CESifo WORKING PAPERS

7368
2018

November 2018

Lying and Reciprocity

Simon Dato, Eberhard Feess, Petra Nieken

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 7368*
Category 13: Behavioural Economics

Lying and Reciprocity

Abstract

Recent literature has shown that lying behavior in the laboratory can well be explained by a combination of lying costs and reputation concerns. We extend the literature on lying behavior to strategic interactions. As reciprocal behavior is important in many interactions, we study a theoretical model on reciprocity where a player's altruism depends on her perception of the other player's altruism towards herself. We analyze a sequential two-player contest and vary the second mover's information on the first movers lying behavior. This allows us to derive predictions on the second mover's behavior which we test empirically in a large scale online experiment and in the laboratory. In both experiments, the second mover's lying propensity does not depend on whether the first mover has (possibly) lied or not. This robust behavioral pattern provides strong evidence that reciprocity does not play a role for lying behavior in our setting.

JEL-Codes: C900, D820, D910.

Keywords: private information, lying, reciprocity.

Simon Dato
University of Bonn
Department of Economics
Adenauerallee 24-42
Germany – 53113 Bonn
simdato@uni-bonn.de

Eberhard Feess
Victoria University of Wellington
School of Economics and Finance
23 Lambton Quay (Pipitea Campus)
New Zealand – Wellington
eberhard.feess@vuw.ac.nz

*Petra Nieken**
Karlsruhe Institute of Technology
Institute of Management
Chair of Human Resource Management
Kaiserstr. 89
Germany – 76133 Karlsruhe
petra.nieken@kit.edu

*corresponding author

Declarations of interest for all authors: none.

1 Introduction

A vastly growing experimental literature has shown that many individuals report their private information at least partially truthfully in settings where lying increases the own payoff and can neither be detected nor punished (Gneezy, 2005; Charness and Dufwenberg, 2006; Mazar et al., 2008; Fischbacher and Föllmi-Heusi, 2013; Abeler et al., Forthcoming). Recently, it has been demonstrated that lying behavior observed in the laboratory can well be explained by a combination of (internal) lying costs and (external) reputation costs for perceived dishonesty (Abeler et al., Forthcoming; Khalmetski and Sliwka, 2017; Gneezy et al., 2018). Most of the literature, however, focuses on isolated individual decision-making or group settings where subjects can lie in concert to maximize their joint payoff (Muehlheusser et al., 2015; Kocher et al., 2017), while many real-world applications are characterized by strategic interactions with negatively correlated payoffs. For instance, competing for a job can be seen as a contest where applicants may exaggerate their own experience or capabilities, thereby improving their chances at the expense of others. Similarly, people applying for scarce rental property may pretend that their own preferences and customs are in line with the landlords' requests. Also tax evasion, often used as an example for lying in individual decision making, can be seen as a public good game with private information where the own declaration may be influenced by the whole population's behavior.

In addition to internal lying costs and reputation concerns, the propensity to lie in strategic interactions is likely to be influenced by reciprocity considerations: Lying in competitive settings may be seen as a hostile act which triggers negative reciprocal behavior by responding to a lie with a lie. Lying also violates a social norm, and as reciprocity is often seen as a "key mechanism for the enforcement of social norms" (Fehr and Gächter, 1998, p. 854), responding to a lie with a lie may be justified as a socially beneficial act that harms those who deserve it. Furthermore, lying in response to a lie may be perceived as a countermeasure that restores fairness.[1]

To study the impact of reciprocity on lying behavior, we consider a sequential two-player contest. Each player privately observes the realization of a binary lottery with the outcomes *low* and *high*, and sequentially report their outcome. The second mover (SM) always observes the report of the first mover (FM) before reporting the own outcome. If only one player reports high, this player receives the winner prize and the other player the loser prize. In

[1]For instance, Lance Armstrong justified his behavior by arguing that doping in professional cycling is widespread, so that it provides a "level playing field" that allows the most capable athlete to win (see e.g. https://www.usatoday.com; November 28, 2017), and similar justifications are brought about by Siemens after the bribery scandal in Argentina (Wenzhong and Limin, 2012).

case of both players reporting the same outcome, a fair coin flip determines the winner. Importantly, this setting ensures that the monetary incentive to lie is independent of the other player's report as lying always increases the probability of receiving the winner prize by 50%. Thus, players who care only about their monetary payoff and face, in addition, an internal disutility from lying would not condition their behavior on the other player's report. Therefore, our setting allows us to isolate the impact of reciprocity.

In our model, we follow the literature by assuming that players incur internal lying costs (see, e.g., Gneezy et al. (2018)), but add the type-based reciprocity model proposed by Levine (1998). In this model, players also care about the other player's payoff and are either altruistic or spiteful. Reciprocity is then incorporated by assuming that a player is more altruistic the more altruistic the other player is towards herself. We study how the report of the FM influences the SM's lying behavior, depending on whether the FM's report allows the SM to up-date her belief on the FM's altruism type.

Specifically, we study three scenarios. The SM always has the possibility to lie but the scenarios differ for the FM. In the *Single Lie* scenario, the FM has no possibility to lie. As her report then provides no signal of altruism for the SM, the SM's lying frequency is independent of the FM's report in any equilibrium. In the *Both Lie* scenario, both players have the possibility to lie, so that the report of the FM provides a stochastic signal on her type. We show that in any equilibrium, the SM's propensity to lie is higher if the FM has reported the high outcome than if she has reported the low outcome. The reason is that, if some FM types lie in equilibrium, a low report by the FM is a sufficiently informative signal of high altruism. A high report, by contrast, signals low altruism and increases the SM's propensity to lie (negative reciprocity).[2] Then, there exist SM types which tell the truth after a low outcome if and only if the FM has reported the low outcome.

In addition to *Both Lie* and *Single Lie*, we consider the *Observable Lie* scenario, in which lies are identifiable with certainty as the actual outcome of the FM's lottery is also observable to the SM. Since a fake high report of the FM is then a clear signal of low altruism, such a report maximizes the SM's propensity to lie. However, the SM's lying propensity is also higher after a true high report compared to a true low report, because all FM truthfully report a high outcome, while only high altruism types truthfully report a low outcome.

We test our predictions with a large scale online experiment on Amazon Mechanical Turk (MTurk). We implement three treatments based on the scenarios above. As an additional

[2]Our theoretical predictions are not specific to the type-based reciprocity approach. Predictions generated with the intentions-based reciprocity approach (Rabin, 1993; Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006) are qualitatively identical, thereby strengthening our theoretical predictions.

robustness check to the online experiment, we conduct a controlled lab experiment. We find no evidence at all that reciprocity affects lying behavior, neither on MTurk nor in the lab, and neither for within- nor for between treatment analyses. In all three treatments, the lying frequency of the SM does not depend on whether the FM has reported a high or a low outcome. While this is expected for *Single Lie* as the FM cannot influence her own report anyway, it is surprising for *Both Lie* where a low report signals altruism, but is nevertheless not rewarded by the SM. This is even more striking for the *Observable Lie* treatment where we again find neither a significant impact of the FM's report nor of the FM's actual outcome on the behavior of the SM. Thus, even when a SM knows with certainty that a high report by the FM is a lie, this does not increase her own willingness to lie. All of our between treatment-comparisons also turn out to be insignificant. For instance, the lying frequency of SM after high FM report is independent of whether the SM knows that the report is true (in *Single Lie*), possibly a fake (in *Both Lie*), or certainly a fake (in *Observable Lie*).

