



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 Two

Ethical Maneuvering: Why People Avoid both Major and Minor Lies

This chapter was published as Shalvi, S., Handgraaf, M.J.J. & De Dreu, C.K.W. (2011a). Ethical maneuvering: Why people avoid both major and minor lies. *British Journal of Management*, 22, s16-s27.

Both within and outside organizations, people face ample opportunities to maximize personal profit by bending the rules and lying. Freelancers filing tax reports may be tempted to underreport income, employees may be tempted to over-report work-related expenses back to their employer, and salespeople may consider not mentioning hidden defects in their products. In fact, most of us are able to recall an instance where we bent the rules in our own favor, where we lied and deceived and got away with it. Early work on crime and punishment suggests that we perform such unethical behaviors as a function of an external cost benefit analysis: Lying in this view is the outcome of consciously balancing the potential profit that will be generated from the lie against the likelihood of getting caught and the magnitude of the punishment that might follow (Becker, 1968; for a similar argument in the literature on tax evasion see, Alingham & Sandmo, 1972).

Interestingly, however, and in contrast to this idea, most of us are also able to report instances in which we “could get away with it,” yet decided to follow an honest course of action – we acted ethically and morally appropriate at a cost to our self interest. Such introspection is at odds with the idea that unethical behavior is purely and solely the product of external considerations such as the likelihood to get caught and the punishment that might follow. It is, however, quite in line with psychological science showing that when making judgments and decisions, people are also driven by moral considerations and sentiments (e.g., Haidt, 2007; Bowles, 2008). In the current research, we proceed on the basis of the general idea that even under conditions of complete anonymity, when the likelihood to get caught is eliminated, people sometimes act honestly, and sometimes do lie and deceive, that they sometimes act morally, and sometimes bend the rules to serve their immediate self-interest. We were specifically interested in “ethical maneuvering” – the tendency for people to compromise between their desire to be honest and morally appropriate on the one hand, and serving their (financial) self-interest on the other. We develop and test the proposition that such ethical maneuvering manifests itself in a tendency for people, even under complete anonymity, to neither lie “all the way,” nor to engage in minor lies. We thus move away from the assumption that people balance external factors such as the likelihood to be caught and the punishment that will follow (Becker, 1968), and instead propose that ethical maneuvering is an internal cost-benefit analysis, which results in seeing some lies as more legitimate than others (for similar approach see Gneezy, 2005; Mazar, Amir & Ariely, 2008; Gino, Ayal & Ariely, 2009).

Lying, Deception, and Self-Concept Maintenance

Ethical behavior is broadly defined as behavior that is “both legal and morally acceptable to the larger society” (Jones, 1991, p. 367). Such a broad definition contrasts the moral act of being honest with such socially unacceptable behaviors as lying, deceiving and stealing. Whereas considerable work on lying and deception has considered interpersonal exchange (e.g., bargaining and negotiation; Boles, Croson, & Murnighan, 2000; Koning, Van Dijk, Van Beest & Steinel, 2010; Moran & Schweitzer, 2008; O’Connor & Carnevale, 1997; Schweitzer, Hershey & Bradlow, 2006; Steinel & De Dreu, 2004), the present work considered (dis)honesty in individual settings where the social system – institutions such as a government or work organization – rather than an interaction partner is taken advantage of (e.g., Schweitzer, Ordóñez & Douma, 2004). As such, our work squarely fits the behavioral ethics approach exploring the circumstances that lead individuals from all walks of life to commit unethical behavior (De Cremer, 2009). We seek to contribute to psychological theory about honesty and morality (e.g., Haidt, 2007; Bazerman & Banaji, 2004; Chugh, Bazerman & Banaji, 2005), and to provide advice to managers and employees seeking “accurate assessment of their best beliefs, careful considerations of diverse social effects, plus conscious adjustments to counteract our natural biases” (Gibson & Murnighan, 2009, p. 287).

Within the psychological and decision sciences, a growing stream of research indicates that people’s tendencies to deviate from ethical and honest behavior toward committing lies is rather easily influenced. For example, promoting a belief in free will reduces lying (Vohs & Schooler, 2007), increasing a sense of anonymity by darkening the room increased deceitful tendencies (Zhong, Bohns, & Gino, 2010), and uncertainty about the real world state of affairs increases the likelihood that people will bend the rules in their own favor (Schweitzer & Hsee, 2002).

