Urs Fischbacher
Franziska Heusi

Lies in Disguise
An experimental study on cheating

THURGAU INSTITUTE
OF ECONOMICS
at the University of Konstanz

LIES IN DISGUISE

An experimental study on cheating

Urs Fischbacher[1] and Franziska Heusi[1]

November, 2008

Abstract

In this paper we present a new design which allows us to draw inferences on the distribution of lying behavior among the population. Participants received a dice in order to determine their payoff anonymously. Whatever they reported to have rolled, they received as payoff. 39% of the subjects were honest and maximally 22% of them were lying completely. Interestingly we found subjects who lied but who did not maximize their income by doing so. Using additional experiments, we can show that a compelling explanation for this behavior is the desire to maintain a favorable self-concept, including honesty and non-greediness.

Key words: Lie detection, honesty, deception, experimental design
JEL Classifications: C91, D63, H26.

[1] Thurgau Institute of Economics, Hauptstrasse 90, 8280 Kreuzlingen, Switzerland.
fischbacher@twi-kreuzlingen.ch; heusi@twi-kreuzlingen.ch
University of Konstanz Department of Economics, University of Konstanz, 78457 Konstanz, Germany

1

1. INTRODUCTION

If people can make themselves better off by lying, why do they not always lie? Is it just a question of monetary incentives? Is it a matter of morals? Or do people not want to be perceived as liars? Whose judgment is it they fear? In this paper, we present a new and very simple experimental design that allows detecting lies where subjects face no threat of being revealed individually. Instead we can draw inferences on the population´s overall behavior. We find a surprisingly robust pattern of lying behavior and evidence that people not only care about their income but also about maintaining favorable traits such as honesty and non-greediness. People try to disguise their lies, and they believe it works!

Every child learns that lying is condemnable. Nevertheless, lies and honesty are part of our daily lives. Bella M. DePaulo et al. (Bella. M. DePaulo et al., 1996) asked people how often they lied and found that people lie in 20 to 31 percent of their social interactions. This number seems high but we do not know the motives behind these peoples' lying behavior. Do people lie whenever it is in their interest to lie or do people only lie when it would be very harmful for them to be honest?

Honesty and especially the precise patterns of lying behavior matter in many economically relevant situations. For example, whether and how agents report their performance is relevant for auditors. Whether people report their income correctly to tax authorities is important for the design of tax regulation and for the frequency of inspections. How easily a politician becomes corrupt is relevant for voters. In most situations, the ability to accurately assess the potential lies of others is crucial for success. Standard economics assumes that people lie when it is in their material interest to do so. However, social norms and several regulatory mechanisms prohibit lying in many situations and create formal and informal incentives for honesty. We present an experiment where no formal incentives for honesty exist. Any abstention from lying therefore must be interpreted as a reaction to – possibly internalized – social rules. In this experiment, subjects are informed that they must roll a dice, which will determine a payoff, the payoff equaling 1, 2, 3, 4, and 5 CHF for the corresponding dice figure and zero if the figure is 6. Since the experimenter cannot see what number was rolled,

the subjects can report whatever number they want. Nevertheless, the distribution of the reported numbers reveals information about patterns of lying behavior within the population.[1]

Questions about humans' honesty have already attracted researchers from the field of psychology (for reviews, see Ray Hyman, 1989, John A. Podlesny and David C. Raskin, 1977) and recently also from neuroscience (for reviews, see Kamila E. Sip et al., 2008, Spemce A. Spence et al., 2004). Economists have investigated lying and honesty in games in which players can announce future moves or can reveal (not verifiable) private information. In an overview on coordination, signaling and bargaining games, Rachel Croson (2005) collected experimental evidence indicating that people often act according to the prediction of the standard economic models. She found that people lied when this increased their profit. However, in contrast to this result, she also found evidence that people tended to believe others' lies. A more applied form of research on honesty and lies has addressed the issue of tax evasion. Tax honesty can be modeled as a decision made under uncertainty (Benno Torgler, 2002). Different to the study mentioned above, Torgler found a surprisingly large number of honest people who assessed their liability correctly. There is evidence that people are more honest than models that take audit rates, audit rules and penalty rates into account predict (see e.g. James Alm and Michael McKee, 1998, James Andreoni et al., 1998).

