

UC San Diego

UC San Diego Previously Published Works

Title

Measuring lying aversion

Permalink

<https://escholarship.org/uc/item/39s6f22v>

Journal

Journal of Economic Behavior and Organization, 93(C)

ISSN

0167-2681

Authors

Gneezy, U
Rockenbach, B
Serra-Garcia, M

Publication Date

2013-09-01

DOI

10.1016/j.jebo.2013.03.025

Peer reviewed

Measuring lying aversion

Uri Gneezy[§], Bettina Rockenbach[†] and Marta Serra-Garcia^{*}

February 24, 2013

Abstract

We introduce a new method for measuring the decision to lie in experiments. In the game, the decision to lie increases own payment independent of the counterpart's decision, but potentially at a cost for the counterpart. We identify at the individual level the decision to lie, and measure how individuals react to different incentives to lie. Furthermore we investigate how lying behavior changes over time. Our method allows us to classify people into types, including those who never lie, those who always lie, and those who react to incentives to lie. We suggest this method as a useful instrument for examining factors that influence the decision to lie.

Keywords: deception; lying; lying aversion; experiment.

JEL-codes: D03; D83; C92.

[§]Rady School of Management, University of California, San Diego and CREED, University of Amsterdam

[†]Department of Economics, University of Cologne

^{*} Corresponding author: Department of Economics, University of Munich. E-mail: marta.serragarcia@lmu.de. Phone: +49 89 2180 9738. Address: Department of Economics, Geschwister-Scholl-Platz 1, 80539 Munich, Germany.

1. Introduction

Accumulating evidence suggests people are averse to lying in economic interactions. This evidence is in contrast with the classic approach in economics, which assumes people are selfish and that lying in itself does not carry any cost (Crawford and Sobel, 1982).

Evidence in support of the positive costs of lying comes primarily from experiments that can measure behavior in controlled settings. The ability to measure lying costs and to compare the effect of manipulating the environment on these costs is important to our understanding of the factors that influence the decision to lie. Two main experimental procedures were developed in the last decade. Each method has its pros and cons. In this paper, we propose a third approach, which eliminates some of the problems with the existing methods.

The first method uses a two-player deception game in which a sender has private information and the receiver takes an action (Gneezy, 2005). The sender sends a message to the receiver, and payoffs to both players depend on the action chosen, not on the message. Gneezy (2005) showed incentives impact the decision to lie: when the sender earns more money from lying, she is more likely to lie. Moreover, increasing the receiver's loss from a lie reduces the probability that a sender will lie. Studies have adopted the game to show, for example, men are more likely to tell a lie that helps them but hurts the other (Dreber and Johannesson, 2008), and the cost of lying for some is high enough that they are unwilling to lie even when doing so would have helped both players (Erat and Gneezy, 2012).¹ A problem with this procedure is the decision to lie depends on the sender's beliefs regarding whether the receiver will follow her message. Some senders may choose to tell the truth because they expect the receiver not to follow the message (Sutter, 2009).

¹ Cappelen, Sorensen and Tungodden (2013) study the non-economic dimensions of the decision to lie, Erat (2013) examines the decision to delegate deception, and Angelova and Regner (2013), Danilov et al. (2013) and Ismayilov and Potters (2013) use this setup to study lying within the context of financial advice.

The second method for measuring lying behavior uses a non-strategic procedure to avoid the strategic considerations associated with measuring lies. In this method, a participant takes an action for which only she knows the outcome, and then reports it to the experimenter. The reported outcome, which the experimenter cannot verify, determines the participant's payoff. This method also allows the participant to lie without the experimenter's knowledge.

In Fischbacher and Heusi (2008), participants are asked to roll a six-sided die in private and then report the resulting outcome. The participants are paid an amount equal to the number they report, unless the number is six, in which case they are paid zero. Although the experimenter cannot verify the outcome of the die roll, the distribution of the reported numbers can reveal the extent of lying in the population in the way that distribution differs from the expected distribution of outcomes from a fair die roll (see also, e.g., Greene and Paxton, 2009, Fosgaard, Hansen and Piovesan, 2013, Jiang, 2013, Ploner and Regner, 2013, and Shalvi and Leiser, 2013).