No matter how interesting and important these and related studies are, they do not provide much insight into ethical maneuvering – what types of lies do people use or avoid, and to what degree do people engage in lying and deception (for an exception see, Steinel & De Dreu, 2004). However, a possible answer to these questions derives from the well-established notion that people are motivated to adopt, maintain and protect a view of themselves that is generally positive, conveys to them an image of an honest, fair, competent, and likeable person (Baumeister 1998; Bem, 1972). People accordingly affirm their self-view through a variety of cognitive and behavioral tactics and strategies including promotion and enhancement of the self (Sedikides & Strube, 1997). The habitual ability or inability to enhance one’s self image and feel generally good

about oneself is reflected in someone's self-esteem—an attitude towards oneself that associates with numerous self-beliefs pertaining to the self as a whole, or to its particular attributes (Sedikides & Gregg, 2003), in a tendency to overvalue one's possessions (i.e., mere ownership, Beggan, 1992; De Dreu & Van Knippenberg, 2005), an increased sense of entitlement (Thompson & Loewenstein, 1992; Van Dijk & De Cremer, 2006), and overly optimistic assessments of one's own skills and abilities (De Dreu, Nauta, & Van de Vliert, 1995; Shalvi, Shenkman, Handgraaf & De Dreu, 2011; Svenson, 1981; Van Lange & Rusbult, 1995; for review, Dunning, Heath & Suls, 2004).

To protect one's positive self-view, people tend to attribute failure externally and successes internally (Nisbett & Ross, 1991) and, importantly, they may shy away from unethical and morally inappropriate behavior that, however beneficial materially, may hurt one's social standing and private self-image. This latter notion is captured in Self-Concept Maintenance Theory advanced by Mazar and colleagues (2008). It proposes that people value honesty and try to maintain this aspect of their self-view. Thus, the aversive notion of acknowledging that one's own behavior is unethical prevents people from acting in ways that will force them to update their otherwise honest self-view. As long as one's behavior does not require moving from seeing oneself as honest to seeing oneself as a dishonest person, people may lie for profit. But, whenever behavior requires a change in self-view from being decent and honest to being less decent and dishonest, people refrain from lying for profit. This points to what we termed ethical maneuvering: people navigate between the materialistic (e.g., money) and psychological (e.g., increased status if not caught) benefits the lie provides with its materialistic (e.g. punishment) and psychological (e.g. inability to see oneself as an honest) costs. In the current work we focus on the most stringent test of ethical maneuvering, by studying if people lie for money in strictly anonymous settings where the tradeoff is between material gain and the private self-concept.

The Psychological Cost of Major Lies

Initial tests of the basic idea that people balance material gain and psychological cost of lying and dishonesty focused on the tendency of people to avoid lying in interpersonal settings. In a classic study on the role of consequences on deception, Gneezy (2005) found that people were willing to forego personal profit to avoid deceiving another person. The likelihood with which participants deceive their partner not only varied as a function of the personal profit generated by the lie but also by the loss it would yield the partner. On the basis of this and related results, Gneezy proposed that people may be rather lenient towards lies that are directed at large institutes (e.g.,

tax authorities) but not towards lies that are directed towards other individuals, as in the latter situation the consequences of dishonesty hurt a specific individual. However, even in the case of the lies that are not directed at other individuals, Mazar and her colleagues (2008) found that people avoid lying to a major extent but instead prefer using minor lies. Mazar et al. had participants take a short quiz composed by twenty matrices in which participants had to find in each matrix the two numbers that add up to exactly ten (e.g., 3.48 and 6.52). Importantly, participants were either graded by the experimenter or graded their own performance and received pay accordingly. Thus, participants who graded their own work could gain money by lying. Five studies showed that participants with an opportunity to lie reported solving more correct matrices than their peers who were graded objectively. Importantly, this over-reporting was not done to the full extent possible. The authors showed that the higher mean number of correctly solved items among participants who graded themselves was not the result of a few bad apples lying all the way, but rather by an upward shift in the distribution of outcomes by two to three items. Put differently, participants who had the opportunity to lie did not report solving correctly all 20 quiz items but rather they typically over-reported their performance by stating that they solved two to three items more compared to participants in the control condition who could not lie.

Based on these and other findings, Mazar et al. (2008) proposed that people use small lies to benefit themselves financially because these do not require changing one's self-view as an honest person, which major lies would. This tendency to avoid major lies is consistent with other work using different paradigms. Specifically, Fischbacher and Heusi (2008) employed a die rolling paradigm allowing people to report the outcome of a private roll to receive pay. Comparing the distribution of outcomes to the distribution predicted by chance, results were in line with the self-concept maintenance idea – a significant amount of people over reported the second highest option that was available to them, even when they might as easily have reported the highest option.