Overall, or results strongly suggest that lying in contests is driven by lying costs and an unconditional desire to win rather than by reciprocity considerations. This is in line with Benistant and Villeval (2017) who analyze a two-player simultaneous real-effort tournament and find that lying behavior is neither affected by group identity nor by whether lying increases the own or decreases the opponent's final score.[3] Several other recent experiments also document the low volatility of lying behavior. Abeler et al. (Forthcoming) show that the lying behavior in an individual decision-making context is not affected by the belief about the lying propensity of other subjects in an unrelated experiment, and conclude that individuals do not care about social norms when deciding to lie. Similarly, Gächter et al. (2017) find that the behavior of peers in a dictator game strongly influences what is perceived as fair, but has little effect on the actual behavior.

In addition to studies about the identification of preferences that shape lying behavior, our paper is also related to the more general literature on unethical behavior in contests (Harbring and Irlenbusch, 2008; Carpenter et al., 2010), which has identified various treatment effects. Belot and Schröder (2013) and Faravelli et al. (2015) find that lying frequencies in contests are larger compared to a set-up with piece rate payments, which can be attributed to larger monetary incentives due to the discontinuity of payoffs. The same holds for the finding by Conrads et al. (2014) that the lying frequency in a die-roll contest increases in

[3]Doğan and Roggema (2016) consider simultaneous real-effort tournaments where either both or just one of the contestants could lie and find no treatment effect. This, however, follows already from the fact that hardly anyone lies, which is most likely due to the feature that subjects knew that their results would be evaluated by others.

the prize spread. Charness et al. (2013), however, find that rank-order tournaments induce misreporting and sabotage even when payments are independent of ranks, which indicates that subjects draw a positive non-monetary utility from winning a contest.

Our contribution is twofold: First, we add a model investigating the effects of reciprocity on lying behavior in strategic interactions. Second, we test the behavioral predictions of the model in a lab and a large scale online experiment.

The paper is organized as follows: section 2 introduces the model whereas the experimental design is presented in section 3. We present the results in section 4, and conclude and discuss potential future research in section 5.

2 The model

2.1 Game structure

In our model, two players $1, 2$ take part in separate lotteries. For player $i \in \{1, 2\}$, lottery L_i yields a high outcome $e_i = h$ with probability p_i and a low outcome $e_i = l$ with $1 - p_i$. After observing the own outcome, players sequentially report their outcome $r_i \in \{h, l\}$. The payoff distribution depends on both reports: if the two players report different outcomes, then the one who reports h receives the winner prize w_h, while the other player receives the loser prize $w_l < w_h$. If both report the same outcome, then each of them receives w_h with probability $\frac{1}{2}$, in which case the other player receives w_l.

We consider three different scenarios, which have in common that player 2 (the second mover, SM) privately observes her outcome e_2 and the report r_1 of player 1 (the first mover, FM), and can then report any result irrespective of her true outcome. Thus, the SM can always condition her report on the FM's report and has the opportunity to lie. The scenarios differ with respect to (i) the action space of player 1 (FM) for a given outcome e_1 and (ii) the degree of information the SM receives before reporting her outcome: in the *Single Lie (*SL*)* scenario, it is common knowledge that the FM cannot lie ($r_1 = e_1$), so that the SM learns the FM's outcome from her report. In the *Both Lie (*BL*)* scenario, the FM privately observes e_1 and then reports $r_1 \in \{l, h\}$, i.e., lying ($r_1 \neq e_1$) is possible and cannot be detected. In the *Observable Lie (*OL*)* scenario, both players observe e_1 and the FM reports $r_1(e_1) \in \{l, h\}$, so that the SM knows whether the FM lied or not. The informational structure is common knowledge in all scenarios.

2.2 Utility functions

A player's overall utility emerges via three possible channels: first, she receives a monetary utility from either the winner or the loser prize. We assume a utility function that is increasing in the prize such that $u(w_l) = u_l < u_h = u(w_h)$. The utility difference is denoted by $\Delta \equiv u_h - u_l$. Second, a player i who reports $r_i \neq e_i$ incurs (internal) lying costs of c. Lying costs are identical for all types. Third, the player has social preferences in the sense of Levine (1998): If player $i \in \{1, 2\}$ obtains monetary utility u^i and player $j \in \{1, 2\}, j \neq i$ receives monetary utility u^j, the utility of player i is

$$v_i = u^i + \frac{a_i + \lambda a_j}{1 + \lambda} u^j - Ic,$$

where I is an index variable that takes value one if $r_i \neq e_i$ and zero otherwise.

The overall utility of player i depends on the monetary utility of player j in two ways. First, player i may be altruistic or spiteful towards player j, which is captured by $a_i \in [-1, 1]$: the higher a_i, the more player i cares about player j's monetary utility. Second, how much player i cares about player j may also depend on player j's attitude towards player i. The more altruistic player j is towards player i, the more player i cares about player j. This reciprocity effect is captured by λa_j, where $\lambda \in (0, 1]$ measures the strength of reciprocity relative to pure altruism. As $-1 \leq \frac{a_i + \lambda a_j}{1+\lambda} \leq 1$, no player has an incentive to reallocate money from herself to the other player. The altruism coefficient is drawn from an iid cdf $F(a)$ with mean zero, which is common knowledge.

The realization of a_i is private information but players may extract valuable information from observed actions. Consider the SM, who anticipates the strategy of the FM. After observing report r_1 of the FM, the SM updates her expectation of the FM's type, $E_2[a_1|r_1]$.[4] Her expected utility of reporting r_2 is then

$$v_2(r_2|r_1) = E[u^2|r_1, r_2] + \frac{a_2 + \lambda E_2[a_1|r_1]}{1 + \lambda} E[u^1|r_1, r_2] - Ic.$$

In the same way, the FM anticipates the strategy of the SM. Hence, for a given report r_1, she anticipates the probability that the SM lies after observing the low outcome. We denote this probability by $q_2^{r_1}$. Moreover, she associates only those types with a given report r_2 for which this report is the optimal action. Hence, the regard for the SM depends on the anticipated report of the SM and is determined by an updated belief $E_1[a_2|r_1, r_2]$ about the

[4] In the *OL* scenario, the SM observes the outcome of the lottery of the FM, e_1, in addition to the reported outcome r_1. The expectation about the type of the FM is then given by $E_2[a_1|r_1, e_1]$.

SM's altruism type. The expected utility of reporting r_1 is then given by

$$v_1(r_1) = [p_2 + (1-p_2)q_2^{r_1}]\left[E[u^1|r_1,h] + \frac{a_1 + \lambda E_1[a_2|r_1,h]}{1+\lambda}E[u^2|r_1,h]\right] + (1-p_2)(1-q_2^{r_1})\left[E[u^1|r_1,l] + \frac{a_1 + \lambda E_1[a_2|r_1,l]}{1+\lambda}E[u^2|r_1,l]\right] - Ic.$$

2.3 Analysis

As players are not willing to incur lying costs for reallocating money to the other player, they *always* report truthfully when the outcome is high. Thus, we can safely restrict attention to the behavior after low outcomes. We start with the *Single Lie* scenario. Following backwards induction, we first consider the SM, who anticipates the FM's strategy in equilibrium. Comparing the expected utilities from reporting high and low shows that

$$v_2(h|r_1) \geq v_2(l|r_1) \quad \Leftrightarrow \quad a_2 \leq (1+\lambda)(1-\frac{2c}{\Delta}) - \lambda E_2[a_1|r_1] \equiv \tilde{a}_2^{r_1},$$

such that for a given report r_1, the SM lies if and only if her degree of altruism is sufficiently low. Her critical altruism type, $\tilde{a}_2^{r_1}$, decreases in lying costs c and increases in the prize spread Δ. Moreover, $\tilde{a}_2^l < \tilde{a}_2^h \Leftrightarrow E_2[a_1|l] > E_2[a_1|h]$. Thus, the SM is less likely to lie after observing a low than a high report by the FM if and only if a low report of the FM signals a higher degree of altruism. As the FM's report provides no signal about the altruism type in the Single Lie scenario, the behavior of the SM does not depend on the FM's report either:

Proposition 1 *In the SL scenario, $\tilde{a}_2^l(SL) = \tilde{a}_2^h(SL)$: the willingness to lie of the SM is independent of the FM's report.*

All proofs are relegated to the appendix.