The authors found some participants were honest (reported zero profits) and that significantly more than one sixth reported either a four or five. In a similar setup, Mazar, Amir, and Ariely (2008) designed an experiment in which participants were asked to answer a test with 20 math tasks, and were paid according to the number of correct answers. To establish the benchmark, in the first treatment, the experimenter checked participants' answers. In another treatment, the participants checked their tests themselves and then shredded them, preventing the experimenter from verifying the reported number of correct answers. Mazar, Amir, and Ariely (2008) find participants on average reported about 10% more questions solved when they could cheat (see also, e.g, Pascual-Ezama, Prelec and Dunfield, 2013).

This method has two notable drawbacks. First, as described above, the inference made is not based on results from individual participants but on statistical distributions, because the

experimenter does not know whether a given participant lied. A recent paper by Gibson, Tanner and Wagner (2013) avoids this drawback by informing participants, in the role of CEOs, about the true value of their firm’s earnings and asking them to report the earnings of their company to a passive market. Gino, Ayal and Ariely (2013) also avoid this drawback by asking participants to throw away their tests in a recycle bin, from which they recovered them afterwards. However, the latter papers and the papers mentioned above face a second drawback. The “victim” of the lie is not another participant, but rather the experimenter. The identity of the victim might affect decision making.

In this paper, we propose a new method that attempts to overcome the problems of the two methods described. The sender’s payoff does not depend on the receiver’s decision but only on her message, and yet has consequences for the receiver. We observe lying on an individual basis, which is crucial for obtaining individual-level results. Observing individual decisions however comes with a cost. The participants cannot disguise their behavior from the experimenter; this may evoke behavior more compliant to social norms.

2. Experimental Design

2.1. The Decision Task

Consider a pair of two participants A and B . A randomly determined integer number $1 \leq s \leq 6$, the *state*, is assigned to the pair. Each number is equally likely. Participant A is informed about the number assigned, and sends a message about this number to Participant B . Her message must be one of the following: “The assigned number is r ,” with $r \in \{1, 2, 3, 4, 5, 6\}$. Participant B receives this message and decides whether to follow it.

A ’s payoff increases linearly with the number *reported* in the message and neither depends on the state nor on B ’s decision. More precisely, her payoff is

$$\pi_A = 10 + 2 \cdot r.$$

B 's payoff depends on whether he follows A 's message, and if he does follow it, whether the reported number corresponds to the state:

$$\pi_B = \begin{cases} 10 & \text{if } B \text{ follows and } r = s \\ 0 & \text{if } B \text{ follows and } r \neq s \\ 3 & \text{if } B \text{ does not follow.} \end{cases}$$

If A only cares about her own monetary payoff, she will always report $r = 6$, independent of s . Reporting 6 yields a payoff of 22, whereas being truthful yields $10 + 2 \cdot s$. Hence the lower s is, the higher A 's monetary gain from lying. Note the behavior of participant A is our main interest in this paper because it represents the decision to lie.

The advantage of this game over the existing methods described in the introduction is that the payoff of participant A does not depend on the reaction of participant B , and yet it affects the payoffs of B and not the experimenter. If A cares about B 's payoff, her decision to lie will not only depend on her profit but also on B 's harm from a lie. Importantly, our method (as described below) allows us to get detailed data about an individual's decision to lie and how that decision depends on incentives.

2.2. Experimental Procedures

We ran the experiment at the Munich Laboratory for the Experimental and Social Sciences (MELESSA), at the University of Munich. Using ORSEE (Greiner, 2004), we recruited 72 participants, 36 of whom we randomly assigned to the role of A , and 36 to the role of B . The experiment was conducted using z-Tree (Fischbacher, 2007).

Upon arriving at the lab, participants were asked to draw a number and sit at the computer terminal with that number. The experimenter provided a print version of the

instructions and read them aloud. The English translation of the instructions can be found in the Appendix.

The experimental task was played for 24 periods, with random re-matching in groups of 12 subjects (6 *As* and 6 *Bs*). We informed participants they would meet the same counterpart four times during the experiment but would not know when. We used the strategy method in each period: *A* was asked to decide, for each possible number assigned to the pair, which message to send. Hence we get a rich data set regarding *A*'s decision to lie, as well as how this decision changes over the 24 periods. We also asked *B* to decide whether to follow each possible message sent by *A*. At the end of each period, participants received feedback about the assigned number, the message sent by *A*, *B*'s decision regarding whether to follow it, and the payoff of *A* and *B*.

