Uncertain lies: How payoff uncertainty affects dishonesty

Celse, J.; Max, S.; Steinel, W.; Soraperra, I.; Shalvi, S.

DOI
10.1016/j.joep.2018.09.003

Publication date
2019

Document Version
Final published version

Published in
Journal of Economic Psychology

Citation for published version (APA):
Celse, J., Max, S., Steinel, W., Soraperra, I., & Shalvi, S. (2019). Uncertain lies: How payoff uncertainty affects dishonesty. *Journal of Economic Psychology*, *71*, 117-125. https://doi.org/10.1016/j.joep.2018.09.003

Journal of Economic Psychology 71 (2019) 117–125

Contents lists available at ScienceDirect

Journal of Economic Psychology

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

Uncertain lies: How payoff uncertainty affects dishonesty

Jérémy Celse[b,*], Sylvain Max[b], Wolfgang Steinel[c], Ivan Soraperra[a], Shaul Shalvi[a]

[a] CREED, Faculty of Economics and Business, University of Amsterdam, Netherlands
[b] Burgundy School of Business, Université Bourgogne Franche-Comté, France
[c] Institute of Psychology, Leiden University, Netherlands

ARTICLE INFO

JEL classification:
D63

Keywords:
Dishonesty
Lies
Uncertainty

ABSTRACT

In this paper we experimentally explore how lying changes when its consequences are not certain. We argue that, when consequences are not certain, lying is morally less costly because the action of lying does not mechanically result in the obtainment of the benefit and this produces a lower feeling of responsibility in case the benefit is obtained. Moreover, we argue that the smaller the impact of lying on the probability to obtain the benefit the lower is the feeling of responsibility. We test our predictions using a modified die-under-the-cup task where misreporting, rather than delivering a higher payoff, increases the likelihood to get a prize. Overall we have four treatments where the reported outcome affects the probability to get a prize to a different extent. Contrary to our prediction, we do not observe any treatment difference suggesting that lying is independent to the extent to which it increases the probability to get a benefit. This result suggests that the willingness to lie to secure a benefit and the willingness to lie to marginally increase the probability to obtain a benefit are very similar.

1. Introduction

Lying is common in various circumstances and among people from all walks of life (De Paulo, 2004; De Paulo, Kashy, Kirkendol, Wyer, & Epstein, 1996). People lie for various reasons: to avoid hurting the feelings of someone else, to avoid paying taxes, to get hired, to win a competition, etc. In some of these cases the liar can secure a favorable outcome by lying. In others the lie does not secure the outcome and can, at most, increase the likelihood to obtain the outcome. Think of applicants submitting their curriculum vitae for a job or salesman/workers competing for a bonus. In these cases one can make up the CV or overstate the performance, i.e., lie about the number of sales or about the hours worked, but one does not have the certainty to achieve the goal. By lying one can only increase the likelihood to get the job or the bonus.

In this paper we look at lying when the final consequences of lying are uncertain. In particular, we look at how the moral cost of lying changes when lying does not secure a benefit. Are people tempted to lie to a higher extent to secure a benefit or to increase the probability of getting the benefit? Are people lying more when they can increase the probability of getting the benefit by a little or by a lot? We argue that the moral cost of lying changes with the feeling of responsibility for the final outcome. Specifically, we argue that, when lying permits to increase the likelihood of, but not to secure, the obtainment of a benefit, people are more prone to lying. This because lying leaves room for chance to determine the final outcome and, consequently, increases self-justification and reduces the feeling of responsibility. Additionally, we argue that the lower the impact of lying on the probability to get the benefit, the lower is the moral cost of lying. This because lying has a smaller impact on the chance to get the benefit and this increases the room for self-justification and reduces the feeling of responsibility even further.

* Corresponding author.
E-mail address: jeremy.celse@bsb-education.com (J. Celse).

https://doi.org/10.1016/j.joep.2018.09.003
Received 13 October 2017; Received in revised form 3 September 2018; Accepted 14 September 2018
Available online 15 September 2018
0167-4870/

In our experiment we find that, contrary to our predictions, lying is not affected by whether its final consequences are uncertain or not. Participants lie to the same extent when lying provides a sure benefit and when it simply increases the likelihood to obtain a benefit. Such result suggests for a stable cost of lying when moving from certain to uncertain environments and rules out the hypothesis that leaving room for chance to determine the final outcome may help reducing the feeling of responsibility for the obtainment of the benefit.

The paper is organized as follows: after discussing the relevant literature (Section 2), we illustrate the experimental design and outline our hypotheses (Section 3). Afterwards, we presents the results (Section 4), we discuss the main findings and the drawbacks of our design (Section 5) and we conclude (Section 6).

2. Related literature

The rational model of crime and punishment (Becker, 1968) assumes that an agent's decision to lie, or broadly speaking to engage in unethical behavior, depends, on the one hand, on the trade-off between the expected benefit derived from the action and, on the other hand, on both the probability of being caught and the sanction incurred when being caught.

Recent studies, however, have challenged this pure economic view. These studies highlighted the presence of moral costs of cheating and the importance of psychological variables in the economic trade-off of moral and financial costs and benefits (Abeler, Becker, & Falk, 2014; Ariely, 2012; Mazar, Amir, & Ariely, 2008; Shalvi, Dana, Handgraaf, & De Dreu, 2011). In particular, the decision to lie seems to be affected by the possibility to find a moral justification for it. Evidence shows that seeing the desired counterfactual (Shalvi et al., 2011), being exposed to social norms and other people's unethical behavior (Gächter & Schulz, 2016), and having a third party benefiting from lying (Conrads, Irlenbusch, Rilke, & Walkowitz, 2013) seem to increase its justifiability. Moreover, the propensity to lie is influenced by the goal that one needs to achieve: the need to win a tournament (Conrads, Irlenbusch, Rilke, Schielke, & Walkowitz, 2014; Schwieren & Weichselbaumer, 2010), the need to exert collaborative efforts (Weisel & Shalvi, 2015), and the need to protect oneself from losses (Grolleau, Kocher, & Sutan, 2016), they all increase the propensity to lie.