Next, consider the *BL* scenario. As the FM may misreport a low outcome as high, her report may affect the SM's willingness to lie via her updated expectation about the FM's degree of altruism, $E_2[a_1|r_1]$. We show that, in any equilibrium strategy of the FM, there exists a cut-off $\tilde{a}_1$ such that $v_1(h) \geq v_1(l) \Leftrightarrow a_1 \leq \tilde{a}_1$: if some type $\hat{a}_1$ refrains from lying, then all types $a_1 > \hat{a}_1$ do not lie as well. Thus, three types of strategies might be supported in equilibrium: (i) the FM always reports truthfully ($\tilde{a}_1 < -1$), (ii) the FM lies if and only if $a_1 \leq \tilde{a}_1$ ($\tilde{a}_1 \in [-1, 1)$), and (iii) the FM always lies ($\tilde{a}_1 \geq 1$).

Consider first the strategy in which the FM always reports truthfully. Then, $\tilde{a}_2^l(BL) = \tilde{a}_2^h(BL)$ as in the SL scenario. Second, if $\tilde{a}_1 \in [-1, 1)$, then some altruism types of the FM lie and some do not. As the SM associates only high altruism types with a low report, she now cares about the FM's report. Formally, $E_2[a_1|l] > E_2[a_1|h]$ and therefore $\tilde{a}_2^l(BL) < \tilde{a}_2^h(BL)$. Third, suppose the FM always lies. As the SM only observes high reports in equilibrium, the report of the FM is not informative such that $E_2[a_1|h] = E[a]$. The expectation about the type of the FM conditional on observing a low report is off-equilibrium and can therefore take any value. Applying the D1-criterion (Cho and Kreps, 1987) requires to associate only that type with a deviation to the low report for which this deviation is most profitable. As this is the most altruistic type, $E_2[a_1|l] = 1$. Then, $E_2[a_1|l] > E_2[a_1|h]$ and therefore $\tilde{a}_2^l(BL) < \tilde{a}_2^h(BL)$. To sum up, after observing a high report, the SM is always at least as much inclined to lie than after observing a low report. In addition, whenever the high report of the FM signals a lower degree of altruism than the low report, the SM is strictly more inclined to lie after observing a high report.

Proposition 2 *In the BL scenario, $\tilde{a}_2^l(BL) \leq \tilde{a}_2^h(BL)$: the willingness to lie of the SM after observing a high report weakly exceeds the willingness to lie after observing a low report. The relation is strict as long as at least one of the reports of the FM is informative about her type.*

Moreover, comparing the results from the *SL* and the *BL* scenario yields the following Proposition:

Proposition 3 *Comparing $\tilde{a}_2^{r_1}(BL)$ and $\tilde{a}_2^{r_1}(SL)$, we get:*

a) $\tilde{a}_2^l(BL) \leq \tilde{a}_2^l(SL)$: the willingness to lie of the SM after observing a low report is lower in BL than in SL.

b) $\tilde{a}_2^h(BL) \geq \tilde{a}_2^h(SL)$: the willingness to lie of the SM after observing a high report is higher in BL than in SL.

Next, in the *OL* scenario, the SM knows with certainty whether the FM has lied or not. Her strategy then specifies a report for each combination of r_1 and e_1, which depends on the updated expectation about the FM's altruism type, $E_2[a_1|r_1, e_1]$. If $e_1 = h$, the FM always reports high which hence provides no signal, $E_2[a_1|h, h] = 0$. As the SM associates the most altruistic type with a deviation of the FM to the low report, $E_2[a_1|l, h] = 1$, she is more inclined to lie after a truthful high report than after a downward lie, $\tilde{a}_2^{h,h} > \tilde{a}_2^{l,h}$. If $e_1 = l$, the FM again either always lies, lies if and only if $a_1 \leq \tilde{a}_1$, or always reports truthfully. Similar to *BL*, the SM associates a higher altruism type with a low report in any equilibrium and is, for a given outcome of the FM's lottery, less inclined to lie if the FM reports low: $\tilde{a}_2^{h,e_1} > \tilde{a}_2^{l,e_1}$, with $e_1 \in \{l, h\}$.

Consider next, for a given report of the FM, the impact of the FM's actual outcome on the SM's willingness to lie. A high report of the FM is not informative about her type if $e_1 = h$, whereas it signals (at least weakly) spitefulness if $e_1 = l$. In the same way, a true low report yields no up-date whereas a downward lie signals altruism. In sum, for a given report of the FM, the SM is always less inclined to lie if the actual outcome of the FM's lottery was high: $\tilde{a}_2^{r_1,h} \leq \tilde{a}_2^{r_1,l}$, with $r_1 \in \{l, h\}$.

Finally, conditional on truth-telling, the SM is less inclined to lie if the FM has reported low: $\tilde{a}_2^{h,h} \geq \tilde{a}_2^{l,l}$. As the FM had an incentive to report high, her low report is a stronger signal of altruism than a truthful high report. In the same way, a downward lie signals altruism whereas an upward lie (weakly) signals spitefulness: $\tilde{a}_2^{h,l} \geq \tilde{a}_2^{l,h}$.

Proposition 4 *In the OL scenario,* $\tilde{a}_2^{l,h}(OL) \leq \tilde{a}_2^{l,l}(OL) \leq \tilde{a}_2^{h,h}(OL) \leq \tilde{a}_2^{h,l}(OL)$.

Finally, comparing the results of the *OL* and *SL* scenarios yields the following Proposition:

Proposition 5 *Comparing the results of the OL and SL scenarios, we get* $\tilde{a}_2^{l,l}(OL) \leq \tilde{a}_2^{l}(SL)$*: the willingness to lie of the SM after observing a truthful low report is lower in OL than in SL.*

3 Design and treatments

Based on the theoretical analysis presented in section 2, we designed three treatments: *SL* (*Single Lie*), *BL* (*Both Lie*) and *OL* (*Observable Lie*). In all treatments, two contestants take part in a sequential contest. Each subject has to report the outcome of a die roll to determine the winner of the contest. If the result of the six-sided die roll is 1 to 4, the subject should report "low." If the outcome is 5 or 6, the subject should report "high." In case one constant has reported "high" while her opponent has reported "low," the contestant with the high report is the winner. In case of a tie, a random draw decides who wins. The winner prize is \$1.10 and the loser prize \$0.10. Given the sequential structure of the game, one subject is in the role of the FM and the other subject is in the role of the SM. In all treatments, the SM rolls the die in private after observing the FM's report. As the die roll is made in private, nobody except the SM knows the true outcome of the die roll. The SM, therefore, can deviate in her report from the true outcome.

The three treatments follow the three scenarios analyzed in our model and are executed online via an experimental software. In the *SL* treatment, the FM does not roll a die in private. Instead, the die roll is simulated by the experimental software and the report is filled in automatically. Hence, the FM cannot lie. In the *BL* treatment, the FM rolls a die in private and fills in the outcome of the die roll such that she is able to lie. The SM is only informed about the FM's report and therefore cannot detect a lie. In the *OL* treatment, the FM's die roll is simulated by the experimental software as in the *SL* treatment. However, the report is not filled in automatically. Instead, the FM can fill in the report herself after observing the true outcome of the simulated die roll. The SM is then informed about both, the actual outcome of the die roll and the FM's report and can thus detect a lie in the OL treatment. In each treatment, it is common knowledge whether the FM has the possibility to lie and whether the SM is able to detect a lie.