Sessions lasted approximately 45 minutes and participants received their accumulated earnings, on average 15.74 Euro (sd=7.1), including a show-up fee of 4 Euro.

3. Results

In what follows, we first examine the messages sent by *As*. We then examine *Bs*' decisions regarding whether to follow the messages. Finally, we provide a classification of *As* into types, depending on the lying decisions. We perform all tests at the matching-group level. Because our comparisons will be within treatment, we use two-sided Wilcoxon signed-rank tests throughout, unless otherwise mentioned.

3.1. Messages sent by A

Figure 1 displays a scatter plot showing what message r (y axis) was sent for a given state s (x axis). The size of each circle reflects the relative frequency of each message for a given state. This frequency value is shown in the lower part of each circle.

The frequency with which A sends the true message depends on s . When the state is 1, As reveal it truthfully 36% of the time. As s increases, truthful reporting increases and reaches 60% when the state is 5. The reverse occurs with respect to reporting 6. The probability of reporting 6 decreases as the state increases, as can be seen from the estimated marginal effects in Table 1. It drops from 52% to 48% to 44% to 37% to 33%. In all states except 6, the likelihood of reporting 6 is significantly lower than in state 6, controlling for the period in the experiment and the individual's age and gender (see Table 1). Non-parametric tests (p -value=0.028 in all pairwise comparisons between $s \neq 6$ and $s = 6$) also confirm this result. Thus, in low states, in which the gain from lying is higher, lying is more frequent than in high states. This finding is in line with Gneezy's (2005) finding that the decision to lie depends on the incentives, but appears contrary to arguments that lying costs increase in the "size of the lie," measured as the distance between the reported and the true state.

Also note that for states 1 to 3, in at least 10% of cases, the As do not lie "fully." They report a number higher than the true state but do not report a 6. This finding is in line with Fischbacher and Heusi (2008), though the frequency with which it occurs in our setting is relatively small.

<Insert Figure 1>

<Insert Table 1>

Lying increases over time for the low states 1 or 2, whereas no significant trend in lying occurs for states 3, 4, or 5 (taking 6 as a baseline); see column (2) of Table 1. As shown in Figure 2, in period 1, the frequency with which A reports 6 if the state is *not* 6 is between 28% and 33% in all cases. However, lying more than doubles by period 24, when the state is 1 or 2, up to 58%, whereas it remains around 40% if the state is 4 or 5 (42% and 39%, respectively). This difference in the time-trend of lying across states could be due to As becoming more

aware of the gains from lying and choosing to take advantage of them when the payoff from telling the truth is low. Alternatively, telling the truth may require some self-control. In particular, Mead et al. (2009) and Gino et al. (2011) show that depleted individuals, after completing standard tasks such as the Stroop task, are more likely to cheat. If repeatedly telling the truth, as in our experiment, is also depleting, this could also explain our results.

<Insert Figure 2>

If *A* cares about *B*'s payoff her lying decision should not only depend on her own profit, but should also be influenced by *B*'s following behavior. From the detailed feedback on *B*s' decisions, *As* learn about the likelihood of *B*s' following. Table 2 examines the dynamics of lying, by regressing the likelihood of lying in period *t* on the use of a lie in *t-1* and other controls. In particular, we focus on the effect of a lie in *t-1* that was followed by *B* (Lied in *t-1* & *B* followed). That the likelihood of lying in *t* decreases significantly if *A* told a lie that hurt *B* in the previous period suggests that *As* also care about *B*s' payoffs. Note that also in this regression, where we control for *B*'s past behavior, the probability of lying decreases with the state, as shown in Table 1.

<Insert Table 2>

3.2. Following by B

If *A* sends a message lower than 6, *B*s follow it in most cases. On average, messages 1 to 5 are followed in 85%, 88%, 89%, 87%, and 81% of the cases, respectively. In contrast, *B*s follow 6 only 46% of the times. This difference in following between 6 and all other numbers is statistically significant (p-value=0.028, in all cases). The regression analysis provided in Table 3 also confirms this finding.

Figure 3 shows that the likelihood with which *B*s follow messages 1 to 5 over time increases from between 75% and 89% of the time in period 1 to above 91% for all messages

by period 24. Following a message of 6 does not show any trend over time. This difference in the time trend is significant for messages 3, 4, and 5, compared to message 6, when other controls are included (see Table 3). Table 3 also reveals that *Bs* react to being hurt in the past (column (3)). Interestingly, *Bs* do not react to whether they were lied to in the previous period per se. They, however, significantly reduce following behavior if an *A* lied to them in period $t-1$ and they followed.

