.

Comprehension checks

Participants were asked three questions to verify that they understood the task: whether the first roll counts for pay, whether second roll was used to verify that the die was legitimate, and whether when one sees a 3 on the first roll that was also the number one should have reported. Because incorrect comprehension of the task could lead to a false appearance of using counterfactuals – e.g., participants might think that the task demands reporting the highest roll – 17 participants who did not answer these questions correctly were excluded from all analyses. Including these participants did not modify any of the reported results.

Results

Most people stated that they would have reported honestly. Specifically, 69.6% would have honestly reported rolling a 1, 73.9% when rolling a 2, 81.7% when rolling a 3, 88.6% when rolling a 4, 89.2% when rolling a 5, and 94.4% when rolling a 6. An ANOVA predicting the stated likelihood to lie from rolled outcome revealed a linear effect, $F(1, 50) = 21.64, p < .0001, \eta^2 = .30$. The higher the actual roll was, the more likely the participants were to give an honest report. These greater reports of honesty than were seen in Experiment 1 are not surprising in some sense since participants did not gain financially from reporting higher numbers.

Providing further support to Hypothesis 1, among those reporting that they would have lied, the outcome most commonly used was the desired counterfactual, see Table 1. For all dishonest reports, we calculated the proportion of people who reported the second roll outcome conditional on each first roll outcome. We then tested whether this proportion differed from that specified by a null hypothesis in which dishonest reports were chosen randomly from a uniform distribution of higher outcomes. For example, of the 83 dishonest reports after seeing 1 on the first roll, 40 (48.2%) were equal to the value that appeared on the second roll. A binomial test revealed that this proportion was higher than the 20% predicted by our null hypothesis (e.g., 2, 3, 4, 5, or 6),

Table 1
Frequency of estimated reports as a function of the two observed die rolls (Experiment 2).

Rolled	Report?						Total
	1	2	3	4	5	6	
1 then 2	36	5	4	1	2	3	51
1 then 3	36	2	8	1	1	3	51
1 then 4	35	1	2	8	2	3	51
1 then 5	33	1	3	2	9	3	51
1 then 6	32	2	4	1	2	10	51
2 then 3	0	39	5	2	2	3	51
2 then 4	0	35	4	7	1	4	51
2 then 5	0	34	4	2	8	3	51
2 then 6	0	35	4	1	2	9	51
3 then 4	0	0	43	3	2	3	51
3 then 5	0	0	39	3	6	3	51
3 then 6	0	0	37	1	4	9	51
4 then 5	0	0	1	45	2	3	51
4 then 6	0	0	1	42	0	8	51

For brevity considerations, table presents only results of die roll combinations allowing one to lie by using (or not) the desired counterfactual appearing on the second roll. Honest reports appear in *italic* and justified lies appear in **bold**.

$p < .0001$. Similarly, 29 of the 61 dishonest reports after seeing a 2, 18 of the 34 dishonest reports after seeing a 3, and 10 of the 13 dishonest reports after seeing a 4 were 2nd roll values, all proportions larger than expected by chance ($p < .001$, $p = .03$, and $p = .09$ respectively). Taken together, 51% (97 of 191) dishonest reports were justified lies (i.e., equal to the observed desired counterfactual). This proportion was almost twice as high as the proportion of choosing a higher number at random (26%, $p < .01$).¹

Experiment 3

Results of Experiment 2 provided further support for our main prediction that people use the observed desired counterfactual to unethically boost profit more than any other option available to them. In Experiment 1, support was gained by eliminating the ability to observe more than one roll, which reduced lying. Furthermore, in the multiple-rolls condition the distribution of reported outcomes did not differ from the distribution of reporting the highest of three rolls, while in the single-roll condition it did. Taken together we concluded that, as predicted, people seem to be reporting the highest of the rolls they observed. That said, as people's rolls were completely private, these behavioral results provided strong but indirect evidence for our predicted mechanism. The results of Experiment 2 provided more direct evidence that among people who estimated that they would report dishonestly, the most common reported outcome, across all die roll combinations, was the highest observed value.