A common feature of the studies reported in the previous paragraph is that they adopt tasks where the monetary consequences of lying are certain. The most typical task employed in these studies is the die-under-the-cup paradigm (Fischbacher & Föllmi-Heusi, 2013). Within this paradigm participants privately roll a die, look at the outcome, and report the outcome to the experimenter that cannot verify whether the participant lied or not. Finally, participants get rewarded according to the reported outcome with higher reported outcomes translating into higher financial rewards (with the exception of reporting 6 translating to earning nothing). More generally, almost all the studies that look at the moral cost of lying either employ a setting equivalent to the die-under-the-cup, in which the outcome of a random device is privately observed and the outcome is reported to the experimenter, or employ a setting where participants can overstate their performances to gain more money without any possibility to get caught.[1]

Whereas there are many everyday situations in which one's lies deliver a certain benefit, there are circumstance where lying do not provide such certainty as well. These situations are overlooked by the current studies and little is known about how the moral cost of lying is shaped when the consequences of the lie become uncertain. In what perspective lying to obtain a certain outcome differs from lying to increase the likelihood of winning? How does the moral cost of lying change when the outcome is uncertain?.

A clear difference between lying to obtain a benefit versus lying to increase the likelihood of obtaining the benefit is the fact that in the former case the lie has sure consequences while in the latter case the lie does not necessarily have consequences. This may produce a different perception of the responsibility for the final outcome in the two setting and may provide a stronger justification for lying when the consequences of the lie are uncertain.

Evidence on delegation of opportunistic and selfish actions is suggestive that people may feel less responsible for unfair outcomes when there is no direct link between the action and the outcome. Bartling and Fischbacher (2012), for instance, provide evidence that people are more opportunistic when they can delegate their actions to an agent (and face uncertainty about the final outcome) compared to when they choose directly. This because delegation permits to shift the blame and to avoid responsibility. In particular, their treatments without punishment show that there is a higher fraction of selfish choices when people can either directly take the action or delegate compared to when delegation is not possible. This suggests that it is easier to make unfair choices through delegation than directly. Moreover, subjects that delegate are ready to accept some uncertainty about the final outcome (28% of the delegated agents choose the fair allocation) in order to reduce the moral cost of opportunist behavior. Hamman, Loewenstein, and Weber (2010) provide further evidence that people feel less morally responsible for unfair outcomes when delegation is possible compared to when the decision is made directly. This even though the choice of the agent they want to delegate the decision to plays a central role in producing the unfair outcome.

Finally, indirect evidence comes from the result that, when people face the choice between taking a prize for themselves or giving it to another person, they prefer to leave room to uncertainty and to choose a lottery to assign the prize rather than directly make the unfair choice of getting the prize (Brock, Lange, & Ozbay, 2013; Krawczyk & Le Lec, 2010). This may be viewed either as a taste for procedural fairness (Brock et al., 2013) or as avoidance of being responsible for the unfair outcome (Gordon-Hecker, Rosensaft-Eshel, Pittarello, Shalvi, & Bereby-Meyer, 2017) and it suggests that people are happy to leave room for chance and to give up control over the final outcome to feel less responsible for it.

[1] Within tasks falling into this second category, participants are either asked to guess the outcome of a random device (see, e.g., Shalvi, 2012) and then report if they guessed correctly or to self-assess their performances in a real-effort tasks and then destroy their answer sheet (see, e.g., Grolleau et al., 2016).

Similar results are observed when looking at papers studying delegation of cheating in sender-receiver games. In particular, Erat (2013) shows that a significant fraction of people employ an agent when they could lie themselves and, by doing so, they accept some uncertainty about the final outcome. In a similar setting, Kandul and Kirchkamp (2018) ask the senders to make both the decision of the signal to send and the decision whether to delegate or not. They find that senders report honestly about 21% of the times when directly choosing and those senders that report honestly decide to delegate about 50% of the times. This implies that the moral cost of lying is lower when the outcome is obtained indirectly (the senders have a 79% chance to lie when delegating).

All the mentioned results suggest that people feel less responsible for unfair/unethical choices when the outcome is not a mechanic consequence of their actions, this either because they can delegate to another person or because they prefer to leave the decision to chance. On a similar vein, we argue that people may lie more when the consequences of lying do not mechanically deliver a benefit but only increase the likelihood to get it. This because people can more easily justify the lie to themselves and feel less responsible because they let the final decision to chance. More generally, we argue that the smaller the impact on the likelihood of obtaining the benefit the easier is lying to justify. If this conjecture is correct, the propensity to lie should be higher when lying increases the probability by a little compared to when it increases the probability significantly. Note that Bartling and Fischbacher (2012) use a similar argument and propose a measure of responsibility for achieving an unfair outcome that is based on the impact that the agent's decision has on the probability to obtain the outcome.

To the best of our knowledge there are only a couple of studies looking at lying when consequences are uncertain: i.e., Shalvi (2012) and Dugar, Mitra, and Shahriar (2018). The first study (Shalvi, 2012) explores how people dishonestly increase the likelihood of winning a prize by assigning points to different pots. At the end, one of the pots is chosen at random and the amount of points in the pot are converted into lottery tickets to raffle a prize. In the experiment, the pots start with the same number of points but are subject to either losses or gains. Knowing the loss/gain of each pot, subjects have to choose the pot to which they want to add the points. Subjects add points by guessing the outcome of 20 coin flips and reporting if they guessed correctly. The experiment shows that subjects prefer to dishonestly add points to those pots where there are losses that can be offset. While in the experiment participants dishonestly increase the likelihood of winning, the focus of the paper is on the choice of the pot where to add the points. The paper does not explore how cheating changes when there is a low/high probability of winning. Moreover, the raffle introduces strategic considerations that are not present in our setting.