<Insert Figure 3>

<Insert Table 3>

To examine whether following behavior is in line with its expected payoffs, we calculate for each message r the probability that the message is true $p_T(r)$. Table 4 displays this probability as well as the difference between *B*'s expected payoff of following $E\pi_B(\textit{follow}) = 10 \cdot p_T(r)$ and the payoff from not following $\pi_B(\textit{not follow}) = 3$. Except for message 6, this difference is significantly positive, that is, following results in a higher expected payoff than not following (Sign test, p-value=0.03). However, *Bs* do not earn a significantly different expected payoff when following and not following message 6 (Sign test, p-value=0.6). Thus, on average, *Bs* seem to correctly anticipate *As*' incentives to lie and take them into account when deciding whether to follow a message.

<Insert Table 4>

3.3. Classification of *As*: Lying types

In Table 5, we present a classification of the strategies submitted by *As* in each period. We classify individuals' strategies within each period to one of eight types. We then average over all periods to obtain the numbers presented in Table 5.

First, we classify the strategy to be truthful about every state. *As* follow this strategy 33% of the time over all periods (Type 1). Second, we consider the strategy consistent with the standard assumption of self-interest, namely, reporting 6 independent of the state (Type 2). Over all periods, *As* follow this strategy 28% of the time.² Third, we identify a class of four strategies for which the decision to lie depends on the consequences (Types 3 through 6). Particularly, these strategies map into messages that tell the truth about high states and exaggerate low states. In total, *As* follow this strategy 23% of the time over all periods. A strategy that maps into lies independent of the state is used 5% of the time (Type 7), whereas 11% of the time *As* lie in some states and are truthful in others, without any particular pattern (Type 8).

<Insert Table 4>

Next we examine the behavior of each participant *A* over time and check for consistency in the strategy used. If we consider all periods 1 to 24, we find 14% of *As* (i.e., 5 out of 36) choose to always tell the truth independent of state. These participants send a true message in all periods regardless of the cost. By making the monetary cost of telling the truth steeper, future research may investigate the sensitivity of this fraction to incentives.

The same number of *As* (14%) choose to always exaggerate the state if it is not 6. Hence two types can be isolated from the data: those who never lie and those who lie whenever they have a monetary incentive to do so.

The rest of the *As* try different message patterns, especially in the first periods. However, in the final quarter of the experiment (periods 19 to 24), the majority of *As* (53%) consistently follow the same type of strategy. In the final quarter of the experiment, 19% of *As* always tell the truth and 22% exaggerate in all states except 6. Further, 8% of all *As* consistently lie

² Note that this group may also include individuals who are averse to lying, but for whom the costs of lying are very low and cannot be captured within our range of payoffs. The opposite applies to Type 1 as well.

depending on the consequences, that is, behave either as Types 3, 4, 5, or 6 through the last periods. We also find a single participant *A*, who always lies.

4. Conclusion

When and why do people lie in economic interactions? The answer to these questions is important for understanding economic activities. In this paper, we offer a new approach to measuring such behavior via a simple information transmission game. In the game we design, the decision to lie is independent of beliefs regarding the reaction of the other participants. This design allows us to measure the sender's decision to lie.

As mentioned above, when a sender also cares about the payoffs of the receiver, her beliefs regarding his actions are important. If she believes that he will not follow her message, she might be more willing to lie than if she expects him to follow it. Follow up studies could measure these beliefs directly.

The design also allows us to have data on individual behavior over time. Our results show an individual's tendency to lie increases with the gains from lying. When comparing our data to previous studies we find very similar results. If the state is 1, 52% of senders lie by reporting 6 (overall periods, 33% in period 1), creating a benefit of 10 for the sender and a loss of 10 for a receiver who follows. Interestingly, Gneezy (2005) finds the same percentage of lying for the case (+10, -10). If the state is 5, only an average of 33% of senders lie by reporting a 6, which benefits the sender by 2 and costs the receiver who follows 10. Gneezy (2005) finds 17% of senders lie for the case (+1, -10). Our experiment additionally allows measuring the effect of experience. We find that when the gains are high, the frequency of lying increases with experience.