The Psychological Cost of Minor Lies

Whereas the above reveals that people avoid major lies and prefer to settle for a smaller lie they can get away with psychologically, a critical assessment of the magnitude of the lies students used in Mazar et al. (2008) provides another insight concerning the psychological cost of lying. While people are obviously aversive to major lies, they also seem aversive of very minor lies that do not yield sufficient payoff. Participants typically lied for around 15% of the maximal amount of lying possible (2-3 items of the maximum 20; specifically, 6.7% of the total in Study 1; 13.5% in Study 2;

16.5% in study 3; 14.8% in study 4). However, comparing the average size of the lie to the performance of those students who could not lie (answering on average between 3 to 4 items correctly), results reveal that participants were unethically boosting their earnings by over-reporting their achievements by at least 35%, moving from around 3 correctly solved answers to anywhere between 4 and 7 correct answers (specifically, 35% in study 1; 66% in study 2; 77% in study 3; 59% in study 4). While the evidence clearly shows that people avoid major lies, it also seems to indicate that a sufficient amount of profit is needed to bring about unethical behavior. Put differently, we propose that to maintain an honest self-view, people not only avoid lying a lot (forcing a self-view modification due to the major lie) but also avoid lying for very little. Lying in and of itself is psychologically costly and people accept such a psychological cost only if they receive sufficient material profit to compensate for it.

The Current Study

We propose that maintaining an honest self-concept not only requires avoiding major lies (Mazar, et al., 2008), it requires the avoidance of minor lies as well. The boundaries of lying are thus determined by the requirement that outcomes resulting from lying have to be sufficiently large to warrant a lie (the lower boundary), and the requirement that the lie cannot be so large that it becomes problematic to maintain a positive self-concept (the upper boundary). In other words, people will prefer to stay within the 'comfort zone' between these boundaries when considering a lie. We propose that people will only lie when such a comfort zone is present, and will avoid lies if this is not the case, i.e., when the only possible lies are either minor or major. To test this proposition, we employed a paradigm that ensured complete anonymity and allowed participants to lie and increase their pay. Importantly, we experimentally manipulated the range of possible lies available to them to include (or not) a middle-of-the-road lie that is neither major nor minor. By eliminating the opportunity to use a middle-of-the-road lie we confronted participants with a situation which allowed them to lie for profit but forced them to do so by either lying all the way or for very little—it made ethical maneuvering more difficult and if lying is indeed psychologically costly, this limited set of options should reduce lying.

We used a modified version of a die rolling paradigm (Fischbacher & Heusi, 2008). Participants were given an opportunity to anonymously report the outcome of a die roll and receive payment accordingly. As participants were the only ones to know the actual outcome of the roll, they could lie and increase their payoff. To insure complete anonymity we placed the die under a paper cup with a tiny hole in the bottom

through which participants could check their outcome, thus eliminating the possibility that anyone but participants may secretly observe what they rolled. We call this paradigm the “anonymous die-under-cup paradigm”.

To test whether people only avoid major lies or alternatively both minor and major ones, we gave participants the choice between exiting the situation for a fixed amount and anonymously rolling the die (see similar exit methodology in Dana, Cain & Dawes, 2006). We created two conditions: In the first, lying involved a choice between a minor lie (i.e., a lie yielding a little more than the exit option) and a major lie (i.e., a lie yielding the maximum outcome). In the second situation, we added an intermediate lying possibility: a lie yielding more than the minor lie, but not necessitating lying to the maximum extent. We predict that in situations in which the opportunity to increase profit includes only minor or major lies people will lie less compared to situations in which an intermediate option exists. In the latter case, this intermediate option will be used to increase profit. We call this the cost of lying hypothesis (H1).

Justifying lies. According to the cost of lying hypothesis we propose that when the monetary gain of a given lie is sufficient and surpasses the cost that might be experienced by the act of lying, people are likely to lie as long as they can still justify (to themselves) their deviation from ethical behavior. But how would these individuals subsequently justify their unethical behavior? Would they explain the outcome they reported rolling in terms of a desire to increase personal profit? Would they confess that their report was morally dubious or perhaps try to cover-up their unethical action and promote their decision as morally just? Because reports in our design were intentionally private, we could not test whether such post-hoc justifications may statistically mediate participants’ reports (as we do not know if an individual reporting a certain outcome lied or not). However, we can compare the mean levels of reported justifications (desire to maximize profit and moral appropriateness) between the different experimental conditions to better understand what types of justifications people in these settings use. We suspect that participants who will have a middle-of-the-road option to lie, which is neither minor or major, will not only lie more than those who do not have such option, but will also use more post-hoc justifications for their (generally dishonest) behavior (be it a desire to maximize profit, a reinterpretation of their unethical behavior as morally appropriate, or both). We call this the justification hypothesis (H2).