The second study (Dugar et al., 2018) explores lying in a sender-receiver game where the consequences for the counterpart are risky. The authors show that senders lie more when the consequences for the receiver depend on the outcome of a binary lottery compared to when consequences are deterministic and equal to the expected value of that lottery. They argue that uncertainty about the consequences for the receiver allows the sender to hold the self-serving belief of a favorable lottery outcome for the receiver. This study is different from ours in many dimensions: (i) it explores lying in a strategic setting where beliefs about the choice of the receiver play an important role; (ii) the uncertainty is on the receiver's payoff and not on the sender's payoff; and, most importantly, (iii) the impact of lying on the probability of getting the benefit does not change with uncertainty because the receiver ignores the payoffs of the game.

There are other studies that, looking at different research questions, implement uncertainty in the consequences of cheating. Studies looking at the effect of competition on lying, for instance, show that people cheat more when competing in winner-take-all tournaments compared to when they are paid a piece-rate (Faravelli, Friesen, & Gangadharan, 2015; Schwieren & Weichselbaumer, 2010). This result can suggest that competitive pressure increases cheating but it can also suggest that the uncertain consequences of lying in tournaments can provide a moral excuse for cheating. Moreover, there is evidence showing that, when the competitive pressure is either high or low, i.e., there is a high/low number of people that can obtain the prize, there is less lying compared to when the competitive pressure is at an intermediate level Cartwright and Menezes (2014). This because in the first case cheating does not increase much the likelihood to obtain the prize and in the second case the likelihood is already very high and cheating is unnecessary.[2]

All the studies implementing uncertain consequences of lying that have been mentioned so far are different from ours in the following aspects: firstly, they do not systematically compare how lying changes with its potential impact on the probability of obtaining the benefit and, secondly, they explore cheating in strategic settings where beliefs about other participants' behavior play a major role. In the next section we provide a detailed description of the experimental design and we outline how we test our conjecture that, the smaller the impact of lying on the likelihood of obtaining a benefit, the lower the feeling of responsibility.

3. Experimental design and procedures

To test our conjecture, we conduct a series of treatments based on the standard die-under-the-cup paradigm (Fischbacher & Föllmi-Heusi, 2013). Specifically, we consider a setting where participants privately roll a die, observe the outcome, and report the

[2] Note that there are experimental ways to implement incentives that introduce some uncertainty in the consequences of lying. Randomly selecting a subset of participants (see, e.g. Motro, Ordóñez, Pittarello, & Welsh, 2016; Shalvi, Eldar, & Bereby-Meyer, 2012) or a subset of periods for payment introduces uncertainty about the consequences of cheating, i.e., you can cheat but if you are not selected this has no consequences. These cases, however, are very different from our setting because the decisions of the participants do not impact the probability of obtaining the benefit but only the size of it. A different incentive scheme, which is used in psychology, is to raffle a prize where the payoff of participants is converted into lottery tickets for the prize (see, e.g., Pitesa, Thau, & Pillutla, 2013; Shalvi, 2012). This setting implements a tournament among participants that,by cheating, can increase their probability of winning the prize at the expenses of the others.

Table 1
Expected payoffs by reported outcome in the 4 treatments.

Reported	T_{NR}		T_6		T_7		T_{30}		Exp. pay
outcome	Prob.	Pay	Prob.	Pay	Prob.	Pay	Prob.	Pay	All T
1	1	1	$\frac{1}{6}$	6	$\frac{1}{7}$	7	$\frac{1}{30}$	30	1
2	1	2	$\frac{2}{6}$	6	$\frac{2}{7}$	7	$\frac{2}{30}$	30	2
3	1	3	$\frac{3}{6}$	6	$\frac{3}{7}$	7	$\frac{3}{30}$	30	3
4	1	4	$\frac{4}{6}$	6	$\frac{4}{7}$	7	$\frac{4}{30}$	30	4
5	1	5	$\frac{5}{6}$	6	$\frac{5}{7}$	7	$\frac{5}{30}$	30	5
6	1	6	$\frac{6}{6}$	6	$\frac{6}{7}$	7	$\frac{6}{30}$	30	6

outcome to the experimenter. In all our treatments, we adopt the multiple rolls procedure (Shalvi et al., 2011) and ask participants to roll the die three times knowing that only the first roll will determine their payoff. Participants are instructed to report truthfully but they can report an outcome that is different from the one they observed. The payoff of the participants depends on the reported outcome according to the following treatments.

- Baseline or treatment T_{NR}. This is the standard die-under-the-cup task. Participants receive an amount in euro equal to the reported number.
- Risk treatments. In this case participants have the opportunity to win either 0 or k euros. Participants win the k euros by drawing a blue ball from an urn containing k balls that can be either blue or red. Before drawing the ball, participants determine the urn composition by reporting the outcome of a die-roll. The reported outcome defines the number of winning balls (blue) in the urn. The remaining balls are losing balls (red). We consider three versions of the risk treatment where we manipulate the number k that defines both the total number of balls in the urn and the prize won when a blue ball is drawn. We consider treatment T_6 where $k = 6$, treatment T_7 where $k = 7$, and treatment T_{30} where $k = 30$.