Further, we can classify people according to the evolution of their decisions over time. We find some participants are always honest, and a bit fewer senders always maximize monetary

benefits (send the message 6). A third prominent group consists of people who react to incentives—lie when the incentives to do so are high but not when the incentives are low. Gaining experience over time reduces the strictly truthful behavior and increases both the strictly selfish behavior and the rate of lying when incentives to do so are high.

Nonetheless, on an individual level, we observe a considerable aversion to lying, which strategic considerations cannot explain. Although this aversion decays slightly over time, it is still prominent in the final quarter of the experiment and occurs with the same magnitude as the strictly selfish behavior.

References

- Ai, Chungron and Edward C. Norton (2003). "Interaction terms in logit and probit models." *Economics Letters* 80 (1): 123–29.
- Angelova, Vera and Tobias Regner (2013). "Do voluntary payments to advisors improve the quality of financial advice? An experimental deception game." *Journal of Economic Behavior & Organization*, in print.
- Cappelen, Alexander W., Erik O. Sorensen and Bertil Tungodden (2013). "When do we lie?" *Journal of Economic Behavior & Organization*, in print.
- Crawford, Vincent P. and Joel Sobel (1982). "Strategic Information Transmission." *Econometrica*, 50(6): 1431–51.
- Danilov, Anastasia, Torsten Biemann, Thorn Kring and Dirk Sliwka (2013). "The dark side of team incentives: Experimental evidence on advice quality from financial service professionals." *Journal of Economic Behavior & Organization*, in print.
- Dreber, Anna, and Magnus Johannesson (2008). "Gender Differences in Deception." *Economics Letters*, 99(1): 197–99.
- Erat, Sanjiv and Uri Gneezy (2009). "White Lies." *Management Science*, 58 (4): 723–33.
- Erat, Sanjiv (2013). "Avoiding Lying: The Case of Delegated Deception." *Journal of Economic Behavior & Organization*, in print.
- Fischbacher, Urs. 2007. "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics*, 10(2): 171–78.
- Fischbacher, Urs, and Franziska Heusi (2008). "Lies in disguise: an experimental study on cheating." Thurgau Institute of Economics Research Paper No. 40.
- Fosgaard, Toke, Lars Gaarn Hansen and Marco Piovesan (2013). "Separating will from grace:

- an experiment on conformity and awareness in cheating.” *Journal of Economic Behavior & Organization*, in print.
- Gibson, Rajna, Carmen Tanner and Alexander F. Wagner (2013). “Preferences for Truthfulness: Heterogeneity among and within Individuals.” *American Economic Review* 103 (1): 532 – 548.
- Gino, Francesca, Maurice E. Schweitzer, Nicole Mead and Dan Ariely (2011). “Unable to resist temptation: How self-control depletion promotes unethical behavior.” *Organizational Behavior and Human Decision Processes* 115: 191–203.
- Gino, Francesca, Shahar Ayal and Dan Ariely (2013). “Self-serving altruism? The lure of unethical actions that benefit others.” *Journal of Economic Behavior & Organization*, in print.
- Gneezy, Uri (2005). “Deception: The role of consequences.” *American Economic Review*, 95 (1): 384–94.
- Greene, Joshua D., and Joseph M. Paxton (2009). “Patterns of neural activity associated with honest and dishonest moral decisions.” *PNAS* 106 (30): 12506-11.
- 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.
- Ismayilov, Huseyn, and Jan Potters (2013). “Disclosing advisor’s interests neither hurts nor helps.” *Journal of Economic Behavior & Organization*, in print.
- Jiang, Ting (2013). “Cheating in mind games: the subtlety of rules matters.” *Journal of Economic Behavior & Organization*, in print.
- Mazar, Nina, On Amir and Dan Ariely (2008). “The dishonesty of honest people: a theory of self-concept maintenance.” *Journal of Marketing Research* 45: 633–44.

- Mead, Nicole, Roy F. Baumeister, Francesca Gino, Maurice E. Schweitzer, and Dan Ariely (2009). "Too tired to tell the truth: Self-control resource depletion and dishonesty." *Journal of Experimental Social Psychology*, 45(3), 594–97.
- Pascual-Ezama, David, Drazen Prelec and Derek Dunfield (2013). "Motivation, money, prestige and cheats." *Journal of Economic Behavior & Organization*, in print.
- Ploner, Matteo and Tobias Regner (2013). "Self-Image and Moral Balancing: An Experimental Analysis." *Journal of Economic Behavior & Organization*, in print.
- Shalvi, Shaul, and David Leiser (2013). "Moral firmness." *Journal of Economic Behavior & Organization*, in print.
- Sutter, Matthias (2009). "Deception through telling the truth?! Experimental evidence from individuals and teams." *Economic Journal*, 119: 47–60.