In Experiments 3 and 4, we moved forward to test our predictions about the underlying psychological mechanism leading people to report the highest of the rolls they saw. We predicted that reporting an observed value that is not supposed to count would feel less unethical compared to reporting any other inflated value that was not observed. In turn, this reduced level of perceived

¹ We calculated the weighted expected proportion of untruthful reports that would have been predicted if people who reported dishonestly chose a higher number than the number they saw on the 1st roll at random. We multiplied the number of dishonest reports following each of the 1st roll values by the proportion that reflects the likelihood that such report would have been a random choice (i.e., .20 for 1st roll = 1 in which lying could be done by reporting 2, 3, 4, 5, or 6; .25 for 1st roll = 2; .333 for 1st roll = 3; and .50 for 1st roll = 4). The number of dishonest reports expected if people were choosing a higher value at random was thus: $(83 * .20) + (61 * .25) + (34 * .33) + (13 * .50) = 49.57$, or 26% (49.57 of 191) of the total number of dishonest reports. The 99% CI for the observed proportion of justified lies (51%) ranged between 42% and 60% which was well above this expected proportion. This indicates that people were reporting the justified lies more than any other profit boosting values that were available to them.

unethicality will increase one's likelihood to lie. Put differently, seeing a high value on a roll not meant to count makes it more legitimate to lie as long as the observed value is reported. Note that this is all the more interesting, since if one decides to lie, one might as well always report the highest possible outcome, rather than the observed desirable counterfactual. To gauge people's perceptions of what is ethical and what is not, Experiment 3 asked people to evaluate to what extent different dishonest reports were considered to be unethical. Specifically, we presented people with different combinations of die rolls and reports, and asked them to what extent they considered these reports to be lies. We tested whether observing desired counterfactuals makes a dishonest report seem less unethical compared to a dishonest report occurring without observing the counterfactual. That is, whether observing desired counterfactuals justifies lying.

Participants and procedures

One hundred fifteen people (44% females, $M_{age} = 29.98$) were recruited online and paid for participation. We employed a 6 (1st roll: 1–6) \times 6 (2nd roll: 1–6) \times 6 (report: 1–6) within subjects design, except that we omitted combinations in which the report was lower than the first roll (lying downward), resulting in 126 combinations. Each participant was randomly assigned to evaluate 48 of these combinations in one of three different randomly predetermined blocks (the order of presented combinations within blocks was also randomized). Blocking did not moderate any of the described results and is thus not discussed any further. Participants were asked to rate for each roll/report combination the extent to which they considered each report to be a lie (1 = not at all to 7 = very much).

Comprehension checks

As in Experiment 2, participants were asked three questions to verify that they understood the task: whether the first roll counts for pay, whether second roll was used to verify that the die was legitimate, and whether when one sees a 3 on the first roll that was also the number one should have reported. The 32 participants who did not understand these rules were excluded from all analyses. Analyzing the data without those participants was a conservative test to our hypothesis. This is because participants who fail to understand the task were likely to be lenient about reporting the highest of the observed rolls merely because they failed to understand that such behavior violates the task's rules. Including these participants did not modify any of the reported results.

Results

Die roll/report combinations were recoded to reflect whether the person rolling was honest (e.g., rolled 2 then 3 and reported 2) or lied to increase profit (e.g., rolled 2 then 3 and reported 4). Not surprisingly, an ANOVA with honest vs. dishonest reports as within subject factor, revealed that reporting honestly was considered less of a lie ($M = 1.87$, $SD = 1.28$) than reporting dishonestly ($M = 5.72$, $SD = 1.47$), $F(1, 82) = 189.54$, $p < .001$, $\eta^2 = .70$. Honest reports received ranking higher than '1' due to some participants' tendency not to use the full Likert scale available to them. Notwithstanding this tendency, across all truthful combinations the mode was always '1' assuring that most participants evaluated honest reports by choosing the most honest end of the scale.