In the baseline treatment, T_{NR}, there is no uncertainty about the monetary consequences of lying. Participants obtain from a minimum of 1 euro to a maximum of 6 euros depending on what they decide to report. The risk treatments, instead, differ by the number of balls in the urn and by the size of the prize. In treatment T_6, where there are 6 balls in the urn, certainty about the final outcome is still achievable by participants. If participants report a 6, the urn contains only blue balls and participants can obtain 6 euros for sure. In the remaining treatments, i.e., T_7 and T_{30}, where the urn contains more than 6 balls, certainty about the final outcome is not achievable. Here, by reporting a 6, participants can only maximize the probability of getting the prize. In T_7 they can obtain a maximum probability of 6/7 and in T_{30} a maximum probability of 6/30.

Note that misreporting has a different impact on the probability of winning in the three risk treatments. Specifically, the higher the number of balls in the urn, the lower is the impact on the probability of winning by misreporting. Misreporting by 1, for instance, increases the probability of winning by $\frac{1}{6}$, $\frac{1}{7}$, and $\frac{1}{30}$ in treatments T_6, T_7, and T_{30}, respectively. This would imply that incentives to lie are weaker in T_{30} than in T_6. To control for this, we decided to set different monetary prizes across the risk treatments. With the chosen parameters, the expected earnings of participants are the same in all treatments. This is true both in case of full dishonesty, where the expected earnings are 6 euros, and in case of truthful report, where the expected earnings are 3.5 euros. Additionally, the expected payoff gain when observing x and reporting r is exactly $r-x$ euros in all treatments (see Table 1). Therefore, by keeping constant the expected monetary incentive to misreport, we can interpret any difference in the distribution of reported outcomes as a change in the moral cost of lying due to the treatment manipulation.

Building on the argument that lying is easier to justify when the impact on the probability is lower because people feel less responsible for the final outcome, we predict more lying in T_{30} than in T_7 than in T_6. Additionally, we expect a sharp drop in maximal lying when moving from T_7 to T_6 because of the fact that lying to the full extent in the latter eliminates the likelihood of the bad outcome while in the former still leaves room for chance to determine the outcome.

The experiment was conducted at the Laboratoire d'Economie Expérimentale de Paris (LEEP) using a very heterogeneous subject pool of students and non students. After arrival, participants read the instructions, completed the die rolling task, and filled in a short questionnaire in which we collected non-incentivized beliefs about the number reported by other participants, age, and gender.[3] Additionally, one week before the experiment, participants completed an on-line version of the domain-specific risk-taking scale (DOSPERT) questionnaire (Blais & Weber, 2006).

Experimental sessions lasted about 1 h including payments. In total, 374 subjects participated to the experiment with 96 subjects in T_{NR} (mean age = 23.9, SD = 9.8, 56.8% female); 95 in T_6 (mean age = 31.2, SD = 14.4, 56.8% female); 92 in T_7 (mean age = 34.6, SD = 16.4, 68.5% female); and 91 in T_{30} (mean age = 30.7, SD = 14.1, 61.1% female). Data was collected in two waves: in June 2014

[3] Appendix A reports instructions of the die-rolling task translated from French for the treatments T_{NR} and T_6.

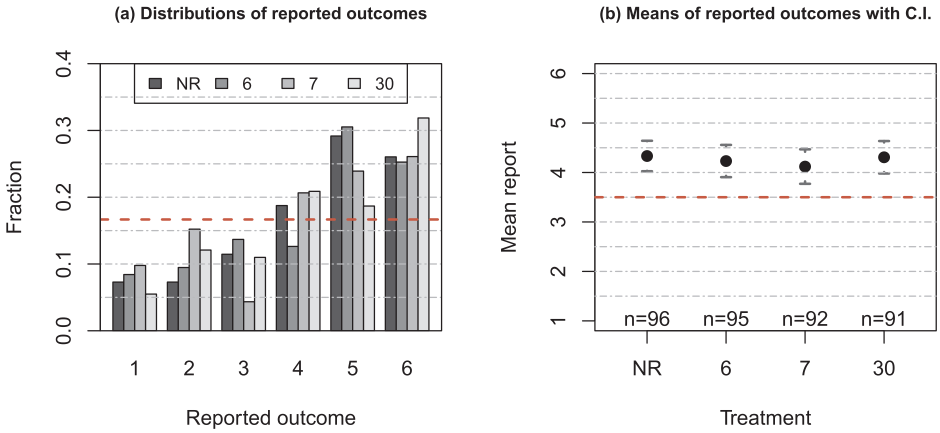

Fig. 1. Distribution of reported outcomes (a) and mean reported outcome with confidence intervals (b).

and in June 2018. Our target was to collect 100 subjects per treatment. We chose this target based on prior commonly used sample sizes in similar studies rather than on a priori power analysis since we had no clear indication of the magnitude of the effect size. Average earnings were approximately €9.2 (inclusive of a €5 show-up fee, SD = 1.6).

4. Results

Panel (a) in Fig. 1 shows the distribution of reported die-rolls across treatments along with the theoretical benchmark under the hypothesis of honest reporting (thick dashed line). Consistently with previous studies using the die-under-the-cup paradigm (Fischbacher & Föllmi-Heusi, 2013; Shalvi et al., 2011), distributions of reported outcomes are skewed to the left with a mode of either 5 or 6. This strongly suggests for the presence of lying in all treatments. Indeed, a Kolmogorov–Smirnov test for discrete distributions strongly rejects the hypothesis that reported outcomes follow a uniform distribution for all treatments (T_{NR}: D = 0.240, $p < 0.001$; T_6: D = 0.225, $p < 0.001$; T_7: D = 0.207, $p < 0.001$; T_{30}: D = 0.214, $p < 0.001$).