Note that assessing die reports of participants who choose to roll may potentially create a self-selection problem into the one of the two experimental

conditions. To address this possibility, we added a control condition in which lying was not possible. In this condition participants who chose to roll, were instructed to call the experimenter to observe their roll determining pay. It is likely that a more profitable exit option will attract more people to opt for that option. Similarly, having an opportunity to lie may attract more people to choose to roll compared to a condition in which no such opportunity exists. Importantly however, to rule out the possibility that selection issues may differentially influence the reporting of outcomes in the two conditions in which participants have a choice to roll the die or accept a fixed amount, we expect that the effect of such an exit option on deciding to roll will not be moderated by having an opportunity to lie.

Experiment 2.1

Participants. One-hundred ninety students (74% females; age $M = 21.97$, $SD = 4.05$) participated in the current study. We employed a 2 (exit option: high vs. low) X 2 (opportunity to lie: yes vs. no) between subjects factorial design. Participants were randomly assigned to conditions and received payment according to their actual decisions.

Procedures and Materials. Participants were seated in a private room. Task instructions were presented on a computer. They learned that they would receive the opportunity to earn money by rolling a six-faced die and reporting the outcome. Reporting rolling an outcome of 1 yielded €1, a 2 yielding €2, etc. However, reporting 6 yielded €0. The expected value of the game thus is $(1+2+3+4+5+0)/6 = €2.5$. To ensure anonymity we placed a paper cup above the die with a small hole enabling only participants to see the die. This low tech anonymous die-under-cup paradigm assured participants that no one but they themselves would be able to see their actual roll.

Participants were instructed to shake the cup once, check the outcome and type it on the computer screen. To assure participants that the die was not fixed or manipulated, they were told that after the first roll (used to determine payoff) they could roll the die as often as they wished. This also assured participants that no one would be able to check the obtained outcome after they left the room. Because the participants were the only ones to know the true value of the first die roll, they could lie about the outcome and increase their payoff, or they could be honest. Most critical to our design these participants could lie to different degrees allowing them to lie to minor or major degrees but importantly also to an intermediate degree.

To test the cost of lying prediction, before rolling we offered participants the option to avoid the die roll and instead accept an exit payoff. Depending on condition, this exit payoff was either €3.5 or €2.5. In the €3.5 exit condition, participants who choose to roll, foregoing a €3.5 exit option, can improve their outcome relative to the exit option only if they report a “4” or a “5” (“6” yields €0). This means either lying to a major extent (reporting a “5” to receive €5), or by lying for a minor extent (reporting “4” to receive €4), thus increasing their relative profit by €0.5. In the €2.5 exit condition participants may also lie to a major extent (“5”) or a minor extent (“3”; marginally increasing profit by €0.5) but can also pick a medium lie (“4”), see Figure 2.1. If people only avoid major lies, they should report “4” to a similar extent in both exit value conditions. However, if the act of lying per-se is psychologically costly, people will report “4” more often in the €2.5 (vs. €3.5) condition in which “4” is an intermediate (rather than minor) lie. Thus participants in both conditions could be honest, and most critical for present purposes, they could lie to a major extent (report a “5”) or engage in minor lying (i.e., report a “3,” in the €2.5 exit condition or “4” in the €3.5 exit condition). In the €2.5 exit condition participants could also lie by using a moderate lie (i.e., report a “4”).

Figure 2.1. Type of available lies (minor, intermediate and major) according to condition

Type of lie	Outcome	Type of lie
	1 = €1	
	2 = €2	
	----- €2.5	
	3 = €3	Minor lie
----- €3.5	-----	
Minor lie	4 = €4	Intermediate lie
Major lie	5 = €5	Major lie
	6 = €0	

As mentioned above, to address self-selection issues, we manipulated whether participants had the opportunity to lie. This was done by instructing participants who


chose to roll in the no opportunity to lie condition, to call the experimenter into the room before rolling and after rolling the die, revealing the outcome to the experimenter by lifting the paper cup and typing the obtained outcome on the computer. After reporting the outcome, participants were asked to indicate to what extent their outcome report was influenced by their “desire to earn as much as possible” and to what extent “reporting X [indicating the value participants reported rolling] was morally appropriate” (1 - not at all, to 7 - very much).