In all treatments, subjects filled out a short survey after stating their report. First, we elicited the subjects' beliefs. FM's were asked to estimate the fraction of (upward) lying SM's. SM's were asked to estimate the fraction of (upward) lying FM's in the *BL* treatment. In the other treatments, this belief is not relevant as there is no uncertainty with respect to the FM's lying behavior. Instead, SM's were asked to estimate the fraction of (upward) lying SM's. Next, all subjects answered questions regarding their values and preferences. We used the subscale Honesty-Humility from the HEXACO to measure fairness, sincerity and greed avoidance (Ashton and Lee, 2009). We also elicited positive and negative reciprocity as well as beliefs about reciprocity (Perugini et al., 2003; Dohmen et al., 2009). In addition, we elicited demographic variables such as gender, age, race and educational background.

Since lying of the SM is never observable, we need a large sample size. Given the simple structure of the game and the short amount of time needed to complete it, we chose to conduct the experiment on Amazon Mechanical Turk (MTurk). On MTurk, business owners and researchers can post simple and short tasks (aka HITs) that need to be executed by humans. Workers that are registered on the platform can browse the HITs and decide whether or not they want to work on the posted task for the offered compensation. MTurk is now a widely used platform for conducting research and we follow the typical guidelines for behavioral research in this environment (Horton et al., 2011; Berinsky et al., 2012; Mason and Suri, 2012; Crump et al., 2013; Paolacci and Chandler, 2014; DellaVigna and Pope, 2018).

Our HIT announced a scientific study that consists of a short survey and a simple game. Each subject earned $1 for answering the survey and knew that she could earn an additional payment for the simple game. All subjects interested in working on our task followed a link taking them to the first page of the survey (hosted on unipark).[5] The survey started with the instructions for the simple die rolling game. If a subject did not want to participate, she was free to reject the HIT without any consequences. If a subject agreed, the die rolling game

[5]The instructions can be found in the appendix.

started followed by the survey questions. Due to the sequential structure of the game and to avoid attrition problems in the matching, we divided each treatment into two parts. The first part encompasses the game for the FM's. Once a session with the FM's was finished, we conducted the session for the SM's based on the reports of the corresponding FM's. All subjects received their fixed payment for answering the survey and the bonus for the die rolling game within two days. The bonus payment resulting from the die rolling game was sent with a message informing the subject about the outcome of the game (winner or loser and the corresponding reports.)

To be eligible to work on our HIT, subjects had to be located in the US. In addition, they had to have at least 500 approved HITs and a HIT approval rate of 95%. We use a between subject design, therefore each subject was only allowed to participate once. We implemented measures to prevent them from participating twice, restart treatments, or self-select into treatments. In total, 2400 subjects participated in our study.[6] On average, they spent 341 seconds on our HIT and received $1.60 including the fixed compensation for the survey of $1. Our participants were on average 36 years old and more than 60% posses an undergraduate or post-graduate degree.

Even though many studies have shown that the behavior of the MTurk sample does not systematically differ from the results obtained in a fully controlled lab experiment, conducting the experiment on MTurk reduces the amount of control over environmental variables (see, e.g., Arechar et al. (2018)). We, therefore, conducted a standard lab experiment to investigate the robustness of the observed behavior.[7] The robustness check contains two treatments *SL_lab* and *BL_lab*, which have a similar game structure as the corresponding treatments above.

We did not use a die roll in the lab experiment to avoid that subjects might feel forced to roll the die in their cubicle to create a sound. Instead, subjects were given a real large envelope which contained 10 smaller envelopes. In each envelope, there was either a green or a pink card. The subjects knew that 7 envelopes contained a pink card, while three contained a green card. In case one subject reported a green and the other a pink card, the subject with the report of the green card won the contest. In case of a tie (same color), a random draw decided who was the winner. The winner prize was 80 token and the loser prize 40 token, corresponding to 8 euro and 4 euro respectively. As in the online experiment, in the *BL_lab* treatment, both, the FM and the SM could lie. In contrast, in the *SL_lab* treatment,

[6] Originally, we collected 2435 observations. We drop 35 observations because they did not finish the survey or the time spent on the survey was zero indicating that some technical device filled out the survey.

[7] The translated instructions can be found in the appendix.

the envelopes were shown on the computer screen for the FM. After selecting one out of ten, the color was revealed and the outcome entered automatically. As attrition is not a problem in the lab, the FM's and the SM's participated in one session. We implemented a between subject design and executed 13 sessions (seven sessions for *SL_lab* and six sessions for *BL_lab*) with 310 subjects in total. All subjects were enrolled at a large European University. The average payment was 7.46 euro and the experiment took approximately 45 minutes. We used hroot for recruiting the subjects (Bock et al., 2014) and ztree (Fischbacher, 2007) as the experimental software.

4 Results

4.1 Main Experiment

We start by analyzing the FM behavior. As the reports of the FM are simulated by the experimental software in the *SL* treatment, roughly 33% of them have a high report. By contrast, 63.75% and 50.62% report a high outcome in *BL* and *OL*, respectively. In both treatments, the fraction of high reports differs significantly from the expected fraction of 1/3 under truth-telling ($p = 0.000$; binomial probability test). High reports occur significantly more often in *OL* than in *BL,* i.e., when the SM cannot detect the lie ($p = 0.000$; Fisher's exact (FE)).[8] This extends findings that subjects lie less when the experimenter can detect lies (see the literature summarized in Abeler et al. (Forthcoming)) to situations where lying can be observed (or stochastically inferred) by other subjects. As a considerable but not too high fraction of FM lie in *BL* and *OL*, we have sufficiently many observations for low as well as for high reports in every treatment to investigate the corresponding responses of the SM.

For each treatment, Figure 1 depicts the fraction of high reports of SM, depending on the FM's report.[9] This fraction significantly exceeds 1/3 for all FM reports and all treatments ($p = 0.000$; binomial probability test).[10] In a next step, we analyze the SM reaction to FM reports. In line with Proposition 1, there is no significant difference in the SM reaction to high and low reports by the FM ($p = 0.911$) in the *SL* treatment:[11]

Result 1 *In the SL treatment, the fraction of high SM reports does not significantly differ*

[8]We report results from two-sided Fisher's exact tests if not stated otherwise.

[9]In the *OL* treatment, we split the dataset by the FM's actual outcome.

[10]Table 1 summarizes the exact fractions of high reports as well as the number of observations by treatment and report of the FM.

[11]Strictly speaking, we cannot test Proposition 1 which states that the FM report has no impact on SM behavior. Still, the fact that the fraction of high reports by the SM is quite similar after low and high reports of the FM supports the theoretical prediction.

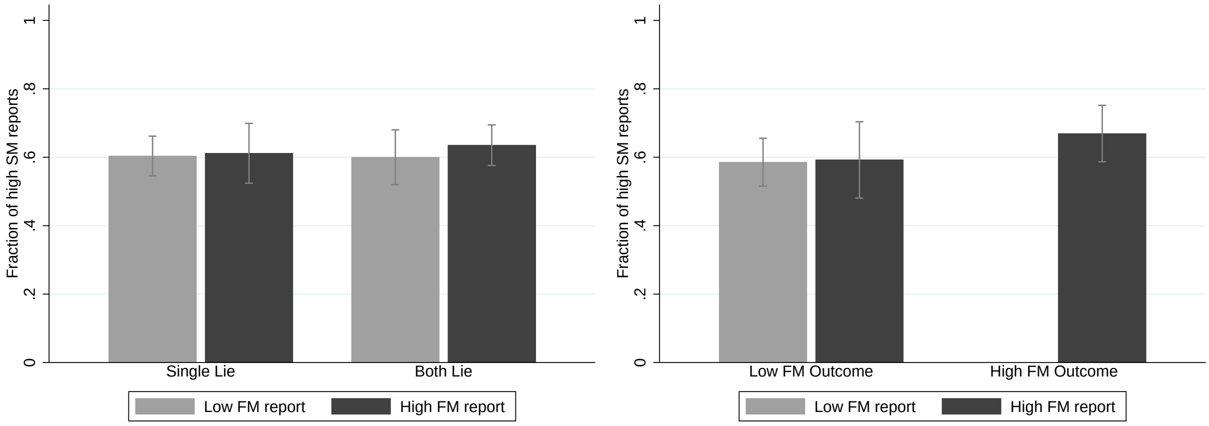

Figure 1: *The left panel shows the fraction of high SM reports for high and low FM reports in the SL and BL treatments. The right panel shows the fraction of high SM reports by report and actual outcome of the FM in the OL treatment.*

for high and low FM reports ($p = 0.911$*).*

In *BL*, a low report of the FM signals a high degree of altruism. Proposition 2 hence predicts that reciprocal SMs lie less often after low than after high reports. Surprisingly, however, this is not observed in the data:

