Who never tells a lie?*

Christoph Vanberg†

Abstract

I experimentally investigate the hypothesis that many people avoid lying even in a situation where doing so would result in a Pareto improvement. Replicating Erat and Gneezy (2012), I find that a significant fraction of subjects tell the truth in a sender-receiver game where both subjects earn a higher payoff when the partner makes an incorrect guess regarding the roll of a die. However, a non-incentivized questionnaire indicates that the vast majority of these subjects expected their partner not to follow their message. I conduct two new experiments explicitly designed to test for a 'pure' aversion to lying, and find no evidence for the existence of such a motivation. I discuss the implications of the findings for moral behavior and rule following more generally.

Keywords: Lying; Deception; Morality; Ethics; Experiments

1 Introduction

One of the most basic moral rules in most cultures is that one should not lie. The importance of this moral principle is reflected in stories told to children, such as the one about the boy who cried wolf. Some moral philosophers follow Immanuel Kant (1785) in suggesting that rules of this kind should be followed in all instances, irrespective of the consequences to be expected in any particular case. According to this position, it would be wrong to lie even in a case where doing so would lead to consequences that would be generally regarded as desirable. Given the importance of this and other basic moral principles, it is interesting to ask whether (or to what extent) common moral attitudes and behaviors coincide with this Kantian position. That is, to what extent do people regard lying as 'wrong' and therefore avoid lying, irrespective of the consequences to be expected from doing so?

A number of authors have experimentally investigated this question using variants of a sender-receiver game originally introduced by Gneezy (2005) (e.g.

*I thank the editor and two anonymous referees for constructive comments. I received valuable encouragement from seminar and workshop participants in Rome, East Anglia, Karlsruhe, and Vienna. Thanks also to Sanjiv Erat and Uri Gneezy for answering my questions in preparation for the replication.

†Department of Economics, University of Heidelberg, Bergheimer Str. 58, 69115 Heidelberg, Germany; Tel: +49 6221 542912; E-mail: vanberg@uni-hd.de

Dreber and Johannesson 2008, Cohen et al 2009, Sutter 2009, Erat and Gneezy 2012, Gneezy et al 2013). In general, the message emerging from this literature has been that most experimental subjects take expected consequences into account when deciding whether to lie. In particular, it appears that people are more likely to lie if the personal benefit to be gained from doing so is larger, and if the harm imposed on others is smaller.[1] With one exception, these experiments have investigated lying in settings where the person telling a lie can benefit at a cost to someone else.[2]

In a recent contribution, Erat and Gneezy (2012, henceforth EG) aimed to investigate cases where lying can be materially beneficial to another person ('white lies'), and possibly to both the liar and another person ('Pareto white lies'). Testing for an aversion to the latter type of lie is particularly interesting. Since the material consequences of telling such a lie are unambiguously desirable, such an aversion could be attributed only to a "pure" cost of lying. Therefore this could provide a good test for the existence of 'Kantian' moral attitudes. As the authors explain,

> "[P]eople who are reluctant to tell Pareto white lies demonstrate lie aversion independent of social preferences for outcomes. Such people refrain from lying not (merely) because of the consequences, but because they simply view lying as a bad act in itself. This provides the best test of a pure cost of lying in line with a moral stand." (ibid. p. 724)

To test for such a 'pure' aversion to lying, EG conduct an experimental game involving 2 players, labeled 'sender' and 'receiver'. The sender is informed that a die has been rolled and the outcome was '2'. He is also told that the receiver does not know the outcome of the die roll. The sender is asked to send a message of the form 'The number rolled was X'. This message is transmitted to the receiver, who must then choose a number between 1 and 6. (I will refer to this number as the receiver's 'guess,' though EG do not use this terminology.) The sender is told that the players' payoffs are determined as follows. If the receiver chooses 2 ('guesses correctly'), both receive 20 dollars. Otherwise (i.e. if she 'guesses wrong'), both receive 30 dollars. Finally, the sender is told that the receiver does not know the payoff consequences of guessing correctly or not.

The result of the experiment is that 35% of senders tell the truth in this situation (choose the message "The number rolled was 2"). The authors interpret this result as showing that "A significant fraction of senders do not lie *even when lying results in a Pareto improvement.*" (ibid. p. 726, emphasis added). This interpretation appears to be widely accepted, and many authors cite EG's

[1] See Hurkens and Kartik (2008) for a critical discussion of this interpretation.