Comparing lying across treatments, panel (a) in Fig. 1 suggests that the distribution of reported die-rolls does not differ across treatments (Kruskal–Wallis test $\chi^2(3) = 0.679$, $p = 0.878$). The mean reported outcomes (panel (b) of Fig. 1) confirm this impression: we observe an average report of 4.333 in T_{NR}; of 4.232 in T_6; of 4.120 in T_7; and of 4.308 in T_{30} and no significant differences are found across treatments (Wilcoxon rank sum test, $p > 0.483$ for all comparisons).

As an additional control, we run OLS regressions testing whether treatments' differences emerge when controlling for beliefs about other participants' reported outcome, age, gender, and risk-taking measured by DOSPERT. Regression results are reported in Table 2.[4] As it is apparent from the table, we cannot reject the hypothesis that the average reported outcome is the same across treatments even when we include controls for gender, age, beliefs about other participants' reported number, and the risk-taking domains.

Looking at the control variables, we do not observe any effect of gender and age. However, we observe a strong positive correlation between the beliefs about other participants' reported number and the outcome reported by the respondent. Unfortunately, we cannot infer anything about the causal direction. It may be that who believes that others report higher numbers feels justified to report higher numbers as well, but it may also be the case that who reports a higher number wants to believe that others do the same as a form of moral cleansing. Another interesting result is that reported outcomes seem to be correlated with the ethical dimension of DOSPERT but not with the other dimensions. This confirms previous results in the literature (Zimerman, Shalvi, & Bereby-Meyer, 2014) and shows that DOSPERT scales collected one week before still have predictive power for lying in the die-under-the-cup task.

In the previous section, we argued that, among the risk treatments, lying is highest in T_{30} and lowest in T_6 because misreporting is easier to justify when its impact on the probability of winning is lower due to the weaker feeling of responsibility for obtaining the final outcome. The observed results, however, are not aligned with this prediction. It seems, on the contrary, that the impact on the final likelihood to obtain the benefit has no effect on the level of lying. Additionally, we found that the level of lying when consequences are uncertain is the same as the level of lying when consequences produce a certain outcome.

A Bayesian analysis of the contingency table of reported outcomes by treatment revealed very strong evidence for the null hypothesis of no treatments difference with a $BF_{01} = 227165$ providing "decisive evidence" (Lee & Wagenmakers, 2014) in support of H_0 (independence of the reported outcomes from the treatment) over H_1 (suggesting dependence of reported outcomes on the treatment).[5]

[4] Since the completion of the DOSPERT questionnaire was not mandatory, we have a smaller sample for the regression using DOSPERT scales as control variables.

[5] The reported Bayes factor tests the independence assumption of the contingency table of reports under the independent multinomial sampling plan where number of participants in the different treatments are the fixed margin.

Table 2
OLS Regressions (Dep var: *Reported outcome*−3.5)

	Mod. 1		Mod. 2	
	Estimate	S. E.	Estimate	S. E.
(Intercept)	0.7316	(0.1636) *	0.5595	(0.7859)
T_{NR}	0.1018	(0.2308)	−0.1304	(0.3363)
T_7	−0.1120	(0.2333)	−0.2475	(0.2737)
T_{30}	0.0761	(0.2340)	−0.1435	(0.2724)
Belief−3.5	–	–	0.3190	(0.0975)
Risk Ethical	–	–	0.3824	(0.1914)*
Risk Financial	–	–	0.1070	(0.1546)
Risk Health	–	–	−0.1158	(0.1927)
Risk Recreational	–	–	−0.2055	(0.1609)
Risk Social	–	–	−0.0257	(0.2141)
Male	–	–	−0.0067	(0.2303)
Age−$\overline{Age}$	–	–	−0.0013	(0.0084)
N	374		233	
R^2	0.0027		0.0935	

Signif. codes: 0 < ⩽ 0.001 < ** ⩽ 0.01 < * ⩽ 0.05 < ° ⩽ 0.10.

Finally, we predicted a drop in the reported number of 6 when moving from T_7 to T_6 because we argued that reporting a 6 in T_6 triggers a stronger feeling of responsibility compared to reporting a 6 in T_7. Looking at the fraction of 6 reported, we cannot reject the hypothesis of no differences across treatments for this measure as well ($\chi^2(3) = 1.301$, p = 0.729).

5. Discussion

Our results reject the hypothesis that lying is affected by its impact on the likelihood to get a benefit. We observed the same lying levels across all treatments, i.e, when consequences of lying are deterministic (T_{NR}), when consequences are uncertain but certainty can be achieved by lying to the full extent (T_6), and when consequences are uncertain and certainty cannot be achieved (T_7 and T_{30}). This seems to suggest that the moral cost of lying is robust to different levels of implied responsibility for the obtainment of a benefit. Lying to obtain a benefit with certainty is not different from lying to increase the likelihood of getting a benefit.

A potential issue of our design is the role of risk aversion. Our hypotheses rely on the assumption that risk aversion does not play a major role in the choices of our participants. Here we tackle the question of what is the impact of risk aversion and show that its effect goes in the opposite direction compared to our responsibility hypothesis.

Consider first the case of risk averse agents without moral costs of lying. In this case, the attitude towards risk does not matter for predictions. As far as utility is increasing in money and expected utility holds, we should observe a report of 6 in all the treatments. This because reporting higher numbers results in lotteries that stochastically dominate the ones obtained by reporting lower numbers.

Predictions are different, however, when considering the interaction between risk aversion and lying costs. To give an idea of how predictions change, we consider two simple cases: risk aversion with a constant cost of lying and risk aversion with a cost of lying that is increasing in the size of the lie. We show that, in both cases, the effect of risk aversion goes in the opposite direction compared to our hypotheses.