Tables and Figures

Table 1. Determinants of A's decision to send message 6

	P(sending "the number drawn is 6")	
	(1)	(2)
s=1	-0.436*** [0.076]	-0.484*** [0.0805]
s=2	-0.477*** [0.071]	-0.532*** [0.0718]
s=3	-0.510*** [0.070]	-0.501*** [0.0703]
s=4	-0.570*** [0.076]	-0.565*** [0.0724]
s=5	-0.607*** [0.074]	-0.603*** [0.0753]
Period	0.003** [0.002]	-0.000676 [0.00165]
Age	-0.011* [0.006]	-0.0103** [0.00451]
Female	0.044 [0.120]	0.0484 [0.122]
s=1*period		0.00797*** [0.00284]
s=2*period		0.00818*** [0.00253]
s=3*period		0.00279 [0.00194]
s=4*period		0.00252 [0.00254]
s=5*period		0.00254 [0.00264]
Constant		1.139*** [0.153]
N	5,184	5,184
Nr. of subjects	36	36
Method	Probit	OLS
(Pseudo-)R2	0.131	0.163
Loglikelihood	-3121	-3302

Note: Column (1) displays the estimated marginal effects from the probit regression. Column (2) displays the estimated coefficients from the OLS regression, due to the difficulties of interpreting interaction terms in the probit model (see Ai and Norton, 2003). Standard errors are presented in brackets and are corrected for clustering at subject level.

Table 2. Determinants of *A*'s decision to lie

	Probability of lying		
	(1)	(2)	(3)
State (<i>s</i>)	-0.066*** [0.015]	-0.0840*** [0.0153]	-0.0902*** [0.0158]
Period	0.002* [0.001]	0.00686** [0.00253]	0.00318 [0.00222]
Age	-0.016** [0.007]	-0.0131*** [0.00401]	-0.00873*** [0.00272]
Female	0.19 [0.129]	0.184 [0.115]	0.123 [0.0743]
State*Period		-0.00136** [0.000631]	-0.00099 [0.000620]
Lied in t-1			0.400*** [0.0526]
Lied in t-1 & B followed			-0.0627* [0.0337]
Constant		0.931*** [0.182]	0.753*** [0.140]
N	5,184	5,184	4,968
Nr. of subjects	36	36	36
Method	Probit	OLS	OLS
(Pseudo-)R2	0.135	0.172	0.306
Loglikelihood	-3095	-3259	-2687

Note: Column (1) displays the estimated marginal effects from the probit regression. Column (2) displays the estimated coefficients from the OLS regression, due to the difficulties of interpreting interaction terms in the probit model (see Ai and Norton, 2003). Standard errors are presented in brackets and are corrected for clustering at subject level. Results remain qualitatively the same if the cases where the state is 6 (and lying is rarely observed) are not considered.

Table 3. Determinants of *B*'s decision to follow

	P(following)		
	(1)	(2)	(3)
Message received = 1	0.290*** [0.044]	0.319*** [0.0800]	0.312*** [0.0823]
Message received = 2	0.318*** [0.044]	0.326*** [0.0823]	0.326*** [0.0845]
Message received = 3	0.340*** [0.037]	0.326*** [0.0818]	0.326*** [0.0846]
Message received = 4	0.305*** [0.036]	0.319*** [0.0768]	0.312*** [0.0792]
Message received = 5	0.245*** [0.034]	0.238*** [0.0661]	0.243*** [0.0708]
Period	0.004* [0.002]	-0.00278 [0.00410]	-0.00201 [0.00389]
Age	0.001 [0.004]	0.000516 [0.00326]	0.000344 [0.00264]
Female	0.034 [0.080]	0.0323 [0.0824]	0.0411 [0.0751]
Period*Message=1		0.00622 [0.00462]	0.00664 [0.00452]
Period*Message=2		0.00763 [0.00471]	0.00763 [0.00466]
Period*Message=3		0.00894* [0.00477]	0.00892* [0.00488]
Period*Message=4		0.00729* [0.00391]	0.00771* [0.00410]
Period*Message=5		0.00913** [0.00399]	0.00884** [0.00415]
Followed in t-1			0.216*** [0.0499]
Received lie in t-1			0.0281 [0.0320]
Followed & received lie in t-1			-0.116*** [0.0408]
Constant		0.455*** [0.118]	0.316*** [0.110]
N	5,184	5,184	4,968
Nr. of subjects	36	36	36
Method	Probit	OLS	OLS
(Pseudo-)R2	0.126	0.15	0.198
Loglikelihood	-2316	-2259	-2021