[2] Rosaz and Villeval (2012) investigate lying in a context where a 'supervisor' reports on another subject's performance in a real effort task. Under some conditions, the supervisors have the opportunity to misreport in a way that increases both subjects' payoffs. They find that a significant fraction of subjects truthfully report in such situations (ibid, Table 3). I will come back to this evidence in the conclusion.

experiment as evidence for a 'pure' aversion to lying (e.g. Cappelen et al 2013, Chen and Houser 2013, Rosaz and Villeval 2012).

Upon reflection, it becomes apparent that this interpretation relies on the implicit assumption that those senders who told the truth in the experiment did so believing that receivers will 'follow' their messages.[3] In order to investigate this assumption, I replicate EG's experiment with the addition of a (non-incentivized) post-experimental questionnaire. I find that 32 out of 106 subjects (30%) told the truth, closely replicating EG's original result. The non-incentivized questionnaire answers suggest, however, that the vast majority of these subjects believed that the receiver would *not* follow their messages, and expected that telling the truth, not lying, would yield better material consequences for both subjects.[4]

Since the questionnaire evidence was not incentivized, the replication study delivers only suggestive evidence that the received interpretation of EG's finding may be mistaken. I therefore conduct two new experiments explicitly designed to test the hypothesis that a significant part of the population avoids lying, even when doing so would result in a Pareto improvement. These experiments differ from EG's in several respects. However, subjects in both experiments ultimately face essentially the same choice as in EG. Both experiments are designed such that senders are likely to be confident that receivers will follow their messages. This is confirmed using a (non-incentivized) questionnaire. As we will see, the results obtained differ substantially from those of EG, and cast doubt on the existence of a 'pure' aversion to lying. More generally, they suggest that subjects' moral attitudes rarely, if ever, coincide with the Kantian position which demands that ethical rules be followed in all instances and irrespective of consequences.

The remainder of the paper is organized as follows. Section 2 presents the replication study, including the results of the non-incentivized questionnaire. Section 3 discusses one of the new experiments and results. Section 4 concludes. Experimental instructions are provided in Appendices 1-3. The second new experiment is presented in Appendix 4.

2 Replication

The replication experiment was designed to replicate, as closely as possible, EG's original experiment. Sender and receiver decisions were elicited in sepa-

[3] Specifically, truth telling is costly (in expectation) only if the probability with which the receiver follows is greater than 1 in 6.

[4] This approach and the findings are reminiscent of Sutter (2009), who points out the importance of sender expectations for the interpretation of truth telling in Gneezy (2005). In that experiment, the receiver could choose between only two options, and interests were opposed. Sutter (2009) replicated the experiment with the addition of a (non-incentivized) questionnaire asking subjects which option they expect their partner to choose. Based on the responses, he concludes that a significant fraction of subjects engage in 'sophisticated' deception: telling the truth in the expectation that the receiver would not follow and thus choose the 'wrong' option.

rate (sequential) sessions conducted in a classroom setting, at the University of Heidelberg, Germany. 106 senders participated (compared to 102 in the original experiment). The instructions used were direct translations, into German, of EG's original (see Appendix 1). The only modification was that payoffs were expressed in Euros rather than Dollars. I.e. both subjects received 20 EUR if the receiver chose the number that was rolled, and 30 EUR if they chose a different number.[5] As in the original experiment, subjects were told that 1 out of every 20 participants would be paid.

Subjects entered their decisions manually onto the instruction sheets and subsequently placed the sheets into envelopes. After these were collected, a questionnaire containing the following three questions was distributed:

(1) What do you expect? Will the other participant choose the number you have transmitted to him, or a different number?

- o The other participant will choose the number I have transmitted to him.
- o The other participant will choose a different number.

(2) What result do you expect in case you are chosen for payment?

- o Both participants receive 20 EUR
- o Both participants receive 30 EUR
- o A different result (please specify):

(3) For what reason did you choose the message you sent?

- o I wanted to receive the maximum payoff.
- o In my opinion, one should always tell the truth.
- o I chose entirely at random.