Result 2 *In the BL treatment, the fraction of high SM reports does not significantly differ for high and low FM reports (*$p = 0.520$*).*

So far, we have compared the SMs reaction to high and low reports within treatments (SL and BL, respectively). Our next results compare the two treatments. As a consequence of up-dating on altruism types, reciprocal behavior predicts that the SM lies less (more) under *BL* compared to *SL* after observing low (high) reports (cf. Proposition 3). However, we find no evidence for either positive or negative reciprocity:

Result 3 *Both for high and low FM reports, there are no significant differences in SM reports between SL and BL treatments (*$p = 1.0$ *for low FM reports,* $p = 0.732$ *for high FM reports).*

In particular the fact that there is no difference in the response to low reports between the two treatments is striking: In the *BL* treatment, the SM knows with certainty that FMs reporting a low outcome refrained from upwards lying, but this clear signal of altruism does not reduce the own willingness to lie at all. Recall that, in the *OL* treatment, the SM knows with certainty whether the FM has lied or not. In line with the theoretical prediction that no FM will lie downwards, false low FM reports occurred only five times (4%). In case of low outcomes, 28.3% of FMs lie upwards, which equips us with enough observations to test Proposition 4.

If the FM reports high, we predict that SMs lie more often when the FM's report is a lie. This, however, is contradicted by the data as there is not only no significant difference ($p = 0.292$), but even the direction is reversed: 59.2% of the SM report high after a lie while 66.9% report high after a true high FM report. Second, for true reports of the FM, SMs lie somewhat less often after low reports, but the difference is insignificant ($p = 0.158$).[12] Thus, SMs do not significantly reward FMs for telling the truth after low outcomes. Next, we compare the cases where FMs truthfully reported a low outcome to those where they lied by reporting a high outcome. However, we again find no difference at all. We view the results for the OL treatment as the most striking findings of our experiment: While SMs can only stochastically update their beliefs on the FMs type in the *BL* treatment, in *OL* they know with certainty whether FMs have lied or not. But even with observability, there is no indication of reciprocity:

Result 4 *In the OL treatment, the SM reports do not differ significantly for different combinations of true outcome and report of the FM.*

In a last step, we compare the SM reports of the *SL* and the *OL* treatments after observing a true low FM report. Again, we find no treatment effect ($p = 0.703$). Given a true low report, the propensity to lie of SM does not depend on whether the FM has had the opportunity to lie or not. Thus, prediction 5 is rejected.

Result 5 *In case of a true low report, the SM reports do not differ between SL and OL treatments (*$p = 0.703$*).*

Summing up, every prediction that postulates an effect of reciprocity on the behavior of the SM is rejected by our data such that we do not find a significant impact of reciprocity on lying behavior in our experiment. In the following, we will demonstrate the robustness of this result by (i) implementing econometric analyses using regressions controlling for possible influence factors to explain SM behavior, (ii) conducting a power analysis, and (iii) analyzing the data from the additional lab experiment.

[12] Note, however, that this effect would become weakly significant under a one-sided test.

4.2 Robustness

4.2.1 Econometric Analysis

As lying is likely to be influenced by the own attitudes regarding, e.g., fairness and reciprocity as well as by the personal background, we now use regressions for testing the robustness of our findings for the SM behavior. The main explanatory variable is the share of high FM reports, or the treatment dummy when estimating between treatment comparisons. Where appropriate, we first add the belief about the percentage of upwards lying FM or SM, respectively, as an independent variable. Then, we add our proxies for fairness, sincerity, and greed avoidance. Finally, we combine the beliefs with measures for reciprocity. In addition, in all specifications but the first, we add the usual controls for demographic factors and time spend on MTurk. We estimate probit models with robust standard errors clustered at session level and report the results in Tables 2 to 8.

The econometric analysis supports the findings of the non-parametric tests as the coefficient for the FM report is insignificant throughout see Tables 2 to 3 and 5 to 7 for the within treatment comparisons). Also, when pooling the data from *SL* and *BL*, the *BL* treatment dummy remains insignificant for low as well as for high FM reports in all specifications (Table 4).

In the *OL* treatment, we also find no evidence that SM condition their report on whether a high FM report was a lie or not (Table 5), or whether the FM had an incentive to lie in case of a truthful report (Table 7). Notably, when we consider only observations where the actual outcome was low in the OL treatment (Table 6), the coefficient for the FM report is significant in specification 3 (controlling for fairness, sincerity, greed avoidance) and specification 4 (controlling for reciprocity). This provides some weak support that SM are more inclined to lie after a false high than after a truthful low report. However, effect sizes are rather small and our measures for reciprocity do not indicate a significant impact. In Table 8, we compare the treatment effect between *SL* and *OL* for true low and high FM reports. Again, the FM report is not significant.

Next, we take a closer look at the controls for the beliefs and the measures for social preferences elicited after the main part of the experiment. Recall that SM state a belief about the lying propensity of FM's only in the *BL* treatment.[13] Table 3 shows that the likelihood that SM lie indeed increases in the lying probability assigned to FM in the BL treatment. In the two other treatments, SM instead announce their belief about the lying propensity of their fellow SM. In all specifications we observe that a higher lying propensity

[13]In the *SL* treatment, FM coud not lie. In the *OL* tretament, lying was observable to SM.

attributed to other SM positively correlates with the own likelihood to state a high report. These findings, however, need to be interpreted with caution as the elicitation was not incentivized, and because subjects might state rather high beliefs to sugarcoat their own lie.

Next, we investigate our controls for fairness, sincerity, and greed avoidance. All constructs encompass two to three items measured on a 5-point likert scale. For each subject, we compute the standardized mean response to all items for each construct, so that higher scores reflect higher levels of the respective constructs. Overall, we find no systematic effects on the lying behavior. The same holds for our measures of positive and negative reciprocity and for the belief about reciprocity. Those measures are also based on survey items, now with a 7-point likert scale, and we compute the standardized mean response to all items for each subject. Summing up, the econometric analysis reinforces our finding that SM neither respond to the FM report nor to the different information on the FM's altruism provided by these reports in the three treatments.

4.2.2 Power Analysis

The combination of non-parametric tests and regression analysis supports the robustness of our findings, but one might still be concerned about the statistical power of the analyses. We, therefore, conducted ex-post power calculations to estimate the minimal detectable effect sizes (MDE) for each prediction. When adjusting $\alpha = 0.05$ and $1 - \beta = 0.8$ and using a two proportions test we are able to detect differences in the shares of high reports between 12.5 and 17.8 percentage points. For a detailed overview about MDE for each prediction see Table 9.