Results

Roll or Not? A logistic regression with exit option (€2.5 vs. €3.5) and opportunity to lie (yes vs. no) predicting the decision to roll, revealed a main effect for exit option, $B = -1.91$, Wald's $\chi^2 (1, N = 190) = 11.74, p = .001$. Thirty eight percent of the participants in the €3.5 exit condition chose to roll while in the €2.5 exit condition, 79% chose to do so. The main effect for the opportunity to lie did not reach statistical significance, $B = .73$, Wald's $\chi^2 (1, N = 190) = 3.49, p = .06$. The trend revealed that when there was an opportunity to lie anonymously, more participants chose to roll (62%) than when no such opportunity existed (47%). Importantly, ruling out the possibility that selection into these conditions was affected differentially, the interaction between exit option and opportunity to lie did not approach significance, $B = .00$, Wald's $\chi^2 (1, N = 190) = .00, ns$.

Lie or Not? Results concerning lying behavior focus on participants who had an opportunity to lie. Following the methods of Fischbacher and Heusi (2008) results supported the cost of lying hypothesis (H1): participants in the €3.5 exit condition reported outcomes that did not deviate from the uniform distribution expected by chance, Kolmogorov-Smirnov $Z = 1.2, ns$. These reports resulted in a mean payment of €3.00 per participant, an outcome that did not differ from the expected value of the game (€2.5), $t (24) = 1.55, ns$. In contrast, participants in the €2.5 condition reported outcomes that deviated from the uniform distribution expected by chance, Kolmogorov-Smirnov $Z = 2.19, p < .0001$. As can be seen in Figure 2.2, this deviation was due to 40% of participants reporting rolling a 4. This percentage exceeded chance level (16.67%; 95%CI = 22.47%, 57.53%) and resulted in a mean payment of €3.30 per participant. This is higher than the expected value of the game (€2.5), $t (29) = 2.81, p < .01$.

Figure 2.2 Distribution of reported die outcomes as a function of exit option (€3.5 versus €2.5)



Note: The horizontal dashed line represents the equal distribution (16.67%). The black bars differ from chance at significant levels (two-tailed): * $p < .05$, ** $p < .01$.

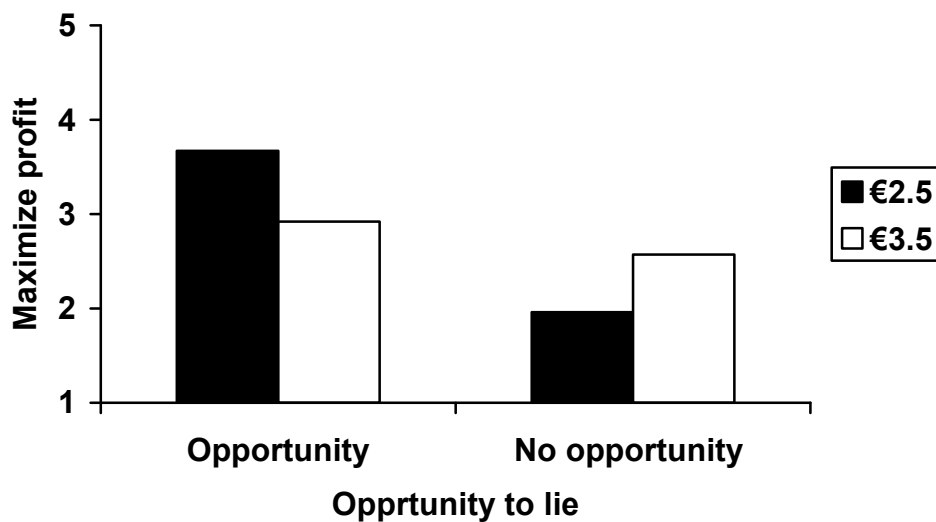
We tested whether the distributions of outcomes reported in the two conditions differed because of an inflation of “4” reports in the €2.5 condition. Reporting the intermediate option (“4”) was computed to hold the value of 1 in case of a “4” report and the value of 0 in case of reporting any other value. As predicted, a chi-square analysis revealed that the proportion of “4” reports in the €2.5 exit condition (40%) was higher than the proportion of “4” reports in the €3.5 exit condition (16%), $\chi^2(1) = 3.81, p = .05$.