Note: Column (1) displays the estimated marginal effects from the probit regression. Column (2) displays the estimated coefficients from the OLS regression, due to the difficulties of interpreting interaction terms in the probit model (see Ai and Norton, 2003). Standard errors are presented in brackets and are corrected for clustering at subject level.

Table 4. *B*'s expected excess payoff from following

Message r	$p_T(r)$	$E\pi_B(\text{follow}) - \pi_B(\text{not follow})$	N
$r=1$	0.80	4.95	6
$r=2$	0.82	5.21	6
$r=3$	0.79	4.86	6
$r=4$	0.82	5.22	6
$r=5$	0.79	4.87	6
$r=6$	0.31	0.15	6

Table 5. Classification of *A*'s strategies

Strategy of Participant <i>A</i>	Over all periods	First 6 periods (1-6)	Last 6 periods (19-24)
(1) Always truthful	33%	37%	31%
(2) Always report 6 ^a	28%	26%	31%
(3) Exaggerate 1; truthful about 2,3,4,5,6 ^b	1%	2%	0%
(4) Exaggerate 1, 2; truthful about 3, 4, 5, 6 ^b	6%	1%	9%
(5) Exaggerate 1, 2, 3; truthful about 4, 5, 6 ^b	10%	10%	8%
(6) Exaggerate 1, 2, 3, 4; truthful about 5, 6 ^b	6%	6%	6%
(7) Always lie	5%	4%	6%
(8) Other	11%	14%	7%

Note: ^a The percentages remain similar if we also include *As* who exaggerate all states except 6 but do not necessarily send 6 in each state (change by less than 1%). ^b The percentages remain similar if we also include *As* who underreport when lying (change by less than 1%).