4.2.3 Lab Experiment

To further strengthen the robustness of our results, we complement the large MTurk dataset with a lab study including 310 subjects. We executed the *SL_lab* and the *BL_lab* treatments with a similar structure as the corresponding MTurk treatments. First, we analyze the FM behavior. In *SL_lab*, the FM could not lie as the report was filled in automatically. As expected, 30% of the reports are high.[14] In line with the MTurk data, 66.2% of the FM report high in *BL_lab* treatment, which is significantly different from 30% (i.e. the expectation without lying, binomial probability test $p = 0.000$). Again, we have enough variation in FM behavior to study the SM reaction to a FM report.

[14]Remember that subjects could select one out of ten envelopes. 3 contained green cards meaning high and 7 pink cards meaning low.

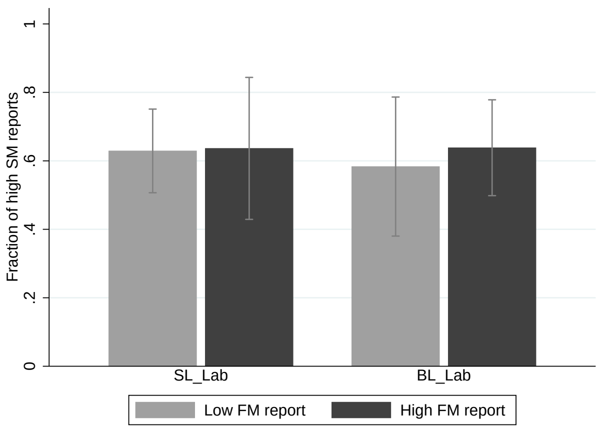

Figure 2: *Fraction of high SM reports for high and low FM reports in the SL_lab and BL_lab treatments.*

In both treatments, the SM could privately draw an envelope and state a report. As depicted in Figure 2, 63.1% of SM in the *SL_lab* treatment and 61.2% of SM in the *BL_lab* treatment state a high report, which differs significantly from 30% ($p = 0.000$ for both treatments). As on MTurk, the SM reports do not differ for high or low FM reports in the SL_lab treatment ($p = 1.0$). The pattern in the *BL_lab* treatment is quite similar ($p = 0.652$). The comparison of both treatments also reveals no differences in SM behavior ($p = 0.696$ for the comparison after low FM report and $p = 0.998$ after high FM report).

Summing up, the lab results are in line with those on MTurk as we do not find any evidence for reciprocity to affect lying behavior: SM subjects neither condition their behavior on the observed report of the FM in *BL_lab* nor on whether a FM report could serve as a signal on the degree of altruism.

5 Conclusion

We have developed a theoretical model to study the impact of reciprocity on lying behavior in a sequential two-player contest. As the marginal benefit of lying is the same when the FM reports high or low, the lying frequency of SM whose utility depends only on money and on internal lying costs is independent of the FM report. By contrast, our model of reciprocity predicts that SM are more inclined to lie if a high FM report signals a lie, and thereby lower altruism. We consider three settings: In *SL*, the FM cannot lie, so that her report provides no update on her type. In *BL*, the FM can lie, which allows for a stochastic up-date after observing her report. Finally, the SM can observe the actual outcome of the FM's lottery in *OL*. These settings allow us to derive predictions that compare the lying frequency of SM among treatments, depending on the FM's report.

To test our predictions on the impact of reciprocity, we have performed a large scale online experiment, amended by a controlled Lab experiment as a robustness check. Surprisingly, we find no evidence at all that reciprocity affects lying behavior. In all treatments, the SM lying frequency does not statistically differ for high and for low reports of the FM. This is most striking in *OL* where SM do not lie more often when they know with certainty that a high FM report was a lie. In addition, the lying frequency of SM does not differ between treatments, neither after a low nor after a high FM report. All results are qualitatively the same for the online- and the lab experiment. We hence conclude that reciprocity plays no role in our sequential contest with the possibility to lie.

Our model is based on reciprocity, but our experimental results also shed light on other possible drivers of lying behavior. First, fairness considerations have been reported to affect lying behavior (Houser et al., 2012). Under procedural fairness concerns (Trautmann, 2009; Krawczyk, 2011), players compare their expected payoffs. In our game, both players have an equal chance of winning, and thus identical expected payoffs, only in case of identical reports. Procedural fairness concerns thus provide an incentive for the SM to mimic the FM's report in all scenarios. Such a behavioral pattern, however, is not observed in any of our treatments.[15]

A second possible driver of lying behavior is social conformity, where the concerns of violating norms decrease in the (anticipated) violation frequency by others (Charness and Dufwenberg, 2006; Gibson et al., 2013; Diekmann et al., 2015). This motive generates similar predictions as the reciprocity model: when observing a high report of the FM in *BL*, the SM should be more inclined to lie than after observing a high report in *SL*. Furthermore, SM should lie more often after a fake than after a truthful high report in *OL*. None of these predictions is supported by our experimental results.

As a third potential motive, recall that the literature has shown that lying behavior in non-strategic situations can well be explained by a combination of internal lying costs and reputational concerns towards the experimenter (Abeler et al., Forthcoming). In settings with multiple participants, it thus seems reasonable that reputation towards other subjects

[15]Outcome-based fairness concerns that arise from inequity aversion (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000) have no bite in our setup. First, the payoff structure is identical between the scenarios. Second, the decrease in inequity costs associated with a lie instead of a truthful low report does not depend on the FM's report.

should also play a role, and that these reputational concerns should be larger towards subjects perceived as honest. Again, however, such a model would lead to predictions similar to our model based on reciprocity, and is hence also not supported by the data.

Summing up, in line with Benistant and Villeval (2017) who find that the lying behavior in contests is neither affected by group identity nor by whether lying increases the own or decreases the opponent's final score, our results suggest that lying behavior in contests is driven by the desire to win and by internal lying costs but neither by reciprocity, fairness considerations and social conformity nor by reputational concerns vis-à-vis honest contestants. A straightforward question for future research is if this finding is limited to situations such as contests, where the subjects' interests are strictly opposing, of if it extends to settings where their interests are aligned, e.g., where teams of two compete against other teams. Experiments on lying behavior find that subjects are more inclined to lie in groups (Muehlheusser et al., 2015; Kocher et al., 2017), which suggests that the own willingness to lie may, in such a setting, still be higher after observing a lie by the own partner.

Acknowledgments

We would like to thank Florian Kerzenmacher for his excellent research assistance and Eric Ritte for programming the experimental software.

6 Appendix

Proof of Proposition 1. The proof proceeds in two steps. First, we will show that in any equilibrium, for any strategy and a given report of the FM, the SM follows a cutoff-strategy: if some altruism type $\hat{a}_2$ refrains from lying, then all types $a_2 > \hat{a}_2$ also refrain from lying. In a second step, we will show that the cutoff does not depend on the report of the FM.

Step1: If, for a given strategy and report of the FM, $\exists \hat{a}_2 \in [-1, 1)$ *such that* $v_2(h, r_1) \leq v_2(l, r_1)$*, then* $v_2(h, r_1) \leq v_2(l, r_1)$ $\forall a_2 > \hat{a}_2$.