Let first assume that lying produces a fixed dis-utility c independent of the reported outcome. This implies that, when lying is profitable for some report r, lying to the full extent must be strictly preferred. Therefore, in treatment T_{NR} an agent observing x will lie if the gain in utility reporting $r = 6$ is bigger than the dis-utility from lying, i.e., $U(6)–U(x) > c$. In case of treatment T_k the same agent will lie if the gain in expected utility reporting $r = 6$ is bigger than the dis-utility from lying, i.e., $U(k)\frac{6}{k}–U(k)\frac{x}{k} > c$.[6]

Comparing the condition for T_{NR}, i.e., $U(6)–U(x) > c$, with the condition for T_6, i.e., $U(6)–\frac{x}{6}U(6) > c$, we conclude that a subject observing x and lying in T_{NR} must be lying in T_6 as well. This because $U(x) > \frac{x}{6}U(6)$ due to the concavity of $U()$. Therefore, with risk aversion we should expect more lying in T_6 than in T_{NR}. Comparing then two risk treatments T_{k_1} and T_{k_2} where $k_2 > k_1$, we have that an agent lying in T_{k_2} will also lie in T_{k_1} because $U(k_2)\frac{6-x}{k_2} < U(k_1)\frac{6-x}{k_1}$ due to the concavity of $U()$. This lead to the prediction that lying is the highest in T_6 and the lowest in T_{30}.

Now assume that the cost of lying is an increasing function $f()$ of the distance between the reported and the observed outcome, i.e., the distance between x and r.[7] In treatment T_{NR} the agent reports the number r that maximizes $U(r)–f(r, x)$ which gives the following F.O.C.

[6] Without any loss of generality we assume that $U(0) = 0$.

[7] We assume that $f(r, x)$ is a well-behaved convex function with a minimum in $r = x$. Formally, $f_r'(r, x) < 0$ for $r < x$ and $f_r'(r, x) > 0$ for $r > x$;$f_r''(r, x) > 0$; and $f(x, x) = 0$.

$$U'(r) = f'_r(r, x)$$

This condition implies that the optimal report is $r^* \in [x, 6]$ and that reporting r^* equates the marginal cost of lying by r^*-x and the marginal benefit of obtaining r^*. [8]

In T_k the agent chooses to report the number r that maximizes $U(k)\frac{r}{k} - f(r, k)$ which gives the following F.O.C.

$$\frac{U(k)}{k} = f'_r\left(r, x\right)$$

In this case the optimal report is the one that equates the marginal cost of lying by r^*-x with the average utility per euro $\frac{U(k)}{k}$. As before $r^* \in [x, 6]$.

The two conditions for T_{NR} and T_k make apparent that in this case it is not possible to draw clear conclusions on whether lying is higher in the baseline or in the risk treatments. However, it is still possible to derive predictions for lying levels across the different risk treatments T_k. Because of risk aversion, $\frac{U(k)}{k}$ is decreasing in k. Therefore, for a marginally increasing cost of lying we should observe less lying when k is big compared to when k is small. This means that lying is higher in T_{30} compared to T_7 and T_6.

In both the cases we discuss, the effect of risk aversion is to reduce lying when the impact of lying on the probability to get the benefit becomes smaller. This goes in the opposite direction compared to our predictions based on a moral cost of lying changing with the responsibility for the obtainment of the benefit. This means that, if we had found the predicted effect we would have confidently concluded that the moral cost of lying is increasing when moving from T_6 to T_{30}. The null result, however, does not allow for such conclusion. It may be that the cost of lying is constant and participants are risk neutral or that the moral cost of lying is decreasing with the impact on probability and risk aversion offsets this effect.

A different natural choice of parameters for our task would have been to keep constant the size of the prize across different risk treatments, e.g., keep a fixed 6 euro prize. This would have induced different expected gains from lying across treatments but it would have ruled out the effect of risk aversion when comparing the risk treatments. While acknowledging the potential concern for the effect of different prize sizes in our design, we would like to highlight the fact that different monetary incentives seem to have little impact on the likelihood of lying (Abeler, Nosenzo, & Raymond, 2016). This would suggest similar results when keeping the prize sizes constants across risk treatments and when keeping the same expected gain from lying constant across treatments.

A recent study shows that lying is less pronounced when the ex-ante likelihood to get a benefit is small (Gneezy, Kajackaite, & Sobel, 2018). The authors argue that this is because the smaller the likelihood to honestly obtain the benefit, the more suspicious and costly is reporting to have obtained it. This cost, which the authors relate to social identity concerns, has a larger impact on lying behavior compared to the costs measuring the size of the lie, i.e., the distance between the reported and actual outcome and the monetary gain obtained by lying. In our setting we keep constant the ex-ante probability to obtain each of the die numbers and, therefore, we argue that social identity concerns remain constant across conditions. In other words, we argue that, since the experimenter knows the outcome space, the loss in reputation when reporting a 6 remains the same in all treatments. The fact that social identity concerns are the main component of lying costs and that they override other components may explain why, in our design where reputation concerns are constant across treatments, we do not find differences in the level of lying.

A final point worth discussing relates to the implications of our result for the studies showing that cheating is higher in tournaments compared to other incentive schemes such as piece-rate (Faravelli et al., 2015; Schwieren & Weichselbaumer, 2010). The fact that we do not find any difference in lying when moving from a setting where final consequences are deterministic—which is similar to a piece-rate incentive scheme—to a setting where lying only increases the probability to obtain the benefit—which is similar to a winner-takes-all tournament—supports the idea that the main driver of the higher level of cheating in tournaments is competitive pressure rather than the lower feeling of responsibility for winning the tournament by lying.