As emphasized above, the questionnaire was not incentivized.[6] After all participants completed the questionnaire, 6 decision sheets were randomly drawn to determine which subjects would be paid. The corresponding messages were manually entered onto the instruction sheets for 6 receiver subjects. In a second (separate) session, 6 subjects played the role of receivers.[7] Their instructions were also direct translations from the original. Naturally, the behavior of receivers is irrelevant for the purpose at hand, since what matters is only senders' *expectations* concerning that behavior.

[5] On the day of the experiment (11/9/15), the exchange rate was 1 EUR = 1.07 USD.

[6] As noted by the editor, one might feel especially suspicious of non-incentivized evidence in a study about lying. In particular, we may be especially worried that answers *consistent* with pure lying aversion may be false. That is, a subject may send a true message for instrumental reasons and subsequently lie in a way that makes him appear to be lying averse. In contrast, answers *inconsistent* with pure lying aversion are less worrisome, since a lying averse subject would (presumably) not lie on the questionnaire. Thus, the number of answers *consistent* with pure lying aversion probably constitutes an upper bound on the true number of lying averse participants.

[7] This procedure mirrors the one employed by EG (personal communication).

2.1 Results

Table 1 summarizes the main results from the replication experiment. The first number in each cell represents the number of subjects who told the truth / lied and expected the recipient to follow / not to follow. The pair of numbers in brackets represent the number of subjects within each cell who expected a payoff of 20 or 30 EUR. For example, 42 subjects lied and expected the receiver to follow. Of these, 3 expected a payoff of 20 EUR, and 39 expected a payoff of 30 EUR.

Table 1. Truth telling and declared expectations†

	expect follow	expect not	$\sum$
truth	8 (8/0)	24 (0/24)	32
lie	42 (3/39)	32 (2/30)	74
$\sum$	50	56	106

† The first number in each cell represents the number of subjects who told the truth / lied and expected the recipient to follow / not to follow. The pair of numbers in brackets represent the number of subjects within each cell who expected a payoff of 20 or 30 EUR.

Overall, 32 of 106 subjects told the truth. At 30%, this replicates EG's finding of 35% truth telling.[8] Among these 32 subjects, 24 declared that they expect the receiver *not to follow* and to receive a payoff of 30 EUR, while 8 stated that they expect the other participant to *follow* and to receive a payoff of 20 EUR. The remaining 74 subjects lied. Of these, 42 expected the receiver to *follow*. Among those who lied, all but 5 subjects said that they expected a payoff of 30 EUR.[9]

To assess what these data reveal about the extent of 'pure' lie aversion, we may start by considering the 24 subjects who tell the truth and expect the receiver *not* to follow. Of these subjects, 19 declared that they chose their message in order to maximize their payoff. This indicates that these subjects felt *confident* that the receiver will not follow.[10] Since (subjectively) these subjects are not in a position to tell a Pareto white lie, the corresponding observations must be discarded. The remaining 5 subjects in this category declared that they felt one should always tell the truth. Their answers are therefore consistent with the possibility that they felt *uncertain* about the receiver's behavior, such that truth telling was costly.[11] In order to bias our analysis in favor of 'pure' lie aversion, we may count these 5 observations as revealing such a motivation.

[8] A Chi-square test for differences yields $p = .43$.

[9] The 5 subjects who lie and expect 20 EUR appear to have been confused, as that outcome would occur only if the receiver does not follow and coincidentally chooses '2', which occurs with *at most* probability 1/5.

[10] Recall that truth telling is (expected) payoff maximizing if the receiver follows with a probability below 1/6.

[11] That is, these subjects *may* have believed that the receiver will follow with a probability greater than 1/6.