A lot of experimental research has been done on honesty in cheap-talk, bluffing or signaling situations, where subjects have the possibility to communicate under varying experimental situations. In these interactive and therefore social settings, various explanations for why subjects sometimes abstain from lying were found. Uri Gneezy (2005) for example found that lying behavior depended on the costs it imposed on both the liar and on the one who is lied to. Other experimental evidence was interpreted as guilt aversion (Gary Charness and Martin Dufwenberg, 2006, Martin Dufwenberg and Uri Gneezy, 2000). This means that some people tell the truth because they do not want to let others down or because they would feel guilty if others' expectations were disappointed. Evidence for the existence of preference for promise-keeping or truth-telling per se was found by Santiago Sánchez-Pagés and Marc Vorsatz (2007) and Christoph Vanberg (Forthcoming).

Our experiment is most closely related to the experiments conducted by Gerald Pruckner and Rupert Sausgruber (2006) and Nina Mazar, Amir On and Dan Ariely (2008) , who interpreted

---

honesty as compliance with a given rule. Pruckner and Sausgruber (2006) collected field data on how many customers pay for a newspaper when this is sold out of a box with payment into a cash-box. They found that more than 30% of the people paid something for the newspaper. On average, people paid one third of the price requested. Nina Mazar, On Amir, and Dan Ariely (2008) also addressed the question of subjects' honesty towards the experimenter. In their experiment, subjects had to complete a test with 50 multiple-choice questions. They were paid according to the number of correct answers. In one treatment, the number of correct answers was checked by the experimenter. In other conditions, the subjects themselves corrected their sheets. The treatment conditions differed with respect to how easy it was for the experimenter to detect a potential fraud. In the most extreme condition, it was impossible because the subjects had to shred their original test. On average, subjects reported about 10% more questions solved when they had the possibility to cheat. As in our experiment, subjects did not cheat maximally. None of the subjects reported to have solved all questions. These results also indicate that the maintenance of a self-concept as an honest person might play a role when deciding to lie or not (see e.g. Roland Bénabou and Jean Tirole, 2002, Tore Ellingsen and Magnus Johannesson, 2004).

Our design adds to this research and continues the search for a more general pattern of lying behaviour. It not only allows detecting lies, but furthermore assesses the distribution of lying behaviour in a given population. Only about one fifth of the people lies fully and acts as payoff maximizers. About 39% of the subjects seem to resist the monetary incentives to lie and remain honest. Another 20% of the subjects obviously do not tell the truth but do not maximize their payoff either; we refer to this behavior as incomplete lying. We alter stakes and consequences of the lie and control for effects of learning and anonymity. The basic patterns of lying remain the same. Most interestingly, the evidence for incomplete lying is robust in every control treatment we ran. In order to understand this finding in more detail, we additionally ran two similar experiments. We find evidence that greed aversion and the belief that it is possible to disguise a lie can explain this finding of incomplete lying.

The remainder of this paper is organized as follows. In section 2 we present our experimental design and the procedure. In section 3 we present our main results. The results are discussed in the light of additional experiments in section 4. In section 5 we conclude.

2. EXPERIMENTAL DESIGN

Our experiment is a one shot single decision making situation. It took less than 10 minutes to conduct it. This is why we did not recruit subjects for this experiment only. Instead we asked other experimenters whether we could add our experiment at the end of their sessions. At the end of such an experimental session we distributed six-sided dices among the participants. They were informed not to touch the dice until requested to do so. The experimenter then told the participants that the following very short experiment had nothing to do with the experiment they had just participated in and that instructions would be given on the screen.

Subjects then read these instructions and were informed that they were going to receive an additional payoff for filling in a questionnaire and that this payoff would be different for each participant. To determine their individual payoff, the participants were requested to roll a dice. The payoff would equal 1, 2, 3, 4, and 5 CHF if the dice came up with the corresponding number and 0 CHF if the dice came up with a 6. Participants were explicitly called to roll the dice more than once in order to check whether the dice was fair. It was highlighted on every screen that only the first throw was relevant for the payoff and therefore should be kept in mind. Next, participants were requested to roll the dice and to memorize the figure rolled. On the last instructing screen, participants reported the number rolled together with the resulting payoff. Appendix A contains an English translation of the screenshots of the baseline treatment.

In this experiment, lying means reporting a different number than the one actually rolled in the first throw. It was impossible to detect lying on the individual level. The consequence of lying is only that this person will receive a different – usually higher – payoff than deserved by the rules of the game. Since the experimenter cannot see what number was rolled, subjects easily can be dishonest. We made it as obvious as possible that it was impossible for the experimenter to find out what number a subject actually rolled. First, we asked subjects to throw the dice more than once. We told the subjects that this was a possibility for them to make sure that the dices were not loaded. They were not only allowed to roll the dice twice but as many times as they wanted to. Secondly, the experimenters were not in the same room during the experiment. It was not possible for the experimenters to walk through the lab and to note the actual figures rolled. Thirdly, we wanted to make it easier for them not to tell the truth. It would be easier to report a number that was actually rolled in one of the following throws, even if it was not the payoff relevant first number. Still, it was explicitly mentioned

on every screen that the first number was the relevant number and that they had to keep it in mind throughout the experiment. Another excuse was the number six. We wanted to avoid that subjects would act according to gambling heuristics, therefore the usually wining number six was payoff minimizing.