We further recoded "justified lies" as untruthful reports equal to the desired counterfactual (e.g., rolled 2 then 5 and reported 5) and "unjustified lies" as untruthful reports that were different from the desired counterfactual (e.g., rolled 2 then 3 and reported 4). Supporting Hypothesis 2, an ANOVA with justified vs.

unjustified lies as within subject factor revealed that an untruthful report that was justified (i.e., when the report was equal to the 2nd roll value) was considered less of a lie ($M = 5.57$, $SD = 1.59$) compared to an unjustified (i.e., when the report was not equal to the 2nd roll value) untruthful report ($M = 5.87$, $SD = 1.47$), $F(1, 82) = 10.52$, $p = .001$, $\eta^2 = .11$. This effect was primarily driven by unjustified lies which were higher than the second roll outcome ($p < .01$), which is the key comparison for our reasoning as such combinations provide absolutely no justification for the dishonest report.

Experiment 4

Results from Experiment 3 supported Hypothesis 2. When considering whether a report was ethical, people took into consideration the outcome obtained on a roll not meant to determine pay. People considered dishonest reports to be less unethical with (rather than without) a possible justification for a lie. In Experiment 4, we manipulated the number of observed desired counterfactuals by providing information about the first roll that counted for pay, and two additional rolls assessing the die's legitimacy. Participants had to evaluate the extent to which they consider another person's lying as unethical after this person saw no (vs. one vs. two) desired counterfactuals. They were also asked to estimate the extent to which they were likely to act the same (i.e., lie).

Participants

Sixty-nine undergraduate students at an eastern US university completed a "die roll survey" after participating in another experiment.

Design and procedures

Participants filled out a survey with 25 three die roll combinations. Participants learned that a person reported rolling a '6' after rolling a lower number than '6' on the first roll and rolling no '6' (vs. one '6' vs. two '6's') in two subsequent rolls determining the die's legitimacy. We employed a 5 (1st roll: 1–5) \times 3 (combined number of '6's' appearing in the second and third rolls: 0 vs. 1 vs. 2) within subject design. Participants indicated for each combination the extent to which they considered the report a lie and the likelihood that they would have done the same (both on a 1 'not at all' to 7 'very much' scale). All participants reported that they understood the rules.

Results

Is this a lie? We recoded all combinations to indicate the number of '6's' they include in the two subsequent rolls (0 vs. 1 vs. 2).² Supporting Hypothesis 2, an ANOVA with number of '6's' (0 vs. 1 vs. 2) as within subject factor predicted the extent to which reporting a '6' was considered a lie: reporting a '6' without seeing a '6' was considered a greater lie ($M = 5.76$, $SD = .56$) than doing so after seeing one '6' ($M = 5.29$, $SD = 1.08$), which in turn was considered a greater lie than doing so after seeing two '6's' ($M = 4.91$, $SD = 1.59$), $F(1, 68) = 23.98$, $p < .001$, $\eta^2 = .26$ (simple effects, $F_s > 18.78$, $p_s < .001$).

Would you do the same? Supporting Hypothesis 1, an ANOVA with number of 6's (0 vs. 1 vs. 2) as a within subject factor predicted the extent to which participants reported that they would have done the same (i.e. lie). Participants indicated that they were less likely to have reported a '6' without seeing a '6' ($M = 1.92$,

$SD = 1.17$) than doing so after seeing one '6' ($M = 2.51$, $SD = 1.47$) and after seeing two '6's' ($M = 2.82$, $SD = 1.72$), $F(1, 68) = 32.58$, $p < .001$, $\eta^2 = .32$ (simple effects, $F_s > 14.79$, $p_s < .001$),³ see Fig. 3. Further analyses revealed that the modification in ethical perceptions partially mediated the effect of observing a desired counterfactual (a '6') on participants' reported likelihood that they would lie as well.⁴

Experiment 4 discussion

Experiment 4 provided further support for hypotheses H1 and H2: Participants evaluated an untruthful report as less unethical as a function of whether a desired counterfactual was observed. They further indicated that they were more likely to lie when observing such desired counterfactuals. Additional analyses provided support for the idea that the relationship between seeing a desired counterfactual and one's likelihood to lie was partially mediated by one's modification in the perceived unethicality of the act.

One limitation to the results of Experiment 4 relates to the fact that since it was a paper and pencil survey in which order of presented die roll combinations was not randomized, we could not rule out order effects. The possibility that order effects were driving any of the reported results of Experiment 4 seems unlikely however, as these results replicate the results obtained in Experiment 3, in which randomized combinations were presented.