Suppose otherwise, i.e., suppose $\exists a_l, a_h$ with $a_l < a_h$ such that $v_2^{a_l}(h, r_1) \leq v_2^{a_l}(l, r_1)$ and $v_2^{a_h}(h, r_1) > v_2^{a_h}(l, r_1)$. First, let $r_1 = h$. Given the low outcome has realized for the SM, her utility from reporting low and high, respectively, amounts to:

$$v_2(h|h) = u_l + \frac{\Delta}{2} + \frac{a_2 + \lambda E_2[a_1|h]}{1+\lambda}\left(u_l + \frac{\Delta}{2}\right) - c$$
$$v_2(l|h) = u_l + \frac{a_2 + \lambda E_2[a_1|h]}{1+\lambda}(u_l + \Delta)$$

Comparing the utilities from reporting high and low shows that the SM, after observing a high report of the FM, lies if and only if

$$v_2(h|h) \geq v_2(l|h) \quad \Leftrightarrow \quad a_2 \leq (1+\lambda)(1 - \frac{2c}{\Delta}) - \lambda E_2[a_1|h] \equiv \tilde{a}_2^h.$$

It then needs to hold that $a_h \leq \tilde{a}_2^h < a_l$, a contradiction. Now suppose the FM has reported a low outcome, i.e., $r_1 = l$. The SM's utilities from reporting high and low, respectively, are:

$$v_2(h|l) = u_l + \Delta + \frac{a_2 + \lambda E_2[a_1|l]}{1+\lambda}u_l - c$$
$$v_2(l|l) = u_l + \frac{\Delta}{2} + \frac{a_2 + \lambda E_2[a_1|l]}{1+\lambda}\left(u_l + \frac{\Delta}{2}\right)$$

Thus, the SM prefers to lie if and only if

$$v_2(h|l) \geq v_2(l|l) \quad \Leftrightarrow \quad a_2 \leq (1+\lambda)(1 - \frac{2c}{\Delta}) - \lambda E_2[a_1|l] \equiv \tilde{a}_2^l$$

Again it needs to hold that $a_h \leq \tilde{a}_2^l < a_l$, a contradiction.

Step 2: $\tilde{a}_2^l = \tilde{a}_2^h$.

If the FM has no decision to take and, hence, can only report truthfully, then the SM can not infer any information about the FM's type from the observed report. In consequence, $E_2[a_1|l] = E_2[a_1|h] = E[a]$ and $\tilde{a}_2^l = \tilde{a}_2^h$. ■

Proof of Proposition 2. As *Step 1* from the proof of Proposition 1 is still valid, we know that the SM follows a strategy determined by a cutoff just like in the *SL* scenario. In a next step, we will show that in any equilibrium, the optimal strategy of the FM is also determined by a cutoff such that: if the FM refrains from lying for some type $\hat{a}_1$, then she will also refrain from lying for all $a_1 > \hat{a}_1$. In a final step we will show that, for any cutoff-strategy of the FM, the SM is weakly less inclined to lie after observing a low report of the SM as compared to a high report. We will show that the relation is strict as long as at least one report of the FM is informative about her type.

Step 1: If, for a given strategy of the SM, $\exists \hat{a}_1 \in [-1, 1)$ such that $v_1(h) \leq v_1(l)$, then $v_1(h) \leq v_1(l)\ \forall a_1 > \hat{a}_1$.

Suppose otherwise, i.e., suppose $\exists a_l, a_h$ with $a_l < a_h$ such that $v_1^{a_l}(h) \leq v_1^{a_l}(l)$ and $v_1^{a_h}(h) > v_1^{a_h}(l)$. Note that the FM anticipates the strategy of the SM. For a given report of herself, she correctly predicts the probability that the SM reports high when realizing the bad outcome, $q_2^{r_1}$, and the expected degree of altruism associated with a given report of the SM, $E_1[a_2|r_1, h]$. Then

$$\begin{aligned} v_1(h) = \ & [p_2 + (1-p_2)q_2^h]\left[u_l + \frac{\Delta}{2} + \frac{a_1 + \lambda E_1[a_2|h,h]}{1+\lambda}\left(u_l + \frac{\Delta}{2}\right)\right] \\ & + (1-p_2)(1-q_2^h)\left[u_l + \Delta + \frac{a_1 + \lambda E_1[a_2|h,l]}{1+\lambda}u_l\right] - c, \end{aligned}$$

and

$$\begin{aligned} v_1(l) = \ & [p_2 + (1-p_2)q_2^l]\left[u_l + \frac{a_1 + \lambda E_1[a_2|l,h]}{1+\lambda}(u_l + \Delta)\right] \\ & + (1-p_2)(1-q_2^l)\left[u_l + \frac{\Delta}{2} + \frac{a_1 + \lambda E_1[a_2|l,l]}{1+\lambda}\left(u_l + \frac{\Delta}{2}\right)\right]. \end{aligned}$$

Comparing the expected utilities shows that the FM lies if and only if

$$\begin{aligned} v_1(h) \geq v_1(l) \Leftrightarrow a_1 \leq (1+\lambda)&\left(1 - \frac{2c}{\Delta}\frac{1}{1-(1-p_2)(q_2^h - q_2^l)}\right) \\ & + \lambda\frac{[p_2 + (1-p_2)q_2^h]E_1[a_2|h,h] - [p_2 + (1-p_2)q_2^l]E_1[a_2|l,h]}{1-(1-p_2)(q_2^h - q_2^l)} \equiv \tilde{a}_1, \end{aligned}$$

where we made use of $[p_2 + (1-p_2)q_2^{r_1}]E_1[a_2|r_1, h] + (1-p_2)(1-q_2^{r_1})E_1[a_2|r_1, l] = E[a] = 0$. It then needs to hold that $a_h \leq \tilde{a}_1 < a_l$, a contradiction. It directly follows that the optimal report of the FM in case of a bad draw depends on her own altruism type in the following way:

$$r_1^l = \begin{cases} h \ \ \forall a_1 & \text{if } \ \tilde{a}_1 \geq 1 \\ \begin{cases} h & \text{if } \ a_1 \leq \tilde{a}_1 \\ l & \text{if } \ a_1 > \tilde{a}_1 \end{cases} & \text{if } \ \tilde{a}_1 \in [-1, 1) \\ l \ \ \forall a_1 & \text{if } \ \tilde{a}_1 < -1 \end{cases}$$

Step 2: $\tilde{a}_2^l(BL) \leq \tilde{a}_2^h(BL)\ \forall \tilde{a}_1$, and $\tilde{a}_2^l(BL) < \tilde{a}_2^h(BL)$ if $\tilde{a}_1 \geq -1$.

In a last step, we will analyze the optimal strategy of the SM for a given strategy of the FM. First, suppose $\tilde{a}_1 \in [-1, 1)$ such that $r_1^l = h \Leftrightarrow a_1 \leq \tilde{a}_1$. Conditional on the observed report of the FM, the SM updates the expectation about the type of the FM such that

$$E_2[a_1|l] = E[a|a > \tilde{a}_1]$$
$$E_2[a_1|h] = \frac{p_1 E[a] + (1-p_1)q_1^h E[a|a \leq \tilde{a}_1]}{p_1 + (1-p_1)q_1^h}.$$

As $E[a|a \leq \tilde{a}_1] < E[a] < E[a|a > \tilde{a}_1]$, it directly follows that $E_2[a_1|l] > E_2[a_1|h]$, and accordingly, $\tilde{a}_2^l < \tilde{a}_2^h$. Second, suppose $\tilde{a}_1 \geq 1$ such that $r_1^l = h \ \forall a_1$. Then $E_2[a_1|h] = E[a]$, whereas $E_2[a_1|l]$ is determined by an arbitrary off-equilibrium belief about the type(s) a_1 associated with a low report. Applying the D1-criterion proposed by Cho and Kreps (1987), which requires to associate that type with a deviation to the low report for which the deviation is most profitable, leads to $E_2[a_1|l] = 1$. Hence, it holds that $E_2[a_1|l] > E_2[a_1|h]$, and accordingly, $\tilde{a}_2^l < \tilde{a}_2^h$. Third, suppose $\tilde{a}_1 < -1$ such that $r_1^l = l \ \forall a_1$. Then $E_2[a_1|l] = E_2[a_1|h] = E[a]$ and $\tilde{a}_2^l = \tilde{a}_2^h$. ■