In order to make the experiment as plausible as possible, we told the subjects that the reason for rolling the dice was to determine the payoff for filling in a questionnaire. It is clearly not very plausible to pay subjects differently for doing exactly the same task. Still, it is more plausible to let them roll the dice in order to determine a payoff for doing something instead of just letting them roll the dice and pay them without any explanation.

CONTROL TREATMENTS – STAKES, CONSEQUENCES, LEARNING & ANONYMITY

As we are interested in the general pattern and stability of lying behavior, we ran several control treatments with only slight variance within the procedures. These control treatments allow us to draw inferences on how robust lying behavior is. In these sessions, half of the subjects played the baseline experiment as explained above and the other half took part in the control treatment. This procedure guarantees that the assignment of subjects to the treatments asserts that different experiences in the preceding experiment are balanced among the control and the treatment group in our experiment. Furthermore, it avoids a recruiting bias as a consequence of different recruiting procedures for different preceding experiments.

The first two control treatments refer to the results of Gneezy (2000), who found that the consequences of lying matter to the person lying. Therefore we altered stakes and the party affected by lying. The first control treatment was a high stake treatment. We wanted to test whether patterns of lying depend on stakes. Do people lie differently when stakes are higher? Do they lie more, because the monetary incentives are higher, or less because the moral concerns are more salient? For the payoff, we only applied a factor of 3 compared to the baseline treatment, which is not very high. Rolling a 1 resulted in a payoff of 3 Swiss francs, rolling a 2 in 6 Swiss francs (= CHF), and 3 in 9, 4 in 12, 5 in 15, and 6 in 0, respectively. Accordingly, the incentives were changed in the following two opposing ways. Lying was financially more rewarding for the participant as reporting a 5 instead of a 2 gave subjects an additional benefit of 9 CHF instead of only an additional benefit of 3 CHF in the baseline situation. On the other hand, reporting a higher number also imposed higher costs on the experimenter when stakes were tripled. The second control treatment allows us to draw inference on whether it matters to subjects to whom they lie. In the externality treatment, we

studied whether it matters who is affected by the lie. In our baseline treatment, only the experimenter is affected by the lie. In this control treatment, a second subject received the rest of the 5 CHF. The situation is similar to a dictator game except that the dictators were instructed to use the dice to determine the distribution of the 5 CHF. As in the other treatments, this treatment was conducted together with a baseline treatment. This means that there were three types of subjects. One third of the subjects were in the baseline treatment, one third was in the role of dictators in the externality treatment and one third consisted of recipients in the externality treatment.

The third control treatment answers the question whether peoples´ lying behaviour is affected by learning effects caused by uncertainty about the true procedure of the experiment. Although we did our best to make sure that there is no detection risk, it might have been some subjects´ concern. We controlled for this by letting some subjects take part in the experiment more than once.

As a fourth and last control, we ran a double anonymous version of the experiment. We were interested in effects of anonymity. As long as it is possible to tell individually what someone reported, it is possible for the experimenter to draw inference on the probability of this subject being honest. Therefore, we ran a treatment where it was not only impossible to tell who *rolled* what number, but additionally absolutely impossible to tell who *reported* what number. Subjects who care about what the experimenter might think about them reporting a certain number could now be sure that it was impossible to reveal their decision on an individual level. To achieve such a situation, we had to alter the procedure in the following way. (This is the only control treatment where also the baseline treatment was changed and processed without computers. The screens were only used to explain the experiment.) At the end of the session, each subject received a dice plus an envelope with five coins each worth 1 CHF and a second neutral envelope inside. Subjects could select an envelope from a box that the experimenter presented to the subjects one after the other. We applied this procedure in order to avoid that people could believe that there are hidden marks on the envelopes that would enable us to identify the decisions. The instructions on the screen were similar to the former baseline treatments. Participants were requested to roll the dice and to take the coins gained out of the envelope, then to put the remaining coins into the spare envelope, seal it and give it back to the experimenter. In the double anonymous treatment subjects had to deposit the sealed envelopes anonymously in a box at the door; in the baseline treatment we requested subjects to leave the sealed envelope on their desk in the laboratory. Thus, in the baseline

situation the experimenter could walk through the lines, collect the envelopes and match the reported numbers with the data after the experiment. Coins and envelopes were prepared in a way to make sure that it was impossible to hear how many coins were taken out or given back. The double anonymous procedure made it as obvious as possible that we had no chance to trace back any decisions on the individual level.