For the remaining 82 subjects, we have more definite indications. Among the 8 who tell the truth and expect the receiver to follow, 5 declared that they felt one should always tell the truth, while 3 said that they chose 'entirely at random'. Again, to bias the analysis in favor of 'pure' lie aversion, we may count all 8 of these observations as indicative of that motivation. Clearly, the 74 subjects who lied are not averse to telling lies, be they white or black.[12]

To summarize, of the 82 subjects for whom a tentative inference is possible, only 8 (10%) exhibit behavior consistent with pure lie aversion. If we include the 5 *uncertain* subjects mentioned earlier, we would arrive at 13 out of 87, or 15%. A similar number emerges if we focus only on the left column of Table 1, where 16% tell the truth. These estimates are intentionally biased in favor of a pure aversion to lying. A more conservative estimate can be constructed by excluding the 24 truth-tellers who expect the receiver not to follow, as well as the 3 who stated that they chose at random. In this case we would be left with 5 out of 79, or 6%. Thus, the data are consistent with between 6% to at most 16% of subjects exhibiting 'pure' lie aversion, well short of the 35% originally reported by EG.

It should be emphasized that the questionnaire is not incentivized, and therefore the data collected provide only suggestive evidence that many senders may not have expected receivers to follow. Indeed, it is interesting to note that a *majority* of participants stated the belief that the receiver would not follow. Although we may draw some tentative inferences, it appears difficult to reliably interpret the behavior observed. Thus, it would seem desirable to conduct a new test of the hypothesis that many subjects indeed exhibit a 'pure' aversion to lying. In order to conduct such a test, we may look at a situation in which senders reliably expect, with a sufficient degree of confidence, that receivers will follow their message. The experiments described in the following section were designed to achieve exactly that, while ultimately presenting senders with essentially the same choice that they make in EG.

3 New Experiments

3.1 Design

I conducted two experiments explicitly designed to test the hypothesis that (many) people avoid lying even when doing so would result in a Pareto improvement. Both experiments involve sender-receiver games similar to the one conducted by EG. Senders in both experiments face the same choice between accurate or inaccurate messages concerning the roll of a die. Other than this, the experiments differ substantially from EG. All modifications are intended to ensure that senders may be quite certain that receivers will follow their messages. I will begin by describing one of these experiments in detail. The other

[12] Those who lie are likely to have believed that the receiver will follow with a probability greater than 1/6. The only alternative explanation would be that they exhibit a pure preference *for* lying. In either case, their behavior is inconsistent with a pure preference for truth telling.

is discussed in subsection 3.4. Results from both experiments are essentially identical.

As in EG, subjects played a sender-receiver game. First, the sender rolled a 6-sided die.[13] He then sent a message of the form 'The number rolled was X'. Finally, the receiver chose a number between 1 and 6. If the chosen number corresponded to the number rolled, both participants received a certain common payoff A. If not, both received another common payoff B. As in EG, only the sender was informed about the concrete values A and B. Unlike in their experiment, however, the fact that both subjects would receive the same payoff was common knowledge. All subjects received the same instructions, such that senders knew perfectly what information was available to receivers.

An important feature of this design is that the players commonly know that their material interests are aligned, though still the receiver does not know whether it is better to guess correctly or not. My expectation was that this should make senders more confident that the receiver will indeed follow their message. (Naturally, there is still no purely *rational* argument for doing so.) In order to emphasize the commonality of interest, the instructions referred to the participants as a 'team'.

In order to further increase the predictability of the receiver's behavior, the game was repeated 8 times in fixed pairs. The payoff values A and B varied from round to round, such that sometimes it was better to 'guess correctly' ($A > B$) and sometimes it was better to 'guess wrong' ($B > A$). The roles (sender/receiver) were switched after 4 rounds. I expected this repetition to further reduce uncertainty as to the receiver's behavior. In order to elicit relevant observations not affected by prior coordination, the first round was set up such that $B > A$, i.e. it was better to guess a wrong number.

Finally, senders were asked whether they expect the receiver to follow their message or not. As in the replication study, this elicitation was not incentivized.[14] Following each round of play, subjects received full feedback indicating the number that was rolled, the message that was sent, the number chosen, and the resulting payoffs. Following the experiment, all subjects were paid in cash for one randomly chosen round.

3.2 Procedures

The experiment was conducted at the University of Heidelberg in Germany. Subjects were students from a variety of disciplines. Participants were recruited using the online recruitment system ORSEE (Greiner 2004). The experiment was computerized using z-Tree (Fischbacher 2007). We conducted two sessions involving 16 and 20 subjects. Upon entering the laboratory, subjects were randomly assigned to visually isolated computer terminals. Instructions (reproduced in Appendix 2) were handed out. After subjects had read the instructions,

[13] The roll of the die was simulated by a computer.