General discussion

The current work contributes to our understanding of people's perceptions of ethicality and how these perceptions translate into unethical behavior. Employing a simple *die-under-cup* paradigm that allowed participants to lie anonymously and gain financially, we found that observing desired counterfactuals leads people to (1) lie more than people who do not observe desired counterfactuals, (2) evaluate unethical acts as more ethical, and (3) assess their own likelihood to commit an unethical act to be higher due to the modification in ethical perceptions.

In Experiment 1, we found that allowing people the opportunity to lie without getting caught led to dishonesty. However, while evidence of lying was clear, the amount of lying was modest. For example, only 34% of participants in the multiple-rolls condition reported a 6, when in fact all of them could have gotten away with reporting a 6. This finding confirms other recent work showing that people limit the magnitude of their lying even when they cannot be caught (Fischbacher & Heusi, 2008; Gino et al., 2009, 2010; Mazar et al., 2008; Shalvi et al., 2011). However, we also show that beyond the magnitude of lying, people rely critically on self-justifications to enable themselves to lie, even when these justifications

³ As the variances in the different levels of the repeated measures were not homogeneous in both reported analyses, F_s are reported using the Greenhouse-Geisser correction for the violation of the sphericity assumption. Analyzing the data using Friedman's non-parametric test for repeated measures designs yielded significant effects ("is this a lie?", $\chi^2(2) = 41.64$, $p < .001$; "do the same", $\chi^2(2) = 29.35$, $p < .001$) which further corroborated the robustness of these effects.

⁴ Following the Judd, Kenny, and McClelland (2001) method for testing mediation in within-group designs, we distinguished between combinations that did not include the desired counterfactual (no '6') and those that did (one or two '6's'). First, as described above, seeing a desired counterfactual was related both to the proposed mediator (ethical perceptions) and the DV (likelihood to lie). Second, the mediator was significantly related to the DV at each level of the IV (both $t_s > 2.9$, $p_s < .001$). Finally, providing support for Hypothesis 3, the effect of seeing a desired counterfactual on the likelihood to lie was significantly reduced when controlling for perceived ethicality, Sobel's $Z = 3.63$, $p < .001$. Since the effect of observing a desired counterfactual was not reduced to zero ($t = 2.80$, $p < .01$), these analyses indicated that perceived lying partially mediated the effect of observing a desired counterfactual on the likelihood to lie.

² The location of the '6' (2nd vs. 3rd roll) in those combination who had only one '6' did not influence any of the reported results.

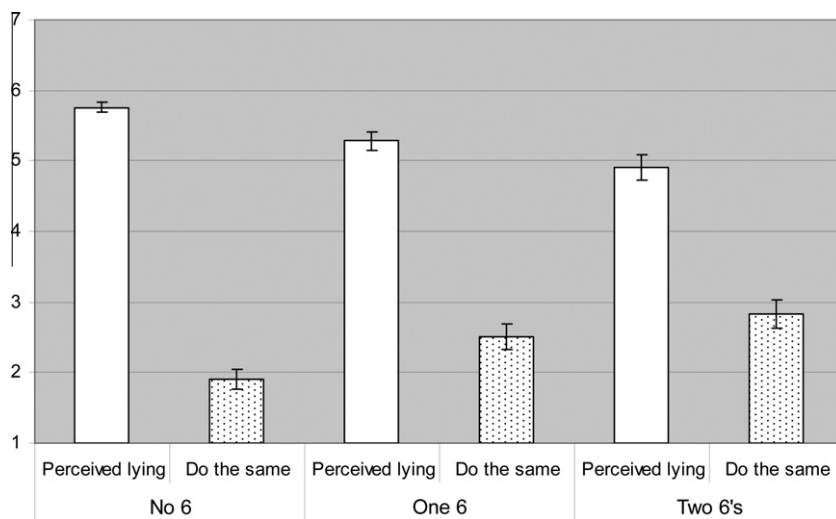


Fig. 3. Mean perceived lying and likelihood to act the same (and lie by reporting a non-existing 6) as a function of the number of (non-relevant) 6's one saw (Experiment 4). Error bars are $\pm 1SE$ around the mean.