PROCEDURE

The participants were students from the University of Zurich and the Swiss Federal Institute for Technology in Zurich. Sessions for this experiment were conducted at the computer laboratory of the Institute for Empirical Research in Zurich from summer 2004 until spring 2007. We had a total of 746 participations. Payments were made in cash, in Swiss francs (CHF; 1 CHF corresponding to about 80 cent at the time) and handed out to the participants immediately after the session.

The experiment was programmed using the software z-Tree (Urs Fischbacher, 2007). Recruiting was partly done by ORSEE (Ben Greiner, 2004). The sessions were run at the end of the sessions of other experiments. The recruiting process was organized for the preceding experiments. As an artifact of the procedure of adding the sessions of this experiment to other experimental sessions, we had to control for multiple participation by checking the identity of the participants after the experiments. This was done by manually comparing their names, surnames and fields of study and generating a personal ID for every person. In this way, we were able to restrict our analysis to the debutants and look at the results of those taking part a second time separately.

3. RESULTS

In this chapter, we first present the results of the baseline treatment and show the main patterns of lying behaviour. Next, we take a look at the results of the control treatments and show that the observed patterns in lying behaviour are robust to treatment variations. In a last step, we discuss in more detail the unexpected result of observed incomplete cheating.

A total of 398 participants took part in the baseline treatment as debutants. Figure 1 shows the resulting distribution of reported payoffs.

Figure 1. Percentage of reported number of subjects in baseline experiment; first participation only (stars display the significance of two-sided binomial test that the observed percentage differs from 16.7% (*=10%-level, **=5%-level, ***=1%-level).)

It is obvious that this distribution is not uniform (a Kolmogorov-Smirnov one sample test is significant at the 1 percent level ($p \ll 1\%$)). Numbers below 4 are significantly less frequently reported than the expected true value of 16.7% corresponding to the fraction of 1/6. The percentages of numbers 4 and 5 are significantly above the expected 16.7% (see binomial tests in table 1). This monotonously increasing distribution implies that some subjects tend to report more than they actually rolled. If we assume that people do not lie to their disadvantage, the positive frequency of subjects reporting zero shows that at least some people are honest. Another interesting observation is that not all lying subjects lie maximally. Significantly more than 1/6 of the subjects report 4. This is evidence that some subjects neither report the truth nor report 5. Instead they choose to report 4. Summing up, we find the following three characteristics in the pattern of behavior:

(1) **Honest subjects:** The fraction of people reporting a 0 is **positive**.
(2) **Income maximizing subjects**: The fraction of people reporting a 5 is **above 1/6**.
(3) **Incomplete cheating:** The fraction of people reporting a **4 is above 1/6.**

Next, we show that observations 1 to 3 are stable with respect to our treatment variations. Table 1 shows the distributions of reported payoffs in the control treatments in comparison to the baseline treatment. For each cell, a one-sided binomial test reporting whether the

percentage is above/below the expected true value of 16.7 % (corresponding to the fraction of 1/6) is given.

Treatment		payoff reported (in percent)					
		0	1	2	3	4	5
Baseline	n=389	6.4 ***	7.2 ***	11.6 ***	12.6 **	27.3 +++	35.0 +++
High Stake	n= 80	11.6	5.0 ***	15.0	8.8 **	27.5 ++	32.5 +++
Externality	n=78	9.0 **	12.8	9.0 **	16.7	25.6 ++	26.9 ++
Learning	n=110	4.6 ***	3.6 ***	5.5 ***	9.1 **	25.5 ++	51.8 +++
Baseline envelope	n= 140	5.7 ***	8.6 ***	10.7 **	17.1	28.6 +++	29.3 +++
Double anonymous	n=137	6.6 ***	8.8 ***	10.2 **	17.5	24.1 ++	32.9 +++

Table 1. Percentage of reported payoffs by treatment. *($^+$) stars (crosses) display the significance of one-sided binomial test that the observed frequency is smaller (larger) 16.7% (*($^+$) =10%-level, ** ($^{++}$) =5%-level, *** ($^{+++}$) =1%-level).

First, we controlled for stake effects. In the high stake treatment, subjects received a payoff which was tripled with respect to the payoff in the baseline treatment. A Wilcoxon-Mann-Whitney ranksum test for equality of distributions reveals that the distribution of reported payoffs in the high stake treatment does not differ from the one in the baseline treatment (Pr>|z| = 0.5079; n=159, high stake n=80, baseline n=79). With respect to our main results, we find that characteristics (1) to (3) still hold. We still observe unconditional honest behavior as well as lying and incomplete lying. Thus, the results remain stable even when the stakes are tripled.