6. Conclusions

The average individual tells two lies a day (De Paulo et al., 1996). Some lies convey direct benefits such as overstating the value of claims to insurance companies or understating our tax revenue to the tax authority. Other lies do not generate direct consequences but only increase the probability of obtaining positive outcomes. Faking a resume and exaggerated personal information on dating websites are illustrations of such lies. In this paper we explored how people lie in the latter situation where the final consequences of lying are uncertain.

We argued that, when consequences are uncertain, the moral cost of lying changes with the feeling of responsibility for the final outcome, Specifically, we hypothesized that, when lying permits to increase the likelihood of, but not to secure, the obtainment of a benefit, people are more prone to lying because it leaves room for chance to determine the final outcome and reduces the feeling of responsibility. We argued as well that the smaller the impact of lying on the probability to obtain the benefit the lower is the feeling of responsibility to obtain the outcome.

We tested our predictions using a modified die-under-the-cup task where misreporting, rather than delivering a higher payoff, increased the likelihood to get a prize. Overall we had four treatments: a baseline replication of the task conducted by Shalvi et al. (2011) and three treatments where the reported outcome affects the probability to get a prize to a different extent. Contrary to our prediction, we did not observe any treatment difference suggesting that lying is independent of the extent to which it increases the

[8] Here we use the assumption that U' is decreasing in its argument and that, for $r > x$ the marginal costs of lying f' is increasing in its argument.

probability to get a benefit. This result suggests that the willingness to lie to secure a benefit or to marginally increase the probability to obtain the benefit are very similar. Job applicants, for instance, would lie to the same extent when knowing that this will get them hired and when knowing that this merely increases the chance to get hired. A consequence of our result is for those organizations seeking to boost honesty among employees applying for promotions and for bonuses. Investing resources to try to reduce the effectiveness of cheating on the probability to get promoted or to get a bonus will not discourage cheaters and will not reduce the number of attempts.

Acknowledgements

Financial support from the Burgundy Regional Council (PARI SHS 6) and from the European Research Council (grant agreement ERC-StG-637915) is gratefully acknowledged. The authors thank Shahar Ayal and Verena Utikal for insightful comments on earlier version of this article.

Appendix A. Instructions

A.1. Instructions treatment T_{NR}

This experiment is about decision-making. Please read the instructions entirely and carefully. The instructions will help you to understand correctly the experiment.

In this experiment, your payoff in this experiment will depend upon your decisions. All your decisions will be anonymous. You will indicate your decisions on a decision sheet that will be given by a monitor during the experiment. The total amount of money earned during the experiment will be paid to you in cash at the end of the experiment. There is no good nor bad answer.

From now and until the end of the experiment, we ask you to remain silent. If you have any questions, raise your hands and a monitor will come to answer your questions privately.

General framework of the experiment

In this experiment, you have to throw a six face die that determines your payoff.

More precisely, you have an opaque cup with a cover. The small hole located in the cover allows you to see the die. You must shake the cup so as to throw the die. Then without removing the cover, take a look through the hole to observe the outcome of your throw. The number displayed by the die corresponds to your payoff in euro for the experiment (the below table indicates the payoffs associated to each possible outcome of the die).

Number displayed by the die	Your payoff for the experiment
1	1€
2	2€
3	3€
4	4€
5	5€
6	6€

The first roll will determine your payoff for the experiment. After the first roll, we ask that you roll the die under the cup at least 2 more times so that you can verify for yourself that the die is legitimate.

In a few moments, a monitor will give you a "decision sheet" as well as the cup so you can roll the die. After rolling the die at least three times, write down on the "decision sheet" the number displayed by the first roll. Give the cup back to the monitor and answer to the other questions of the "decision sheet".

A.2. Instructions treatment T_6

This experiment is about decision-making. Please read the instructions entirely and carefully. The instructions will help you to understand correctly the experiment.

In this experiment, your payoff in this experiment will depend upon your decisions. All your decisions will be anonymous. You will indicate your decisions on a decision sheet that will be given by an experimenter during the experiment. The total amount of money earned during the experiment will be paid to you in cash at the end of the experiment. There is no good nor bad answer.

From now and until the end of the experiment, we ask you to remain silent. If you have any questions, raise your hands and an experimenter will come to answer your questions privately.

General framework of the experiment

In this experiment, you can win a prize of €6. 6 coloured balls, either red or blue, are placed into a bowl. You have to randomly draw one ball. If the ball you draw is blue you win €6 if not you win nothing. At the beginning of the experiment, there are 6 red balls in the bowl. The number of blue balls depends on your roll.

Before drawing randomly a ball, you will have to roll six face die. More precisely, you have an opaque cup with a cover. The small hole located in the cover allows you to see the die. You must shake the cup so as to throw the die. Then without removing the cover, take a look through the hole to observe the outcome of your throw. The number displayed by the die will determine the number of blue balls in the bowl (the below table indicates the number of blue and red balls according to each possible outcome of the die).

Number displayed by the die	Number of blue balls	Number of red balls	Prize
1	1	5	6€
2	2	4	6€
3	3	3	6€
4	4	2	6€
5	5	1	6€
6	6	0	6€

The first roll will determine the number of blue balls located in the bowl. After the first roll, we ask that you roll the die under the cup at least 2 more times so that you can verify for yourself that the die is legitimate.

In a few moments, the experimenter will give you a "decision sheet" as well as the cup so you can roll the die. After rolling the die at least three times, write down on the "decision sheet" the number displayed by the first roll. Give the cup back to the experimenter, the experimenter will set the number of blue balls according to your die report, and you will then draw randomly a ball from the bowl, write down the colour of the ball you draw and answer to the other questions of the "decision sheet".