[14] In the other experiment, the receiver was allowed to *tell* the sender, prior to the choice of message, whether she intends to follow the message or not. After observing that no sender told the truth in that experiment, this feature was eliminated. See Section 3.4.

they were read aloud by the experimenter. Next, subjects filled out a brief questionnaire testing comprehension of the rules (reproduced in Appendix 3). The experiment began after all subjects had correctly filled out the questionnaire.[15] The payoff values A and B were either 12 and 6 EUR or 6 and 12 EUR. Each subject played 4 rounds as a sender, 2 in which $A > B$ and 2 in which $B > A$.[16] After the experiment was completed, one of the 8 rounds was randomly chosen for payment. Subjects filled out a post-experimental questionnaire asking about their attitude towards lying. Finally, participants were paid in private and left the laboratory.

3.3 Results

Each of the 36 participants made 4 choices as a sender, giving us 144 cases overall. In 72 of these cases, both participants would have received a higher payoff (12 instead of 6 EUR) if the receiver correctly guessed the number rolled. In the other 72 cases, they would have received a higher payoff if the guess was wrong. Naturally, these observations are not independent. Given the fixed matching scheme, we have 18 statistically independent observations, this being the number of pairs observed.

When it was better to guess correctly

First, let us look at the 72 cases where $A > B$. In that situation, both the sender and the receiver receive 12 EUR if the receiver guesses correctly and 6 EUR if not. In 68 of those cases (94%), senders said that they expected the receiver would follow the message. Given this expectation, the (joint) payoff maximizing choice for the sender would be to tell the *truth*, i.e. to transmit the number actually rolled.

The left panel of Figure 1, depicts the distribution of messages given the numbers actually rolled. Note that all points are located along the 45 degree line, suggesting a strong correlation between the messages sent and the actual roll of the die. Another way to visualize this result is shown in the right panel, which contains a histogram of the frequencies of lying and truth telling in this condition. Naturally, these data are consistent with the notion that some or all subjects dislike telling lies.

When it was better to guess wrong

Next, consider the 72 cases where $B > A$. In that situation, both the sender and the receiver receive 12 EUR if the receiver guesses *wrong* and 6 EUR if he guesses correctly. In 67 of those cases (93%), senders said that they expected the receiver would follow their message. (This fraction is not statistically different

[15] Only very few subjects made a mistake on the questionnaire. In such cases they were asked to reconsider their answers. All subjects were then able to correct their errors immediately.

[16] The instructions informed subjects that the payoffs would differ from round to round, but no information was given as to the underlying distribution.

Figure 1: Truth telling when A>B (better to guess correctly)

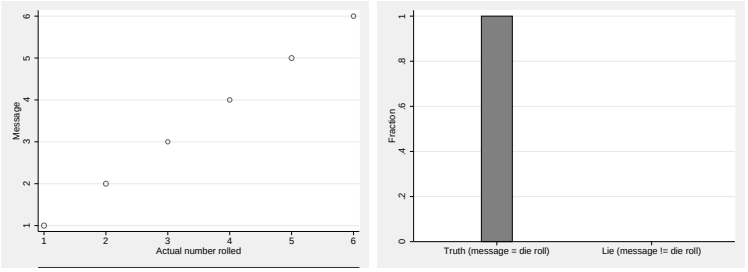

Left panel: The distribution of messages given actual rolls of the die.
Right panel: The relative frequency of truth and lies.

Figure 2: Truth telling when B>A (better to guess wrong)

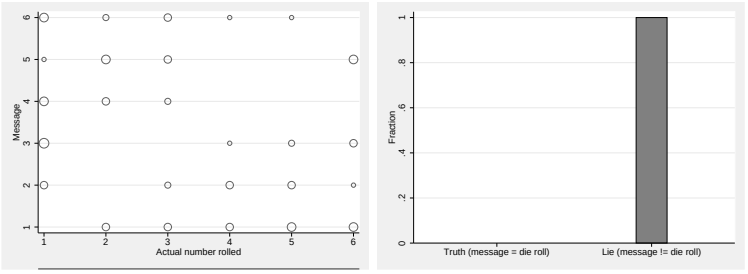

Left panel: The distribution of messages given actual rolls of the die.
Right panel: The relative frequency of truth and lies.

from the expectation in the $A > B$ conditions.) Given this belief, the (joint) payoff maximizing choice for the sender would be to *lie*, i.e. to transmit a number different from the one actually rolled.