Our second robustness test concerns the question whether the consequences of the lie matter in our treatment. Our externality treatment addresses this issue. We implemented this by altering the victim of the lie. Instead of the experimenters, now another subjects´ payoff was affected by the number reported. As we can see in table 1, this does not change the distribution (Wilcoxon-Mann-Whitney ranksum test for equality of distributions Pr > |z| = 0.279; n=158, externality n=78, baseline n=80). A change of the party to be affected by a lie does not alter the results (1) to (3), i.e. people do not lie more or less when other subjects are harmed by their lies.

Third, we controled for effects of learning by checking whether subjects taking part in the same baseline experiment a second time behave differently from the first time. We found that people taking part a second time in the experiment lied differently. People reported higher payoffs when they participated a second time (Wilcoxon matched-pairs signed-ranks test with

the null hypothesis that both distributions are the same leads to a Prob > |z| = 0.000). However, as we can see in table 1, the distribution of reported numbers is robust with respect to the characteristics (1) to (3) above. We still observe people reporting 0 and the fractions for reported numbers 4 and 5 are above 16.7%.

Our last robustness test investigates whether anonymity in the experiment is an issue. We expected that people would lie more the less they had to fear detection. In our baseline experiment, they could not be caught lying. Still, the experimenter could – based on the reported number – update his belief about the subjets' honesty. Our double anonymous treatment excludes this possibility, as it was impossible for the experimenter to find out what number a certain subject reported. Only the subjects themselves knew what number they rolled and what number they reported. Interestingly, our main results (1) to (3) do not change when we implement a double anonymous procedure. The fractions of subjects who choose 0 and 5 as well as the fraction of subjects who choose 4 remain significantly above 1/6. We find that slightly more subjects choose 5 and less subjects choose 4 in the double anonymous treatment compared to the baseline envelope treatment. However, this difference is not significant (Wilcoxon-Man-Withney ranksum test for equality of distributions Prob > |z| = 0.8634; n=277, double anonymous n=137, baseline n=140; Fisher exact test for equal number of choices of 4 in the two treatments, p=0.239; Fisher exact test for equal number of choices of 5 in the two treatments, p=0.305). To conclude, those subjects who reported their payoff by anonymously throwing the remainder of the five francs packed in a sealed envelope into a box did not behave any different from those in the baseline treatment, where the experimenter could match the reported number to the individual.

To sum up, the observed main characteristics (1) to (3) are robust to changes of stake, externality, learning and anonymity. In every control treatment we observe honesty, lying and incomplete lying. How can we understand this pattern of lying behavior? Especially the behavior of incomplete lying made visible by the fact that the number of reported fours remains above 1/6 in every control treatment, leaves us with an open question. Could it be that people try to disguise their lie by reporting a four instead of a five? But why would they do so even in the double anonymous situation? In the next section, we will examine these questions in greater detail.

4. DISCUSSION

The following section investigates several possible explanations why people sometimes abstain from lying and especially why people lie incompletely. As a first argument, we examine whether lying aversion can explain our findings. We can show that this type of model cannot explain incomplete cheating. Therefore, we have to change our line of reasoning. We then argue that people might try to maintain a favorable self-concept by avoiding unfavorable traits like greediness and dishonesty.

LYING AVERSION

Most experimental and theoretical studies on lying explicitly or implicitly assume that people are honest because lying causes bad feelings. For instance, Gary Charness and Martin Dufwenberg (2000) and Martin Dufwenberg and Uri Gneezy (2000) assume that people feel guilty if they harm others by lying to them. Vanberg (2008), on the other hand, shows in an experiment that people dislike the act of lying per se. We will now check whether lying aversion can explain our findings.

In a very simple model of lying aversion, utility would be the difference of income from lying and the subjectively felt cost this act imposes on the liar. It presumes that subjects balance their material payoff against disutility from lying. This disutility would be a function of the amount gained by lying, e.g. the difference of the payoff earned by lying and the payoff earned when being honest. A reasonable specification of such a model of lying aversion can explain the monotonously increasing distribution of reported numbers including characteristic (1) and (2). Assuming that subjects are heterogeneous with respect to their disutility function but that this function is monotonously increasing in the benefit of lying and is either increasing or decreasing in the marginal benefit of lying, it is possible to show that this leads to a monotonously increasing distribution of reported payoffs explaining characteristics (1) and (2), but never (3). The behavior of an income maximizing subject and the behavior of an honest person can both be considered as special cases of this model. An honest person will report the number he or she has rolled because his disutility from lying is too high. The fraction of people who reported a 0 gives us the possibility to estimate the fraction of honest people. Assuming that no person reporting a payoff of zero is lying, we can estimate the percentage of honest people to be as large as 39%.[2] A homo economicus suffers no cost when