will never be available to other people. This conclusion stems from three pieces of evidence. First, in Experiment 1, the distribution of reported outcomes did not differ from the theoretical distribution of choosing the highest of three rolls, the mechanism for lying we proposed. Second, eliminating people's ability to observe more than one roll reduced lying. This solid piece of evidence speaks to the fact that the extra rolls were worth money to participants although these rolls were clearly for their eyes only. In the single-roll condition, people made less money because they could not avail themselves of these justifications. Third, in Experiment 2, participants indicated what they would have reported in all possible die roll combinations. Over all possible combinations, people reported the outcome of the second roll value more than any other profit boosting value available to them. Together, these findings suggest that private justifications translate to larger lies.

We further proposed that self-justifications allowed people to lie more for money while feeling honest. That is, we suggested that justifications make a given dishonest act feel less unethical compared to the exact same dishonest act that occurs without a justification. Results of Experiments 3 and 4 supported this idea. We found that people viewed a lie more severely when it did not match the counterfactual support of an additional die roll. In Experiment 4 we gained support for the notion that the modification in one's perception of ethicality partially mediates the relation between observing a desired counterfactual and lying. Our results do not imply that people do not care about their ability to appear honest in the eyes of others. Obviously, as pointed out in the introduction, there is ample evidence that such considerations do play an important role in determining one's likelihood to lie. The contribution of the current work is that over and above such a desire to appear honest in the eyes of others (which was constant across our experimental conditions), people value feeling honest in their own eyes. Importantly, we found that having a justification to behave unethically, in the form of observing desired counterfactuals, allow people to maintain a feeling of honesty even when lying quite a bit.

An interesting finding in Experiment 4 was the different level of variance observed for both ratings of how unethical it is to lie without seeing a '6', when seeing one '6' and when observing two '6's' as well as how likely people are to lie in these conditions. The variance increased when more justifications were observed. While this issue did not pose a statistical problem in interpreting the data, from a theoretical point of view this finding may be of inter-

est as it seems to suggest that people are more polarized in their opinions when evaluating justified lies compared to when evaluating unjustified lies. Put differently, when justifications are not available people seem to be rather unanimous about the fact that lying is unethical while when justifications are available people are relatively more diverse in their interpretation of the lies. This potentially points out to some individual difference on this issue. Future research may be needed to further explore this possibility in greater detail.

Limitations

The *die-under-cup* paradigm prevented us from assessing who actually lied. Some may consider the fact that we cannot trace individual reports back to observed die rolls as the downside of the paradigm. Indeed, in Experiment 1, we were unable to collect self-report data to assess the mediating process for dishonesty. We believe, however, that the indispensable feature of our approach is that it makes absolutely transparent to participants that their actions could never be observed. In this way, we can be confident in our data interpretation that our participants were not concerned that their actions would be discovered and that our experiment simply tests their propensity to lie as governed by their internal standards (for similar approach see e.g., Greene & Paxton, 2009; Lammers, Stapel, & Galinsky, 2010; Mazar et al., 2008). As such, using justifications in the form of the additional rolls that were observed were absolutely for the self.

The fact that we could not trace back actual reports in Experiment 1 forced us to assess lying indirectly by comparing the distribution of outcomes to different theoretical distributions (i.e., honest, best of three). The significant difference between an observed distribution in the multiple-rolls condition and a theoretical honest distribution provided solid indication for lying. However, the lack of difference between the observed distribution and the theoretical distribution of choosing the highest of three provided indirect support as it did not reject the null (no difference) hypothesis. We used a hypothetical experiment to fill in this gap. Results of Experiment 2 helped in that they clarified that indeed a significant portion of people indicating that they would lie chose the value appearing on the second roll as a way to boost their profit. Together, these pieces of evidence converge to support our

proposition that people use private justifications when deciding the extent to which they can lie.