The left panel of Figure 2 displays the observed distribution of messages given the number rolled. In contrast to Figure 1, we can see that all observations are located away from the 45 degree line, suggesting that the message sent was reliably different from the actual roll of the die. A histogram displaying the frequency of truth telling and lying is provided in the right panel of Figure 2. Naturally, these data are not at all consistent with the notion that any of the participants dislike telling lies when doing so results in a Pareto improvement.

3.4 Additional experiment

Prior to conducting the experiment described above, I conducted two sessions with a slightly different design. The main difference between that design and the one described above was that the receiver was permitted to send a message to the sender, explicitly stating whether she intends to follow the sender's message. 40 subjects participated in this experiment. As in the main experiment, subjects interacted in fixed pairs and each subject was in the role of sender four times. Each subjects faced two situations in which it would have been better to guess

correctly ($A > B$) and two in which it would be better to guess wrong ($B > A$). Unlike the treatment described above, the first round involved $A > B$, i.e. it was better to guess correctly.

With only four exceptions out of 160 cases (discussed in a moment), results from this treatment are identical to those presented above: All receivers stated that they would follow the sender's message. In all cases where it was better to guess correctly, senders told the truth. In all cases where it was better to guess wrong, senders lied. The four exceptions involved receivers who indicated that they would not follow the message. In two of these, it was better to guess correctly, and senders chose to lie. In two others, it was better to guess wrong, and senders told the truth. Thus, *in all cases, senders chose the message that maximized the material payoff given the receiver's stated intention.*

4 Conclusion

Erat and Gneezy (2012) conduct an experiment to test whether people are reluctant to tell a lie even when doing so results in a Pareto improvement. According to the authors, their experiment "provides the best test of a pure cost of lying in line with a moral stand." Based on their results, EG conclude that indeed a significant number of people exhibit such a 'pure' aversion to lying. This conclusion has received substantial attention and appears to be widely accepted.

I conducted a replication experiment with the addition of a non-incentivized post-experimental questionnaire asking subjects about their expectations and the reasons underlying their choices. The results closely replicate EG's original result, with 30% of senders telling the truth. The non-incentivized questionnaire evidence suggests, however, that the vast majority of these subjects expected that telling the truth, not lying, would lead to a Pareto improvement. Depending on how the results are interpreted, the data obtained are consistent with between 6% and at most 16% of subjects exhibiting pure lie aversion. Although the questionnaire was not incentivized, it provides suggestive evidence that the received interpretation of EG's result may have been mistaken.

I therefore conducted two new experiments *similar* to but deliberately *distinct* from EG. My experiments are designed to explicitly test the hypothesis that (many) people avoid telling even Pareto white lies. The main difference is that the material consequences to be expected from lying are rather unambiguous. Although this reduction in ambiguity is achieved by substantial modifications in the overall design, subjects ultimately make the same choice between telling the truth or lying about the outcome of a die roll.

Out of the 72 participants who took part in the two experiments, *each and every* one chose to lie when doing so implied a Pareto improvement in material terms. This evidence suggests that a truly 'pure' aversion to lying is quite rare indeed, and possibly non-existent. More broadly, it indicates that few if any subjects adhere to a Kantian ethic according to which general moral rules must be followed under all circumstances, even when a violation would bring about

consequences that are unanimously preferred.

Before concluding, I should mention two studies, besides EG, that appear to provide some evidence for the existence of 'pure' lie aversion. Rosaz and Villeval (2012) found evidence that subjects are reluctant to lie to the experimenter about another subjects' performance on a real effort task, even when this results in higher earnings for both subjects. In Fischbacher and Föllmi-Heusi (2013) a single subject reports the outcome of a die roll, and her payoff is increasing in the reported number. They estimate that 39% of subjects tell the truth in this context, despite the fact that no other participant would be harmed if they lied.[17] A possible explanation for the difference between these studies and the result obtained herein is that subjects in those experiments would have to lie *to the experimenter*, explicitly making him a party to the interaction that could be harmed by false information. If this is how subjects see the situation, lying would not in fact result in a Pareto improvement. In the experiments reported on here (as in EG), the subject would lie to the other participant, and thus the experimenter is more likely to be ignored.