---

[2] Assuming that unconditional honest people in fact roll a uniform distribution of numbers, it is reasonable to take the number of people reporting a 0 to estimate the percentage of honest people in each number reported. As 6.4% reported a 0, we can estimate the percentage of unconditionally honest people at 6*6.4% = 39%.

lying. Hence, he would always report a 5. Our results indicate that the percentage of people acting as homo economicus can be estimated at maximally 22%.[3] Still, lying aversion cannot explain incomplete cheating. This type of models of lying aversion predict that the fraction of reported 4 is smaller than 1/6. Therefore, we must turn our attention to a different line of arguments to explain why we observe this pattern of behavior with such robustness.

MAINTAINING A FAVORABLE SELF-CONCEPT

A different approach in modeling honesty is based on the idea that people want to appear as having favorable traits. They not only want to appear favorable in front of others but also when thinking of themselves. Examples of models based on this idea are those by Ellingson and Johannesson (2004), who assume that people want to appear generous, and those by Bénabou and Tirole (2002), who assume that people care about their self-image. What are the consequences of such a concept for the subjects' behavior in our experiment? Which favorable traits could be relevant in the situation we are looking at here? Before discussing the most obvious trait, honesty, we discuss another possibly relevant unfavorable and therefore avoided trait, namely greed. Both lines of arguments will be supported by additional experimental evidence.

Greed is supposed to be an unfavorable trait. Subjects could try to avoid appearing greedy by reporting a 4 instead of a 5. If this is an important motive, then it should also occur when subjects are not instructed to report the result of rolling a dice but when they can claim any payoff between 0 and 5. We tested for this motive with a further related experiment. The experiment is very similar to the baseline treatment of the experiment above. The participants were told that they would receive an additional payoff for filling in a questionnaire and that this payoff would not be the same for everybody. Instead of letting them throw a dice to elicit their payoffs they just had a choice of 6 different payoffs (0, 1,2,3,4 or 5 CHF). There was no incentive related to honesty to claim anything else than 5 CHF, as there was no rule telling anyone to take less. Still, only 85% claimed a payoff of 5 CHF. 4 out of 34 claimed 4 CHF, and one person claimed 1 CHF. It seems that 15% of people have a willingness to pay in order not to appear greedy. This finding can help explain the observed high frequency of 4 in our main experiments.

---

[3] In the Baseline treatment, 35% reported a 5. Assuming that nobody who has actually rolled a 5 tells anything else than 5, we can estimate that the maximal percentage of people acting as homo economicus is 22% ((35%-17%)*6/5). The multiplication with 6/5 is necessary to take into account those income maximizers who actually rolled a 5.

Honesty, on the other hand, is a favorable trait. If a 4 is assessed differently than a 5 in respect to honesty, it might be reasonable to lie incompletely and to try to disguise the lie and appear honest. Disguising a lie by only reporting 4 instead of 5 is possible only when participants have certain beliefs about the others' behavior. Strictly speaking, the assumed probability of facing a liar must be smaller when a subject has reported a 4 compared to the situation when it reported a 5. In this way, subjects might effectively claim being honest with a higher probability when reporting a 4 instead of a 5. Additionally, subjects also must feel less as a liar when reporting a 4 instead of a 5, as our results are robust even in the double anonymous setting. This is why we ran another session of our experiment eliciting the subjects' beliefs about the reported distribution. 60 subjects took part in this control treatment. Instead of asking them to roll the dice they read the whole instructions of the baseline treatment and then were asked to guess what they think people had reported. They were paid for this task dependent on the accuracy of their guess. Procedures were as follows: Subjects were informed that they would have to guess the behavior of other participants in a previously run experiment. Then they read the instructions of the baseline treatment and had to guess what percentage of participants earned which payoff. They were paid 5 CHF if they guessed every percentage correctly. Their payoff was reduced by 4 centimes (=0.04 CHF) for every percentage point of deviation from the correct percentages. (Screen shots for these instructions can be found in appendix A.)