Theoretical contribution and boundary conditions

Gneezy (2005) noted that “people have non-consequential preferences in which they treat the same monetary outcome differently, depending on the process that leads up to it” (p. 392). Corroborating this notion, we found that people’s likelihood to lie for a given outcome was influenced by whether the outcome was observed while assessing the die’s legitimacy. It seems that the utility of a given monetary outcome varies depending on whether it is attained by a (non)justified lie. Recent work (Shalvi et al., 2011) reveals that the likelihood that people will lie for a specific amount depends on the additional profit boosting alternatives that are available. People are less likely to lie when the available options are narrowed to include only major lies (lying to the maximum extent) or minor ones (increasing profit only marginally above an exit-option) compared with situations that also include middle-of-the-road lies (increasing profit substantially but not to the maximum). Narrowing the range of potential lies may have reduced people’s ability to justify a potential unethical act, in turn decreasing lying.

The idea that justifications modify perceptions of ethicality further communicates with work on “bounded ethicality” (Bazerman & Banaji, 2004; Chugh, Bazerman, & Banaji, 2005) proposing that implicit biases in people’s awareness systematically prevent them from identifying others’ wrongdoing. Attempting to detect cheaters, people were less tolerant when they observed an abrupt change in performance rather than multiple, gradual changes (Gino & Bazerman, 2009). Our findings may suggest that people’s lack of awareness of others’ misconduct, leading to an ethical slippery slope, may be due to the extent to which the others’ unethical behavior may be justified. Non-obvious lies may be more likely to be perceived in ethically lenient ways. Future research is needed to conclusively address this possibility and determine to what extent our ethicality is bounded or justified.

Our reasoning linking justifications and counterfactual thinking relied on Kahneman and Miller’s Norm Theory, which suggests that norms are established after the fact and based on available information and one’s ability to reinterpret this information. A key factor in this theory is the extent to which the available information is mutable (i.e., modifiable). This is one key aspect that may provide more insight into the boundary conditions of the phenomenon described in our studies. There are several leading candidates for such boundary conditions. First, a mutable event is an event that people may influence and modify, or at least feel that they may. A likely boundary condition to our results is thus reducing people’s control over generating the justification, which should reduce their likelihood to use the justification. Manipulating whether the participant or another person rolls the die may yield interesting data pertaining to this issue.

Second, it is possible that merely imagining the alternative scenarios (rather than actually observing them) may be sufficient to justify lying. Attempting to clarify this issue would help determine whether it is the mental simulation or the actual evidence that is used as a justification for lying. Third, past work found that recent events are more mutable than earlier events (Miller & Gunasegaram, 1990). It is thus possible that manipulating the roll that counts for pay to be the first vs. a later roll would influence people’s likelihood to use the non-relevant for pay rolls as justifications. This line of reasoning suggests that lying about the outcome of roll X (e.g., the 5th roll) using the values that were observed before would be relatively difficult as all the counterfactual evidence that preceded the roll that count is less mutable at the time of reporting. Studying these potential boundary conditions to our results will not only shed light about

what type of information is more likely to be used as justification, but will also contribute to the more general discussion about the relation between mutability and counterfactual thinking.

Concluding thoughts

It seems that Oscar Wilde was right in suggesting that morality, like art, means drawing a line somewhere. The current results suggest that this line is drawn exactly where justifications end. The current work contributes to our understanding of how justifications that are available only for the self shape one’s ethical perceptions and behavior. Using a simple *die-under-cup* paradigm we provided evidence that people report rolling the highest of the rolls they see in order to gain money, while they know that the rules clearly indicate that they should report only the first roll outcome. As such, people seem to drive value from self-justifications and lie for money as long as they can feel honest.

Acknowledgments

We thank Jeff DeWitt for his help in collecting the data, Imre Aniba for posing in the pictures of Fig. 1, and Femke Ten Velden for taking them. Jason Dana is thankful for support from the Goldstone Research Unit at the University of Pennsylvania. Shaul Shalvi is thankful to the European Association of Social Psychology, the Society for Judgment and Decision Making, and the European Association for Decision Making for their help in funding a visit to the University of Pennsylvania.