Materialistic Justifications. Using a two-way ANOVA we tested the justification hypothesis (H2), suggesting that participants in the €2.5 condition will be more likely to use post-decision justifications compared to participants in the €3.5 condition, but only when the situation allows them to lie. Focusing on participants who chose to roll, a 2 (exit option: €2.5 vs. €3.5) X 2 (opportunity to lie: yes vs. no) X 6 (reported outcome: €0, €1, €2, €3, €4, €5) ANOVA was used to predict the justification of desiring to maximize profit. Results revealed that neither of the main effects reached significance; opportunity to lie ($F(1, 102) = 1.36, ns.$), exit option ($F(1, 102) = .19, ns.$), and reported outcome ($F(5, 102) = 1.95, ns.$). Similarly, the interaction between the exit option and reported outcome ($F(5, 102) = .51, ns.$) as well as the three-way interaction ($F(5, 102)$

= 2.11, $p = .07$) did not reach significance. However, as predicted (H2), the interaction between opportunity to lie and exit option was significant, ($F(1, 102) = 5.86, p < .025$).


As can be seen in Figure 2.3, participants in the €2.5 exit condition that had the opportunity to lie reported a higher desire to maximize their payoff ($M = 3.67, SD = 2.18$) compared to those that had no such opportunity ($M = 1.96, SD = 1.15$), $F(1, 102) = 10.89, p = .001$. Participants in the €3.5 exit condition, were not influenced by the opportunity to lie. Those who had the opportunity to lie expressed similar desire to maximize profit ($M = 2.92, SD = 2.53$) compared to those who did not have such opportunity ($M = 2.57, SD = 1.86$), $F(1, 102) = .30, ns$.

Figure 2.3. Desire to maximize profit as a function of exit option and opportunity to lie



In addition, the interaction between opportunity to lie and reported outcome was also significant, ($F(5, 102) = 3.55, p < .01$). As can be seen in Figure 2.4, while participants who had the opportunity to lie reported a higher desire to maximize their payoff according to their reported outcomes ($F(5, 102) = 6.73, p < .0001$), participants who did not have the opportunity to lie, did not vary on their desire to maximize profit as a function of their reported outcomes $F(5, 102) = .22, ns$. In the opportunity to lie condition, participants reported a higher desire to maximize profit the higher their reported outcome was (polynomial linear contrast, $F(5, 102) = 3.04, p < .0001$).

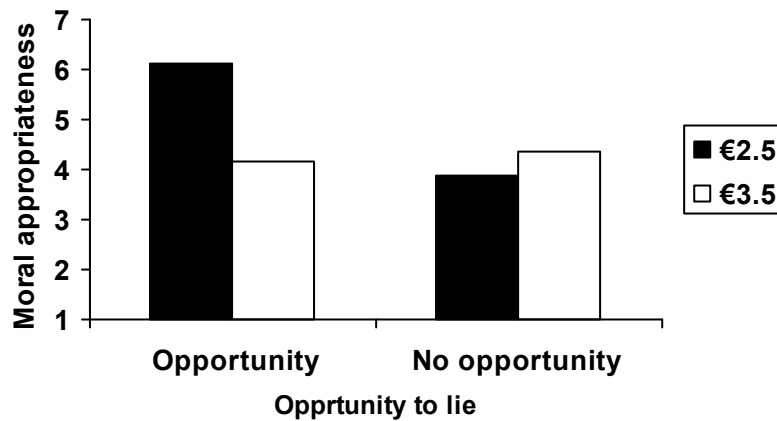
Figure 2.4. Desire to maximize profit as a function of opportunity to lie and reported outcome



Moral Justifications. To assess the extent to which participants ranked their report to be morally appropriate we used the same ANOVA with 2 (exit option: €2.5 vs. €3.5) X 2 (opportunity to lie: yes vs. no) X 6 (reported outcome: €0, €1, €2, €3, €4, €5) to predict moral appropriateness. Results supported the (moral cover-up) justification hypothesis (H2), revealing a main effect for the opportunity to lie, $F(1, 102) = 14.06, p < .0001$. Participants who had the opportunity to lie considered their report to be more morally appropriate ($M = 5.32, SD = 1.87$) compared to those who had no opportunity to lie ($M = 4.00, SD = 2.10$). Similarly, providing further support for the justification hypothesis, the main effect for exit option was also significant, participants in the €2.5 exit condition considered their report to be more morally appropriate ($M = 5.03, SD = 1.93$) compared to those in the €3.5 exit condition ($M = 4.27, SD = 2.21$), $F(1, 102) = 8.81, p < .005$. The main effect for reported outcome was not significant, $F(5, 102) = 2.07, ns$.