41 subjects were inexperienced when taking part in the belief treatment, meaning that this was the first time they took part in this experimental series. 19 subjects were already experienced when asked for their beliefs. We were interested in how many subjects did not anticipate that the fraction of participants reporting a 4 was above 1/6. These participants would certainly believe that participants reporting a 4 are honest. Like this, such a person might disguise her own lie by reporting a 4. She would believe to be the only liar among those people who report a 4 and assume that she appears honest. 61% (25 out of 41) of the naïve participants did not anticipate others' behavior of incomplete lying. They reported a belief of a fraction of less than 1/6 for reporting a 4. Among the experienced subjects this percentage drops to 32% (6 out of 19). Experience obviously changes subjects' belief in respect to the anticipation of incomplete lying. However, the belief treatment shows that it is possible to disguise a lie by reporting a 4 instead of a 5, as more than 60% of the naïve subjects would believe that those reporting a 4 are honest.

Summing up, it is plausible to argue that subjects care about how their action will be assessed. The results of the double anonymous treatment indicate that the behavior of disguising lies is also applied to maintain a favorable self-concept. They do not want to appear greedy and they want to appear honest in front of both themselves and of others.

5. CONCLUSION

Imagine you rolled a 3, what would you report? You would maximize your income by reporting a 5, but this would seem too obvious and you would not want the experimenter to suspect you of being a greedy liar. And what would you tell someone else, at home or at the bus station after the experiment, who asked you what number you had reported? Would you tell this other person the truth, i.e. that you reported a 5 after having rolled a 3? Or would you instead tell that you only reported a 3? Why not just report a 4 and tell everybody that one was lucky enough to roll a 4 and therefore decided to be honest in the experiment? And by the way, wouldn't a reported 4 appear less suspicious to the experimenters? Is it possible to disguise the lie by reporting a 4?

Summing up our observations we can say the following: The pattern of lying does not change when stakes, consequences or anonymity is altered. We always observe liars, honest subjects and some subjects who lie incompletely. A model of lie aversion can explain the monotonously increasing pattern of our results. By assuming some heterogeneity in the subjects' disutility of lying it is possible to model the observed behavior of income maximizing subjects as well as that of honest subjects. Nevertheless the model fails to explain incomplete cheating. Something prevents some people from lying fully. By running two additional treatments, we were able to show that people care about having favorable traits. They do not want to think of themselves as greedy or dishonest although they are ready to lie a little bit in order to make some money. In daily life it is important to tell compelling lies as it is usually unpleasant being detected. Our results indicate that subjects do not adapt to the situation in the lab where it is irrelevant to the experimenters how compelling a lie is. People possibly have internalized this behavior. It might be easier for these people to live with the fact of having told a lie which did not harm the other party maximally. In addition, not lying maximally might leave more space for arguing having been honest. If ever anybody has to tell what he has done in that experiment, the person who can tell having reported 4 with a clear conscience might seem more credible. Further research is necessary to discriminate between these different explanations.

The design presented here is a new type of experiment with the potential to be used in further innovative settings. As we were unable to find a suitable explanation for the observed fraction of reported 4, it also leaves open the interesting question of disguised lies. As mentioned above, more experiments are needed to further investigate the motives to do so. Secondly, it would be interesting to analyze how lying is related to trust and trustworthiness. A third application is in cross-cultural studies. It would also be very interesting to see whether the pattern of disguising a lie differs between cultures. Fourthly, we now know the pattern of lying but we do not know yet how credible these lies are. Since we hypothesize disguise being an important motive for incomplete lying, experiments in which subjects asses the credibility of others' lies are a natural next step.

References

Alm, James and McKee, Michael. 1998. "Extending the Lessons of Laboratory Experiments on Tax Compliance to Managerial and Decision Economics." *Managerial and Decision Economics*, 19(4-5): 259-75.

Andreoni, James; Erard, Brian and Feinstein, Jonathan. 1998. "Tax Compliance." *Journal of Economic Literature*, 36(2): 818-860.

Bénabou, Roland and Tirole, Jean. 2002. "Self-Confidence and Personal Motivation." *Quarterly Journal of Economics*, 117(3): 871-915.

Charness, Gary and Dufwenberg, Martin. 2006. "Promises and Partnership." *Econometrica*, 74(6): 1579-601.

Croson, Rachel T.A. 2005. "Deception in Economic Experiments," In *Deception in Markets: An Economic Analysis.* ed. Caroline Gerschlager, 113-130. New York: Palgrave Macmillan.

DePaulo, Bella. M.; Kashy, Deborah. A.; Kirkendol, Susan. E.; Wyer, Melissa. M. and Epstein, Jenifer. A. 1996. "Lying in Everyday Life." *Journal of Personality and Social Psychology*, 70(5): 979-95.

Dufwenberg, Martin and Gneezy, Uri. 2000. "Measuring Beliefs in an Experimental Lost Wallet Game." *Games and Economic Behavior*, 30(2): 163-82.