Finally, it may be important to reconsider the implications of my result for the question of whether honest behavior is typically motivated by a general disposition or commitment to moral rules. Superficially, the evidence presented here might be taken to favor instead the idea that such behavior is governed entirely by an evaluation of its foreseeable consequences. This would be premature. What the evidence suggests is merely that subjects' attitudes about moral rules are not *completely* insensitive to the foreseeable consequences of rule violations. Under normal circumstances, common moral rules prohibit behaviors that inflict *harm* on others. Though Kantian moral philosophers may argue that such rules should be followed even in exceptional instances where no harm is done, it seems possible that people can be simultaneously *committed* to such rules in general, but ready to make exceptions in instances where it is obvious that the rule does not serve any purpose. In other words, the evidence presented here is inconsistent with a Kantian ethical attitude, but consistent with a 'pragmatic' preference for moral rule following.

References

Cappelen, A., E. Sorensen, B. Tungodden 2013. When do we lie? *Journal of Economic Behavior & Organization* **3**, 258-265.

Cohen, T., B. Gunia, S Kim-Jun, J. Murnighan 2009. Do groups lie more than individuals? Honesty and deception as a function of strategic self-interest. *Journal of Experimental Social Psychology* **(45)**, 1321-1324.

Chen, J. and D. Houser 2013. Promises and Lies: An Experiment on Detecting Deception. *GMU working paper*.

[17] However, Fischbacher and Föllmi-Heusi's main result is that many subjects appear to inflate their number to less than the full extent, indicating that they are concerned about the *appearance* of lying.

Dreber, A. and M. Johannesson 2008. Gender differences in deception. *Economics Letters* **(99)**, 197-199.

Erat, S. and U. Gneezy 2012. White Lies. *Management Science* **58**, 723-733.

Fischbacher, U. 2007. Zurich toolbox for readymade economic experiments. *Experimental Economics* **10**, 171-178.

Fischbacher, U. and F. Föllmi-Heusi 2013. Lies in disguise - an experimental study on cheating. *Journal of the European Economic Association* **11**, 525-547.

Gneezy, U. 2005. Deception: The role of consequences. *American Economic Review* **95**, 384-394.

Gneezy, U., B. Rockenbach, M. Serra-Garcia 2013. Measuring Lying Aversion. *Journal of Economic Behavior and Organization* **93**, 293-300.

Greiner, B. (2015). Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE. *Journal of the Economic Science Association* **1**, 114-125.

Hurkens, S. and N. Kartik 2009. Would I lie to you? On social preferences and lying aversion. *Experimental Economics* **12**, 180-192.

Kant, I. 1785. *Fundamental Principles of the Metaphysics of Morals.*

Rosaz, J. and M.C. Villeval 2012. Lies and biased evaluation: A real effort experiment. *Journal of Economic Behavior and Organization* **84**, 537-549.

Sutter, M. 2009. Deception through telling the truth?! Experimental evidence from individuals and teams. *The Economic Journal* **119**, 47-60.

Appendix 1: Erat and Gneezy (2012) Instructions

The following instructions were provided to senders in EG's experiment. For the replication study, these were directly translated to German. The translation is available from the author upon request.

"Before starting this experiment, we have rolled a 6-sided die, and obtained the outcome 2. The other participant will not be informed that the outcome of the die roll was 2. However, he or she will be told that you have been informed the outcome of the die roll.

We now ask you to send a message to the other participant. The message will correspond to a number from 1 to 6. There are six possible messages:

Message 1: 'The outcome from the roll of the 6-sided die is 1.'
Message 2: 'The outcome from the roll of the 6-sided die is 2.'
Message 3: 'The outcome from the roll of the 6-sided die is 3.'
Message 4: 'The outcome from the roll of the 6-sided die is 4.'
Message 5: 'The outcome from the roll of the 6-sided die is 5.'
Message 6: 'The outcome from the roll of the 6-sided die is 6.'

Your message will be shown to the other participant, and then he or she will be asked to choose a number between 1 and 6.