Figure 1: Messages sent by *A*

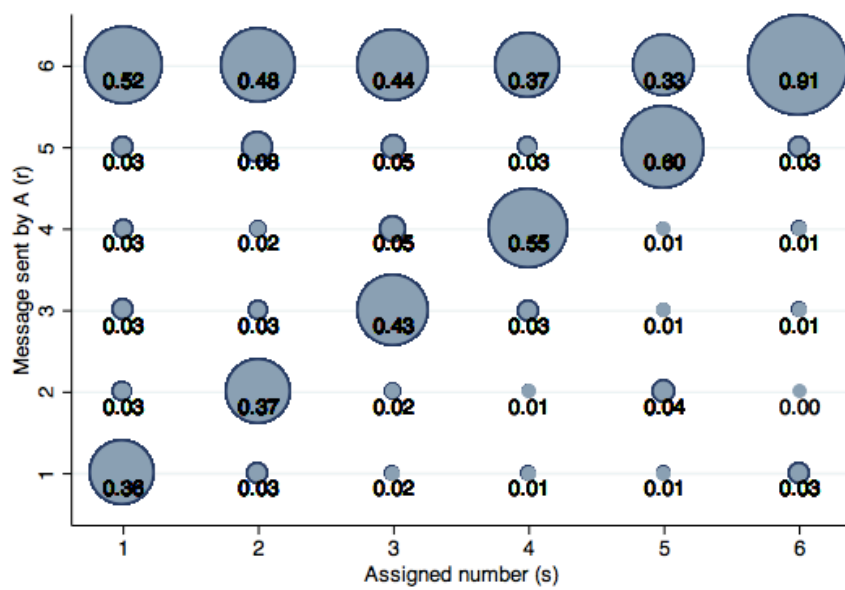


Figure 2: Truthfulness and lying of A over time


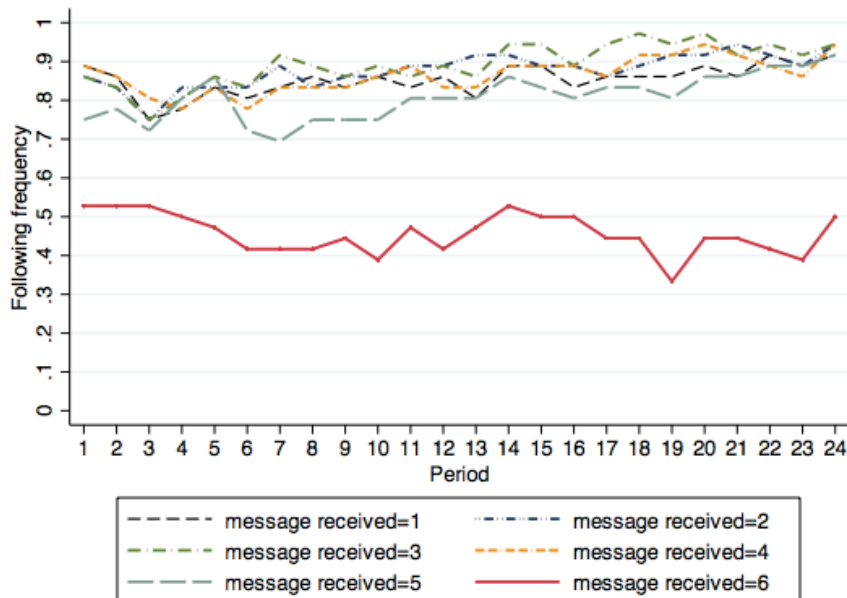


Figure 3: Following behavior by B



Appendix: Instructions

Welcome to the experiment!

Please read these instructions carefully. You may earn a considerable sum of money, depending on the decisions you and the other participants make in this experiment.

Participants, Periods and Pairs

- During the entire experiment you play in a group of 12, consisting of 6 participants A and 6 participants B. Neither participant knows the identity of the other participants.
- At the beginning of the experiment you are informed whether you are participant A or participant B. This will not change in the course of the experiment.
- The experiment consists of 24 periods.
- In every period the computer randomly allocates the participants of your group to six pairs, each consisting of one participant A and one participant B. In the course of the experiment every participant A will be matched to every participant B in exactly four periods (and vice versa).
- The computer will randomly assign the numbers 1, 2, 3, 4, 5, 6 to these six pairs. No number is assigned more than once. Participant A is informed about the number assigned to the pair and sends a message about that number to participant B. Participant B, who is not informed about the number assigned to the pair, has to decide whether he/she will follow the message of participant A. Detailed descriptions of the course of the experiment and the resulting payments will follow.

Decisions

- For every possible number assigned to the pair participant A has to formulate a message to participant B about the assigned number. The message does not have to contain the actual assigned number:

If the assigned number is	1	2	3	4	5	6
then the message to participant B is:						
The assigned number is	_____	_____	_____	_____	_____	_____

- Before participant B receives the message of participant A, participant B has to decide for all possible messages whether or not he/she will follow the message:

- Corresponding to the number assigned to the pair, the message of participant A is sent to

If participant A sends the message	Then your decision is
“The assigned number is 1”	to follow <input type="radio"/> <input type="radio"/> not to follow
“The assigned number is 2”	to follow <input type="radio"/> <input type="radio"/> not to follow
“The assigned number is 3”	to follow <input type="radio"/> <input type="radio"/> not to follow
“The assigned number is 4”	to follow <input type="radio"/> <input type="radio"/> not to follow
“The assigned number is 5”	to follow <input type="radio"/> <input type="radio"/> not to follow
“The assigned number is 6”	to follow <input type="radio"/> <input type="radio"/> not to follow

participant B and the corresponding decision of participant B is carried out.

Payments

- Participant A receives 10 points plus the doubled amount of the number sent in the message in points. That is 12 points if participant A sent the message that the assigned number is 1; 14 points if participant A sent the message is that the assigned number is 2; and so on.
- If participant B follows the message of participant A, participant B receives 10 points in case the message contains the actual number assigned to the pair. Otherwise, participant B receives 0 points.
- If participant B does not follow the message of participant A, participant B receives 3 points.

Information

- At the end of every period you will be informed about the results of your pair: the assigned number, the message of participant A to participant B, the decision whether participant B followed the message of participant A, the payment of participant A and the payment of participant B.
- Moreover, you obtain a table containing the same information of all your plays from the previous periods.

Payment in Euro

At the end of the experiment you will receive 1 Euro per 25 points. This payment and the additional 4 Euros for showing up in time are paid immediately after the experiment ends.