Appendix B. Supplementary material

Supplementary data associated with this article can be found, in the online version, at https://doi.org/10.1016/j.joep.2018.09.003.

References

Abeler, J., Becker, A., & Falk, A. (2014). Representative evidence on lying costs. *Journal of Public Economics, 113*, 96–104.
Abeler, J., Nosenzo, D., & Raymond, C. (2016). Preferences for truth-telling. Discussion Papers 10188 IZA Institute of Labor Economics Bonn.
Ariely, D. (2012). The Honest Truth About Dishonesty: How We Lie to Everyone—Especially Ourselves. Harper Collins.
Bartling, B., & Fischbacher, U. (2012). Shifting the blame: On delegation and responsibility. *The Review of Economic Studies, 79*, 67–87.
Becker, G. S. (1968). Crime and punishment: an economic approach. *Journal of Political Economy, 76*, 169–217.
Blais, A.-R., & Weber, E. U. (2006). A domain-specific risk-taking (dospert) scale for adult populations. *Judgment and Decision Making, 1*, 33–47.
Brock, J. M., Lange, A., & Ozbay, E. Y. (2013). Dictating the risk: Experimental evidence on giving in risky environments. *American Economic Review, 103*, 415–437.
Cartwright, E., & Menezes, M. L. (2014). Cheating to win: Dishonesty and the intensity of competition. *Economics Letters, 122*, 55–58.
Conrads, J., Irlenbusch, B., Rilke, R. M., Schielke, A., & Walkowitz, G. (2014). Honesty in tournaments. *Economics Letters, 123*, 90–93.
Conrads, J., Irlenbusch, B., Rilke, R. M., & Walkowitz, G. (2013). Lying and team incentives. *Journal of Economic Psychology, 34*, 1–7.
De Paulo, B. M. (2004). The many faces of lies. In A. G. Miller (Ed.). *The Social Psychology of Good and Evil* (pp. 303–326). New York: Guilford Press.
De Paulo, B. M., Kashy, D. A., Kirkendol, S. E., Wyer, M. M., & Epstein, J. (1996). Lying in everyday life. *Journal of Personality and Social Psychology, 70*, 979–995.
Dugar, S., Mitra, A., & Shahriar, Q. (2018). Deception: The role of uncertain consequences. Mimeo.
Erat, S. (2013). Avoiding lying: The case of delegated deception. *Journal of Economic Behavior & Organization, 93*, 273–278.
Faravelli, M., Friesen, L., & Gangadharan, L. (2015). Selection, tournaments, and dishonesty. *Journal of Economic Behavior & Organization, 110*, 160–175.
Fischbacher, U., & Föllmi-Heusi, F. (2013). Lies in disguise – An experimental study on cheating. *Journal of the European Economic Association, 11*, 525–547.
Gächter, S., & Schulz, J. F. (2016). Intrinsic honesty and the prevalence of rule violations across societies. *Nature, 531*, 496–499.
Gneezy, U., Kajackaite, A., & Sobel, J. (2018). Lying aversion and the size of the lie. *American Economic Review, 108*, 419–453.
Gordon-Hecker, T., Rosensaft-Eshel, D., Pittarello, A., Shalvi, S., & Bereby-Meyer, Y. (2017). Not taking responsibility: Equity trumps efficiency in allocation decisions. *Journal of Experimental Psychology: General, 146*, 771–775.
Grolleau, G., Kocher, M.G., & Sutan, A. (2016). Cheating and loss aversion: Do people cheat more to avoid a loss? Management Science, 62, 3428–3438.
Hamman, J. R., Loewenstein, G., & Weber, R. A. (2010). Self-interest through delegation: An additional rationale for the principal-agent relationship. *American Economic Review, 100*, 1826–1846.
Kandul, S., & Kirchkamp, O. (2018). Do i care if the others lie? current and future effects of delegation of lying. Journal of Behavioral and Experimental Economics, Forthcoming.
Krawczyk, M., & Le Lec, F. (2010). 'give me a chance!' an experiment in social decision under risk. *Experimental Economics, 13*, 500–511.
Lee, M.D., & Wagenmakers, E.-J. (2014). Bayesian Cognitive Modeling: A Practical Course. Cambridge University Press.
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.
Motro, D., Ordóñez, L.D., Pittarello, A., & Welsh, D.T. (2016). Investigating the effects of anger and guilt on unethical behavior: A dual-process approach. Journal of Business Ethics, (pp. 1–16).
Pitesa, M., Thau, S., & Pillutla, M. M. (2013). Cognitive control and socially desirable behavior: The role of interpersonal impact. *Organizational Behavior and Human Decision Processes, 122*, 232–243.
Schwieren, C., & Weichselbaumer, D. (2010). Does competition enhance performance or cheating? a laboratory experiment. *Journal of Economic Psychology, 31*, 241–253.
Shalvi, S. (2012). Dishonestly increasing the likelihood of winning. *Judgment and Decision Making, 7*, 292–303.
Shalvi, S., Dana, J., Handgraaf, M. J., & De Dreu, C. K. (2011). Justified ethicality: Observing desired counterfactuals modifies ethical perceptions and behavior. *Organizational Behavior and Human Decision Processes, 115*, 181–190.
Shalvi, S., Eldar, O., & Bereby-Meyer, Y. (2012). Honesty requires time (and lack of justifications). *Psychological Science, 23*, 1264–1270.
Weisel, O., & Shalvi, S. (2015). The collaborative roots of corruption. *Proceedings of the National Academy of Sciences, 112*, 10651–10656.
Zimerman, L., Shalvi, S., & Bereby-Meyer, Y. (2014). Self-reported ethical risk taking tendencies predict actual dishonesty. *Judgment and Decision Making, 9*, 58–64.