References

- Alingham, M. G., & Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of Public Economics*, 1, 323–338.
- Batson, C. D., Kobrynowicz, D., Dinnerstein, J. L., Kampf, H. C., & Wilson, A. D. (1997). In a very different voice: Unmasking moral hypocrisy. *Journal of Personality and Social Psychology*, 72, 1335–1348.
- Batson, C. D., & Thompson, E. R. (2001). Why don’t moral people act morally? Motivational considerations. *Current Directions in Psychological Science*, 10, 54–57.
- Batson, C. D., Thompson, E. R., Seufferling, G., Whitney, H., & Strongman, J. (1999). Moral hypocrisy: Appearing moral to oneself without being so. *Journal of Personality and Social Psychology*, 77, 525–537.
- Bazerman, M. H., & Banaji, M. R. (2004). The social psychology of ordinary unethical behavior. *Social Justice Research*, 17, 111–115.
- Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76, 169.
- Chugh, D., Bazerman, M. H., & Banaji, M. R. (2005). Bounded ethicality as a psychological barrier to recognizing conflict of interest. In D. A. Moore, D. M. Cain, G. Lowenstein, & M. H. Bazerman (Eds.), *Conflict of interest: Challenges and solutions in business, law, medicine, and policy* (pp. 74–95). NY, NY: Cambridge University Press.
- Dunning, D. (2005). *Self insight: Roadblocks and detours on the path to knowing thyself*. New York: Psychological Press.
- Dunning, D., Heath, C., & Suls, J. M. (2004). Flawed self-assessment: Implications for health, education, and the workplace. *Psychological Science in the Public Interest*, 5, 69–106.
- Ehrlinger, J., & Dunning, D. (2003). How chronic self-views influence (and potentially mislead) assessments of performance. *Journal of Personality and Social Psychology*, 84, 5–17.
- Epstude, K., & Roese, N. J. (2008). The functional theory of counterfactual thinking. *Personality and Social Psychology Review*, 12, 168–192.
- Fischbacher, U., & Heusi, F. (2008). Lies in disguise, an experimental study on cheating. TWI working paper 40, Thurgau Institute of Economics, University of Konstanz.
- Garner, W. R. (1962). *Uncertainty and structure as psychological concepts*. NY: Wiley.
- Garner, W. R. (1970). Good patterns have few alternatives. *American Scientist*, 58, 34–42.
- Gino, F., Ayal, S., & Ariely, D. (2009). Contagion and differentiation in unethical behavior: The effect of one bad apple on the barrel. *Psychological Science*, 20, 393–398.
- Gino, F., & Bazerman, M. H. (2009). When misconduct goes unnoticed: The acceptability of gradual erosion in others’ unethical behavior. *Journal of Experimental Social Psychology*, 45, 708–719.
- Gino, F., Norton, M., & Ariely, D. (2010). The counterfeit self: The deceptive costs of faking it. *Psychological Science*, 21, 712–720.

- Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, 95, 384–394.
- Greene, J. D., & Paxton, J. M. (2009). Patterns of neural activity associated with honest and dishonest moral decisions. *Proceedings of the National Academy of Sciences USA*, 106, 12506–12511.
- Güth, W., Schmittberger, R., & Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, 3, 367–388.
- Hsee, C. K. (1995). Elastic justification: How tempting but task irrelevant factors influence decisions. *Organizational Behavior and Human Decision Processes*, 62, 330–337.
- Hsee, C. K. (1996). Elastic justification: How unjustifiable factors influence judgments. *Organizational Behavior and Human Decision Processes*, 66, 122–129.
- Judd, C. M., Kenny, D. A., & McClelland, G. H. (2001). Estimating and testing mediation and moderation in within-subjects designs. *Psychological Methods*, 6, 115–134.
- Kagel, J. H., Kim, C., & Moser, D. (1996). Fairness in ultimatum games with asymmetric information and asymmetric payoffs. *Games and Economic Behavior*, 13, 100–110.
- Kahneman, D. (1992). Reference points, anchors, norms, and mixed feelings. *Organizational Behavior and Human Decision Processes*, 51, 296–312.
- Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93, 136–153.
- Kahneman, D., & Varey, C. A. (1990). Propensities and counterfactuals: The loser that almost won. *Journal of Personality and Social Psychology*, 59, 1101–1110.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108, 480–498.
- Lammers, J., Stapel, D. A., & Galinsky, A. D. (2010). Power increases hypocrisy moralizing in reasoning, immorality in behavior. *Psychological Science*, 21, 737–744.
- Lundquist, T., Ellingsen, T., & Johannesson, M. (2009). The aversion to lying. *Journal of Economic Behavior and Organization*, 70, 81–92.
- Mandel, D. R., & Dhali, M. K. (2005). "What I did" versus "what I might have done": Effect of factual and counterfactual thinking on blame, guilt, and shame in prisoners. *Journal of Experimental Social Psychology*, 41, 627–635.
- Markman, K. D., Gavanski, I., Sherman, S. J., & McMullen, M. N. (1993). The mental simulation of better and worse possible worlds. *Journal of Experimental Social Psychology*, 29, 87–109.
- Markman, K. D., & McMullen, M. N. (2003). A reflection and evaluation model of comparative thinking. *Personality and Social Psychology Review*, 7, 244–267.
- Markman, K. D., McMullen, M. N., & Elizaga, R. A. (2008). Counterfactual thinking, persistence, and performance: A test of the reflection and evaluation model. *Journal of Experimental Social Psychology*, 44, 421–428.
- Mazar, N., Amir, O., & Ariely, D. (2008). The dishonesty of honest people: A theory of self-concept maintenance. *Journal of Marketing Research*, 45, 633–644.
- Miller, D. T., & Gunasegaram, S. (1990). Temporal order and the perceived mutability of events implications for blame assignment. *Journal of Personality and Social Psychology*, 59, 1111–1118.
- Morris, M. W., & Moore, P. C. (2000). The lessons we (don't) learn: Counterfactual thinking and organizational accountability after a close call. *Administrative Science Quarterly*, 45, 737–765.
- Mussweiler, T. (2003). Comparison processes in social judgment: Mechanisms and consequences. *Psychological Review*, 110, 472–489.
- Pillutla, M. M., & Murnighan, J. K. (1995). Being fair or appearing fair: Strategic behavior in ultimatum bargaining. *The Academy of Management Journal*, 38, 1408–1426.
- Pillutla, M. M., & Murnighan, J. K. (2003). Fairness in bargaining. *Social Justice Research*, 16, 241–262.
- Reichert, L. K., & Slate, J. R. (2000). Reflective learning: The use of "if only..." statements to improve performance. *Social Psychology of Education*, 3, 261–275.
- Ritov, I. (1996). Anchoring in simulated competitive market negotiation. *Organizational Behavior and Human Decision Processes*, 67, 16–25.
- Ritov, I., & Baron, J. (1995). Outcome knowledge, regret, and omission bias. *Organizational Behavior and Human Decision Processes*, 64, 119–127.
- Roese, N. J. (1994). The functional basis of counterfactual thinking. *Journal of Personality and Social Psychology*, 66, 805–818.
- Roese, N. (1997). Counterfactual thinking. *Psychological Bulletin*, 121, 133–148.
- Roese, N. J., & Olson, J. M. (2007). Better, stronger, faster: Self-serving judgment, affect regulation, and the optimal vigilance hypothesis. *Perspectives on Psychological Science*, 2, 124–141.
- Schweitzer, M. E., & Hsee, C. K. (2002). Stretching the truth: Elastic justification and motivated communication of uncertain information. *The Journal of Risk and Uncertainty*, 25, 185–201.
- Shalvi, S., Handgraaf, M. J. J., & De Dreu, C. K. W. (2011). Ethical maneuvering: Why people avoid both major and minor lies. *British Journal of Management*, 22, S16–S27.
- Snyder, M. L., Kleck, R. E., Strenta, A., & Mentzer, S. J. (1979). Avoidance of the handicapped: An attributional ambiguity analysis. *Journal of Personality and Social Psychology*, 37, 2297–2306.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131.
- Zeelenberg, M., van Dijk, W. W., van der Pligt, J., Manstead, A. S. R., van Empelen, P., & Reinderman, D. (1998). Emotional reactions to the outcomes of decisions: The role of counterfactual thought in the experience of regret and disappointment. *Organizational Behavior and Human Decision Processes*, 75, 117–141.