The choice of the number if the other participant will determine the payments in the experiment.

- If he or she chooses the actual outcome of the roll of the dice, 2, then you will receive $20 and he or she will receive $20.
- If he or she chooses a number different than the actual outcome, you will receive $30 and he or she will receive $30.

Finally, only you will be informed of the particular monetary value connected to each message. The other participant will not be informed of these monetary values. However, he or she will be told that you have been informed of the monetary value connected to each message."

Appendix 2: Instructions for the new experiment

The following is a translation of the printed instructions. Original German instructions are available upon request.

Instructions (Page 1 of 2)

Basic Procedure

Teams

Each participant will be randomly assigned to a **Team**. Each **Team** consists of **2** Participants: one **Roller** und one **Chooser**. At the end of the experiment, the two members of a team will receive the same **common payoff**.

Rounds

The experiment will last for 8 rounds. **The composition of the teams will remain fixed for all 8 rounds**. After 4 rounds, the members of a team will switch roles (Roller and Chooser). At the end of the experiment, **one round** will be randomly chosen for payment. Your payoff will correspond to the common payoff of your team in the round chosen.

Procedure within a round

1. The roller **rolls** (with a die) a number between 1 and 6.
   (The roll of the die will be simulated by the computer.)
   Only the roller will know the number rolled.

2. The chooser **chooses** a number between 1 and 6.
   The common payoff of the team depends upon whether the chosen number corresponds to the number rolled.

The **common payoff** of the team is determined as follows:

If number chosen = number rolled -> **both** participants receive **A EUR**

If number chosen = different number -> **both** participants receive **B EUR**

The concrete numbers **A** and **B** will differ from round to round.

Instructions (Page 2 of 2)

Important Details

What the participants know

Only the **roller** in the team will know the number rolled, and only he will know the concrete payment values **A** and **B**.

The **chooser** in the team will not learn the number rolled and the payment values until the end of the round, i.e. after he has chosen a number.

The opportunity to send messages

After the roller has rolled the die, and **before** the chooser chooses a number, the roller can send one of the following 6 messages:

- "The number rolled is 1"
- "The number rolled is 2"
- "The number rolled is 3"
- "The number rolled is 4"
- "The number rolled is 5"
- "The number rolled is 6"

This message is immediately transmitted to the chooser.

Appendix 3: Questionnaire

The following is an English translation of the control questionnaire.

Control Questions (Page 1/2)

(1) Assume: A = 12 EUR und B = 6 EUR. The number rolled is 3. The chooser chooses the number 2. How large are the payoffs?

ANSWER: The roller receives: ____ EUR

The chooser receives: ____ EUR

(2) Assume: A = 12 EUR und B = 6 EUR. The number rolled is 3.

- Which number would the chooser have to choose in order to receive the largest payoff?

  Answer: ______

- What would the payoffs be in this case?

  Payoff roller: _____

  Payoff chooser: _____

(3) Assume: A = 12 EUR und B = 6 EUR. The number rolled is 4. Suppose that the chooser will pick the number which the roller reports to him. What payoffs would result if the roller sends the following messages? (If the information provided is insufficient to determine the result, mark the field with an "X".)

- „The number rolled is 1." Payoffs: Roller: ____ Chooser: ____
- „The number rolled is 5." Payoffs: Roller: ____ Chooser: ____
- „The number rolled is 4." Payoffs: Roller: ____ Chooser: ____

Control Questions (Page 2/2)

(4) Under what circumstances will the roller receive a larger payoff than the chooser? (Please mark all correct answers!)

- If the chooser chooses a number different from the number rolled.
- If the chooser picks the number that the roller reports to him.
- If the chooser reports the number that he rolled.
- Under no circumstances.

(5) Assume: A = 6 EUR und B = 12 EUR. The number rolled is 4. Suppose that the chooser will NOT pick the number which the roller reports to him. What payoffs would result if the roller sends the following messages? (If the information provided is insufficient to determine the result, mark the field with an "X".)

- „The number rolled is 1." Payoffs: Roller: ____ Chooser: ____
- „The number rolled is 5." Payoffs: Roller: ____ Chooser: ____
- „The number rolled is 4." Payoffs: Roller: ____ Chooser: ____