Ellingsen, Tore and Johannesson, Magnus. 2004. "Promises, Threats and Fairness." *Economic Journal*, 114(495): 397-420.

Fischbacher, Urs. 2007. "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics*, 10(2): 171-78.

Gneezy, Uri. 2005. "Deception: The Role of Consequences." *The American Economic Review*, 95: 384-94.

Greiner, Ben. 2004. "The Online Recruitment System Orsee 2.0–a Guide for the Organization of Experiments in Economics." *Papers on Strategic Interaction, University of Cologne, Working Paper Series in Economics*, 10.

Hyman, Ray. 1989. "The Psychology of Deception." *Annual Reviews in Psychology*, 40(1): 133-54.

Mazar, Nina; Amir, On and Ariely, Dan. 2008. "The Dishonesty of Honest People: A Theory of Self-Concept Maintenance." *Journal of Marketing Research*, 45(6).

Podlesny, John A. and Raskin, David C. 1977. "Physiological Measures and the Detection of Deception." *Psychological Bulletin*, 84(4): 782-99.

Pruckner, Gerald and Sausgruber, Rupert. 2006. "Trust on the Streets: A Natural Field Experiment on Newspaper Purchasing," University of Copenhagen. Department of Economics (formerly Institute of Economics).

Sánchez-Pagés, Santiago and Vorsatz, Marc. 2007. "An Experimental Study of Truth-Telling in a Sender-Receiver Game." *Games and Economic Behavior*, 61(1): 86-112.

Sip, Kamila E.; Roepstorff, Andreas; McGregor, William and Frith, Chris D. 2008. "Detecting Deception: The Scope and Limits." *Trends in Cognitive Sciences*, 12(2): 48-53.

Spence, Spemce A.; Hunter, Mike D.; Farrow, Tom F. D.; Green, Russell D.; Leung, David H.; Hughes, Catherine. J. and Ganesan, Venkatasubramanian. 2004. "A Cognitive Neurobiological Account of Deception: Evidence from Functional Neuroimaging." *Philosophical Transactions of the Royal Society B: Biological Sciences*, 359(1451): 1755-62.

Torgler, Benno. 2002. "Speaking to Theorists and Searching for Facts: Tax Morale and Tax Compliance in Experiments." *Journal of Economic Surveys*, 16(5): 657-83.

Vanberg, Christoph. Forthcoming. "Why Do People Keep Their Promises? An Experimental Test of Two Explanations." *Econometrica*.

Warner, Stanley L. 1965. "Randomized Response: A Survey Technique for Eliminating Evasive Answer Bias." *Journal of the American Statistical Association*, 60(309): 63-69.

APPENDIX A - INSTRUCTIONS

Baseline screen1: Information screen.

For the following questionnaires you will receive a small additional payoff. However, this payoff is not the same for every participant. You determine your own payoff by throwing your dice twice as soon as you are asked to.
Your first throw decides on how much you receive. You can see the exact payoff from the following chart. It will remain on the screen until you have entered your throw.
The second throw only serves to make sure that the dice is working properly. You may of course throw the dice more than twice. However, only the first throw counts.
If you have any questions, please raise your hand. If you are ready, please press OK.

Number thrown:	1	2	3	4	5	6
Resulting payoff	1.00	2.00	3.00	4.00	5.00	0.00

OK

Baseline screen 2: Throw of dice.

Please throw the dice now.
Please keep in mind the first number you have thrown.

If you have thrown the dice, please press ok.

Number thrown:	1	2	3	4	5	6
Resulting payoff	1.00	2.00	3.00	4.00	5.00	0.00

OK

Baseline screen 3: Report of number and payoff.

Now please enter the number , i.e. the first number you have thrown.

Number thrown: [   ]

Resulting payoff: [   ]

Number thrown:	1	2	3	4	5	6
Resulting payoff	1.00	2.00	3.00	4.00	5.00	0.00

OK

Beliefs screen 1: Information

On the following screens you will see the instructions for an experiment where 153 participants had taken part. You will not take part in this experiment. Instead we want you to guess the bahavior of these 153 participants. If your guess is accurat you can earn an additional payoff of up to 5 swiss francs.

OK

Beliefs screen 4: Report of guess.

Please enter your guess on the payoff of the participants. If your guess is correct you earn 5 swiss francs.
For every percentage point you deviate from the correct percentage your payoff will be reduced by 4 Raps.
You will receive a minimum payoff of 50 Raps.

What percentage of participants received a payoff of 0?

What percentage of participants received a payoff of 1?

What percentage of participants received a payoff of 2?

What percentage of participants received a payoff of 3?

What percentage of participants received a payoff of 4?

What percentage of participants received a payoff of 5?

OK