



UvA-DARE (Digital Academic Repository)

Ethical decision making: on balancing right and wrong

Shalvi, S.

Publication date
2011

[Link to publication](#)

Citation for published version (APA):

Shalvi, S. (2011). *Ethical decision making: on balancing right and wrong*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter Three

Justified Ethicality: Observing Desired Counterfactuals Modifies Ethical Perceptions and Behavior

An article based on this chapter is in press: Shalvi, S., Dana, J., Handgraaf, M.J.J & De Dreu, C.K.W. (2011). Justified ethicality: Observing desired counterfactuals modifies ethical perceptions and behavior. *Organizational Behavior and Human Decision Processes*.

Morality, like art, means drawing a line somewhere.

Oscar Wilde (1854-1900)

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 (Becker, 1968; Alingham & Sandmo, 1972). 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, et al., 2011a). Mazar, Amir and Ariely (2008) explain this 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., 2011a; based on Fischbacher and 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; Shalvi, et al. 2011a; Mazar, et al. 2008; Gino, Ayal & Ariely, 2009; Gino, Norton & Ariely, 2010) 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 also find that the extent to which people allow themselves to lie depends on the availability of self-justifications. Participants who were allowed to roll three times to verify that the die was legitimate apparently lied more by reporting the largest of the three rolls, even though they knew that these additional rolls were not supposed to "count" for determining pay. Importantly, our paradigm insures that these justifications are completely private; no one else knows what numbers the participant has seen. Additional studies corroborate this behavioral result by showing that of all possible profit boosting possibilities, participants are most likely to use one of the additional rolls as means to increase their pay. Addressing the psychological driver for this

behavioral pattern, we further found 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. This allowed participants “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 the 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, and 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 cubicle over the other, a majority avoided the handicapped person. This finding further suggests that when people can justify their decision by their preference to 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. 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 one day more than six months instead of one 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 when lying implies that honesty is not perceived as a sharp contrast between lying versus being honest, but rather as a continuum that stretches between these two

ends. 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. Indeed, considering upward counterfactuals "that are evaluatively better than actuality" (Roese, 1997; p. 134), influences people's decisions and behaviors (Ritov & Baron, 1995; Roese, 1994; Reichert & Slate, 2000; Morris & Moore, 2000; Markman, McMullen & Elizaga, 2008; for review, Epstude & Roese, 2008). Markman and McMullen (2003) proposed that when considering such upward counterfactuals people engage in reflective thinking – or "as if" thinking. In this mind set people reflect on the counterfactual "as if" it was real, thereby reducing the contrast between this piece of counterfactual information and factual reality (see also Kahneman & Varey, 1990). People who observe 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. We thus suggest that when an observed counterfactual is desirable compared to factual reality, dishonestly using this counterfactual to boost profit may feel like less unethical compared to using other counterfactuals that were not observed.

Overview and Predictions

In Experiment 3.1 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. Participants in Fischbacher and Heusi's studies rolled the die once to determine pay and then more times to verify that the die was legitimate (also ensuring that no one would be able to know the outcome of the first roll determining pay). Across 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 envelop), 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 this paradigm that was not investigated before, namely the irrelevant rolls used to verify the die's legitimacy and not to determine pay. If indeed reporting a desired counterfactual 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 the possibility that Fischbacher and Heusi's (2008) pattern of results stems from the fact that people saw more than one die roll, we adapted their paradigm to eliminate these extra rolls. Holding all other aspects of the paradigm constant, we introduced a condition we call, single-roll and contrast it to a multiple-rolls condition which is similar to the one used before. 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 more than one roll (even though these extra rolls are not meant to determine pay). This is 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. As reports using the die rolling paradigm were private, in Experiment 3.2 we asked people to indicate what

outcomes they would have reported when rolling different combinations of outcomes. Doing so we assessed if indeed, people over-report the highest of these rolls more than any other profit boosting option available to them.

Experiments 3.3 and 3.4 were designed to provide support for 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 an 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 justification and thus is more likely to be used. Specifically, we tested the following predictions.


Hypothesis 2: Lying by reporting the observed desired counterfactual is considered less unethical than a lie that differs from the observed desired counterfactual.

Hypothesis 3: The modification of ethical perceptions mediates the relationship between observing a desired counterfactual and the likelihood of lying.

Experiment 3.1

Participants rolled a die and earned money according to what they reported rolling (Fischbacher & Heusi, 2008; Shalvi, et al., 2011a; see similarly Lammers, Stapel & Galinsky, 2010). 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 Figure 3.1). The first roll was to determine pay equivalent in US Dollars to the number they reported. Participants were also instructed to roll at least two more times to ensure that the die was legitimate. This also assured them that their first roll could not be observed 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.