Results further revealed a significant interaction between exit option and opportunity to lie, $F(1, 102) = 4.57, p < .05$. Supporting the pattern predicted by the justification prediction (H2), participants who had the opportunity to lie in the €2.5 exit condition considered their report to be more morally appropriate ($M = 6.13, SD = 1.19$) than those in the €3.5 exit condition, ($M = 4.36, SD = 2.08$), $F(1, 102) = 12.35, p = .001$. For participants in the no opportunity to lie condition, results revealed no effect for the exit option that was provided to participants, $F(1, 102) = .87, ns.$, see Figure 2.5.

Figure 2.5. Moral appropriateness of reported outcome as a function of exit option and opportunity to lie



The interaction between reported outcome and opportunity to lie as well as the three-way interaction were not significant (both F 's $< 1.4, p$'s $> .1$). Finally, an unexpected interaction was observed between reported outcome and exit option, $F(5, 102) = 2.57, p < .05$. While participants in the €2.5 did not vary on their reported moral appropriateness according to their reported outcomes, $F(5, 102) = .93, ns.$, for participants in the €3.5 exit option condition, results revealed an effect for reported outcome, $F(5, 102) = 3.05, p < .05$. Post-hoc analyses revealed that this unexpected effect was due to a single comparison between participants who reported rolling 2 ($M = 1.83$) and those who reported rolling 3 ($M = 5.40$), $D = 3.56, p = .05$, which was of minor relevance for our current reasoning.

General Discussion

Early work on crime and punishment (Becker, 1968) predicts people in strictly anonymous settings, where the likelihood to get caught does not exist, to lie all the way and to maximum extent possible. Yet this is not what we found here – people avoid

major lies even in a completely private and anonymous setting. Theories based on these early works further predict that people will lie even if doing so only marginally increases their profits but again, this is not what people do: Even in a completely private and anonymous setting, and even after volunteering into the situation, people avoid lying for very minor profit. Rather, our results show that when unethically boosting personal profit people strive to maintain a positive self-concept and lie only when the materialistic gain exceeds the psychological cost caused by lying. People avoid both major and minor lies but not the moderate lies that provide a reasonable increase in profit without exerting too much damage on the honest self-concept. People sacrifice personal profit to maintain a positive self-concept but there clearly are boundaries to what such sacrifice entails.

Our findings expand the work by Mazar and colleagues (2008) on self-concept maintenance and deceitful behavior in two ways. First, building on previous work (Fischbacher & Heusi, 2008) we placed a die-under-cup, to examine lying in a purely anonymous setting. Second, and more important, we replicated earlier findings of self-concept maintenance theory and provided a boundary condition showing that people not only avoid major lies but also minor ones. People like to consider themselves not only as relatively honest but also require a relatively large incentive before they lie, even if the lie is untraceable. The mere act of lying seems to be psychologically costly.

The finding that lying is psychologically costly resonates with findings concerning deception in interpersonal setting (Gneezy, 2005). Gneezy 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). Our results corroborate this notion, indicating that people’s likelihood to lie for a given outcome (e.g., €4) are influenced by whether this outcome is an incremental versus a more significant increase in profit compared to an alternative these people just decided to forego. As such, our results demonstrate that even when another person’s outcomes are not influenced by the lie (or at least the other is an abstract research funding agency) the utility of a given monetary outcome may vary depending on whether it was attained by a minor, a middle-of-the-road, or a major lie.

The current work further reveals that when people are provided with ‘middle-of-the-road’ opportunities to lie, they not only use them but also feel obliged to justify their behavior. Specifically, we found that participants in the €2.5 exit condition who had an opportunity to lie reported that their roll reports reflected a desire to maximize profit but at the same time were morally appropriate. Knowing that their actual die

rolls are known only to themselves, people in this condition attempted to appear as if they were not the ones lying. Recent evidence point out to the important of such justifications. For example, Gino and colleagues (Gino, et al., 2009) manipulated the social norm concerning the legitimacy of lying while allowing people to lie regarding own success in an exam and getting paid accordingly. When the social norm was lenient toward unethical behavior, by having a confederate display an obvious lie and getting away with it, people lied. In contrast, when the confederate was an out-group member, dressed with a t-shirt of a competing school, people lied less. Without a justification to behave unethically (“this other person obviously lied, so why would not I?”) lying seems to be psychologically costly. Schweitzer and Hsee (2002) found that people are more likely to stretch the truth and lie when they may justify their behavior by their uncertainty about real world state of affairs. When the mileage of a car a person was supposedly going to sell was estimated by a narrow mileage range, thus decreasing uncertainty about the actual value, sellers lied less about the mileage compared to others who received a more uncertain estimation with a wider possible range. Similarly, Lundquist, Ellingson, and Johannesson (2009) examined the effect of cheap talk in bargaining games with one-sided asymmetric information. The seller had private information about her skill and was provided (or not) with an opportunity to communicate this information to a buyer through a written message. Results suggest that individuals have an aversion towards lying about private information and that the aversion to lying increases with the size of the lie and the strength of the promise. Finally, Steinel and De Dreu (2004) showed that such lie-aversion is especially strong among humans with chronic pro-social concerns, and is substantially reduced when the situation justifies lying and deception (e.g., when the partner is depicted as highly competitive).