Figure 3.1 Die under cup



Note to Figure 3.1. 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.


Participants and procedure. 129 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 to 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 all the dice placed in a box and could test that they were legitimate. Participants then individually received a die-under-cup and were instructed 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 experiment before participants began the actual task.

Results

If indeed 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 3 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),

see Figure 3.2. 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 $Z = 1.93$, $p = .05$.

Figure 3.2. Percentage of reported rolled outcomes as a function of condition (multiple-rolls vs. single-roll, Experiment 3.1)




Note to Figure 3.2. The dashed line represents the honest distribution predicted by chance (16.67% per die face). The error bars represents the 95% confidence interval around the proportion.

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 3 rolls (see Figure 3.3), 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 indeed 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.

Figure 3.3. The theoretical distribution of choosing the highest number observed in three die rolls.



Experiment 3.2

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 3 rolls they observed (which was the minimal number of times they were instructed to roll).

The method employed in experiment 3.1 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 3.2, we thus asked people what they would have reported in such ethically challenging situations. This allowed 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, Mage = 26.71) were recruited online and paid for participation. We employed a 6 (1st roll: 1-6) X 6 (2nd roll: 1-6) within subject design. Participants read instructions of the multiple-rolls conditions 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 they 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? Seventeen participants who did not answer these questions correctly were excluded from all analyses.

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.

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 binominal test revealed that this proportion was higher than the 20% predicted by our null hypothesis (e.g., 2, 3, 4, 5, or 6), $p < .0001$. 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).

Table 3.1. Frequency of estimated reports as a function of the two observed die rolls (Experiment 3.2)

		Report?						
Rolled	1	2	3	4	5	6	Total	
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	

Note to Table 3.1. For brevity considerations, Table 1 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 within a double border and justified lies appear in bold.

Experiment 3.3

Results of Experiment 3.2 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 3.1, 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 3.2 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 3.3 and 3.4, 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 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 3.3 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. 115 people (44% females, M age= 29.98) were recruited online and paid for participation. We employed a 6 (1st roll: 1-6) X 6 (2nd roll: 1-6) X 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. 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). 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). We defined "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).

Comprehension checks. As in Experiment 3.2, 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.

Results

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 ($M = 5.72, SD = 1.47$), $F(1, 82) = 189.54, p < .001, \eta^2 = .70$.

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 ($M = 5.87, SD = 1.47$), $F(1, 82) = 10.52, p = .001, \eta^2 = .11$.

Experiment 3.4

Results from Experiment 3.3 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 3.4, 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). This manipulation allowed us to test the mediation hypothesis (H4) suggesting that the influence of observing a desired counterfactual on people's likelihood to lie will be mediated by their modified ethical perceptions.

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) X 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. 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, $F's > 14.79, p's < .001$), see Figure 3.4.

Mediation analyses. 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.

Figure 3.4. 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 3.4)



Note to Figure 3.4. Error bars are +1 SE around the mean.

Experiment 3.4 Discussion

Experiment 3.4 provided further support for hypotheses H1 and H2: Students evaluated an untruthful report as less unethical as a function of observing (or not) a desired counterfactual. They further indicated that they were more likely to lie when observing such desired counterfactuals. Importantly, Experiment 3.4 also provided support for the mediation hypothesis (H3) demonstrating 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.

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 allowing 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 3.1, 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; Shalvi, et al. 2011a; Mazar, et al. 2008; Gino, et al., 2009; Gino, et al., 2010). However, we also show that beyond the magnitude of lying, people rely critically on justifications to enable themselves to lie, even when the justifications are only for themselves. This conclusion stems from three pieces of evidence. First, in Experiment 3.1, the distribution of reported outcomes did not differ than 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 3.2, 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 3.3 and 3.4 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 3.4 we gained support to the notion that the modification in one's perception of ethicality partially mediates the relation between observing a desired counterfactual and lying.

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 3.1, 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., Mazar, et al., 2008; Greene & Paxton, 2009). 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 3.1 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 3.2 helped in that they clarified that indeed a significant portion of people indicating that they would lie choose 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

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., 2011a) 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 incrementally 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 by Bazerman and colleagues (Bazerman & Banaji, 2004; Chugh, Bazerman & Banaji, 2005) proposing that implicit biases in people's awareness systematically prevent them from identifying other's 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 other's misconduct, leading to an ethical slippery slope, may be due to the extent to which the other's 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.

Concluding thoughts

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. In such, people seem to drive value from self-justifications and lie for money as long as they can feel honest.

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.