Taken together, these pieces of evidence suggest that when people have no way to justify an unethical act lying is reduced. Sources of justifications may range from the social norm established by another person’s immoral acts (Gino, et al., 2009) to the level of uncertainty about reality (Schweitzer & Hsee, 2002), but importantly we propose, also by a sufficient profit that may be generated by the lie. People engage in ethical maneuvering, balancing the profits from lying against the (im)material costs associated with it, their positive self-concept being one major resource at stake when lying. Justifications and excuses assist in such ethical maneuvering as they exonerate the self and allows one to improve profit by bigger lies.

One limitation of our paradigm is that it does not allow us to assess if individuals who lied were the ones who used such justifications more often. However, on the aggregate level people who lied more were also more likely to justify their behavior. Some may consider the fact that we can not trace back individual behavior as a downside of our paradigm. Indeed, it prevents us from assessing statistically what factors may mediate the process of reporting dishonestly. We however maintain that our minimalistic approach actually allows a better study of unethical behavior. First, it was transparently clear to participants that their rolls were for their eyes only. Thus we study when people lie without lying to them about how we do so. Second, the low tech die-under-cup paradigm ensures participants, even the most suspicious of them, that no one will be able to know whether they lied or not. Indeed, a growing number of researchers studying ethical decision making adopt a similar approach focusing on assessment of unethical behavior between different situations rather than between specific individuals (e.g., Mazar, et al., 2008). While tracing back individual behavior may allow us to statistically assess mediating factors, it also requires us to deceive participants and figure out ways in which we learn about their actual decisions without their knowledge. Studying situational factors on the other hand, prevents us from assessing statistical mediation but increases our ability to create trustworthy and transparent methods in which participants' true values are tested.

Studying the behavior of people who actively choose to enter the situation opens up a potential self-selection issue that might influence the observed results and create difficulty in interpreting them. Acknowledging that fact, we included in our design control conditions in which people had to make the same entrance decision (take an exit of a certain value or roll the die) but importantly could not subsequently lie. Not surprisingly, results revealed that the higher the exit option is the more people opt for it. Similarly, people were somewhat more likely to enter situations that allowed them to subsequently lie than situations that did not. Importantly, the interaction between these two factors was not significant suggesting that people were not selecting themselves differentially into one of the two exit option conditions. Moreover, the obtained results can not be explained by risk-preferences. If risk preference drives entrance, it should impact the decisions of honest people (as liars do not mind entering and lying in either exit condition). This suggests that the (honest) people opting-out of the €3.5 condition due to risk aversion will inflate lying in this condition compared with the €2.5 condition. Results however, reveal that lying is decreased rather than increased, in the €3.5 condition, thus making this alternative interpretation very unlikely.

Conclusions and practical implications

Taken together, our results contribute to the field of behavioral ethics by demonstrating that otherwise honest people may act unethically when provided with the right opportunity. While these people knew to distinguish right from wrong (as suggested by their tendency to justify their unethical behavior) they still lied under the right circumstances. Understanding that most people may yield to temptation in some situations but less so in others, is crucial for organizational leaders, managers and employees. Letting go of the normative approach proposing what people ought to do and focusing on how humans actually behave, our work provides insight into such conditions and may hold an implicative prospect. It is timely to consider environmental designs that not only maintain freedom of choice but also motivate ethical behavior (Thaler & Sunstein, 2008). Given that people avoid both minor and major lies it would be interesting to think through system designs that provide people with certain degrees of freedom to lie and deceive. For example, reimbursing employees for work related expenses on a business trip can be done by either requesting them to report their expenses or alternatively providing them with per-diem allowance. In well designed systems providing calibrated alternatives (e.g., per-diem allowance) that narrow the range of lying and force people into either major or minor lies – honesty should flourish.