Proof of Proposition 4. As the SM observes e_1 and r_1, the updated expectation about the altruism type of the FM is given by $E_2[a_1|r_1, e_1]$ and a strategy for her specifies a report for every possible pair (r_1, e_1). Step 1 from the proof of Proposition 1 holds for every pair (r_1, e_1) such that, for a given strategy of the FM, the optimal strategy of the SM is determined by the cutoffs $\tilde{a}_2^{r_1,e_1}$, with $r_1 \in \{l, h\}$ and $e_1 \in \{l, h\}$. Regarding the FM, it still holds that a downward lie can never be optimal such that $E_2[a_1|h, h] = 0$ and $E_2[a_1|l, h] = 1$ (by D1). Given the bad outcome of the lottery has been realized for the FM, the insight from Step 1 from the proof of Proposition 2 still holds. Hence, the FM will either always lie, lie if and only if $a_1 \leq \tilde{a}_1$, or always report truthfully. We will now derive the ranking of the critical values $\tilde{a}_2^{r_1,e_1}$ for each possible equilibrium strategy of the FM. First, suppose $\tilde{a}_1 \geq 1$ such that the FM will always lie. It then holds that $E_2[a_1|h, l] = 0$ and $E_2[a_1|l, l] = 1$ (by D1). Accordingly, $\tilde{a}_2^{l,h} = \tilde{a}_2^{l,l} < \tilde{a}_2^{h,h} = \tilde{a}_2^{h,l}$ Second, suppose $\tilde{a}_1 \in [-1, 1)$: it then holds that $E_2[a_1|h, l] = E[a|a \leq \tilde{a}_1] < 0$ and $E_2[a_1|l, l] = E[a|a > \tilde{a}_1]$ with $0 < E[a|a > \tilde{a}_1] < 1$ such that $\tilde{a}_2^{l,h} < \tilde{a}_2^{l,l} < \tilde{a}_2^{h,h} < \tilde{a}_2^{h,l}$. Finally, suppose $\tilde{a}_1 < -1$. Then $E_2[a_1|h, l] = -1$ (by D1) and $E_2[a_1|l, l] = 0$ such that $\tilde{a}_2^{l,h} < \tilde{a}_2^{l,l} = \tilde{a}_2^{h,h} < \tilde{a}_2^{h,l}$. Overall, it is always true that $\tilde{a}_2^{l,h} \leq \tilde{a}_2^{l,l} \leq \tilde{a}_2^{h,h} \leq \tilde{a}_2^{h,l}$ with at least one comparison holding with strict inequality. ■

	SL	BL	OL_low	OL_high	SL_lab	BL_lab
Low FM report	60.4% (275)	60% (145)	58.5% (193)	80% (5)	62.9% (62)	58.3% (24)
High FM report	61.2% (121)	63.5% (255)	59.2% (76)	66.9% (127)	63.6% (22)	63.8% (47)

Table 1: Fraction of high SM reports by treatment and FM report. OL_low (OL_high) indicates that the true actual outcome of the FM was low (high). The number of observations of SM reports for each contingency is denoted in parentheses.

	(1)	(2)	(3)	(4)
Share of high FM reports	0.0206	0.0221	0.0428	0.0123
	(0.204)	(0.211)	(0.182)	(0.193)
Belief prop. lying SM		0.0886^{***}	0.0804^{**}	0.0929^{***}
		(0.0330)	(0.0345)	(0.0331)
Fairness			-0.178^{***}	
			(0.0536)	
Sincerity			−0.0233	
			(0.0438)	
Greed avoidance			0.0912	
			(0.0994)	
Pos. reciprocity				0.0226
				(0.0755)
Neg. reciprocity				0.172^{***}
				(0.0160)
Belief reciprocity				−0.109
				(0.0772)
Controls	*no*	*yes*	*yes*	*yes*
Constant	0.263	-0.332^{**}	0.0664	-0.483^{***}
	(0.190)	(0.156)	(0.448)	(0.173)
N	396	396	396	396
Pseudo R^2	0.0000	0.0398	0.0561	0.0511

Robust standard errors clustered on sessions in parentheses.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Controls: gender, age, race, education, time spent on MTurk

Table 2: Probit Regressions for *SL* treatment with fraction of high SM reports as the dependent variable.

	(1)	(2)	(3)	(4)
Share of high FM reports	0.0926	0.0601	0.0612	0.0613
	(0.134)	(0.125)	(0.163)	(0.132)
Belief prop. lying FM		0.0802***	0.0782***	0.0795***
		(0.0151)	(0.0164)	(0.0149)
Fairness			−0.0309	
			(0.113)	
Sincerity			−0.0353	
			(0.112)	
Greed avoidance			−0.0209	
			(0.0975)	
Pos. reciprocity				0.0246
				(0.0595)
Neg. reciprocity				0.0733
				(0.0722)
Belief reciprocity				0.106*
				(0.0588)
Controls	*no*	*yes*	*yes*	*yes*
Constant	0.253***	0.0221	0.288	−0.0329
	(0.0391)	(0.378)	(0.359)	(0.348)
N	400	400	400	400
Pseudo R^2	0.0009	0.0348	0.0369	0.0437

Robust standard errors clustered on sessions in parentheses.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$
Controls: gender, age, race, education, time spent on MTurk

Table 3: Probit Regressions for BL treatment with fraction of high SM reports as the dependent variable.

	(1)	(2)	(3)	(4)	(5)	(6)
BL treatment (low FM)	−0.00942	−0.0484	−0.0249			
	(0.180)	(0.172)	(0.170)			
BL treatment (high FM)				0.0625	0.0000159	0.0250
				(0.132)	(0.131)	(0.128)
Fairness		−0.127**		−0.110		
		(0.0615)		(0.0804)		
Sincerity		−0.106*		0.0243		
		(0.0591)		(0.0703)		
Greed avoidance		0.0580		0.0227		
		(0.0549)		(0.0778)		
Pos. reciprocity			−0.00293			0.0153
			(0.0775)			(0.0472)
Neg. reciprocity			0.148***			0.0823
			(0.0540)			(0.0642)
Belief reciprocity			−0.0427			0.0700
			(0.0614)			(0.0535)
Controls	*no*	*yes*	*yes*	*no*	*yes*	*yes*
Constant	0.263	1.054***	0.380**	0.283***	0.706*	0.415
	(0.176)	(0.163)	(0.177)	(0.0572)	(0.378)	(0.350)
N	420	420	420	376	376	376
Pseudo R^2	0.0000	0.0263	0.0197	0.0004	0.0314	0.0324

Robust standard errors clustered on sessions in parentheses.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Controls: gender, age, race, education, time spent on MTurk

Table 4: Probit Regressions for *SL* and *BL* treatment comparison with fraction of high SM reports as the dependent variable. Reaction to a given report. Specification 1-3 dataset restricted to low FM reports. Specification 4-6 dataset restricted to high FM reports.

	(1)	(2)	(3)	(4)
Share of high true FM results	0.205	0.209^{*}	0.205	0.183
(true high report)	(0.135)	(0.115)	(0.194)	(0.121)
Belief about SM reports		0.0313^{***}	0.0408^{***}	0.0284^{***}
		(0.00373)	(0.00613)	(0.00619)
Fairness			−0.0503	
			(0.0906)	
Sincerity			−0.143	
			(0.238)	
Greed avoidance			0.193^{***}	
			(0.0702)	
Pos reciprocity				0.0318
				(0.0955)
Neg. reciprocity				0.0107
				(0.0563)
Belief reciprocity				0.126
				(0.0850)
Controls	*no*	*yes*	*yes*	*yes*
Constant	0.233^{***}	0.384^{*}	0.419^{*}	0.377^{***}
	(0.0892)	(0.224)	(0.250)	(0.0422)
N	203	203	203	203
Pseudo R^2	0.0046	0.0133	0.0273	0.0211

Robust standard errors clustered on sessions in parentheses.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$
Controls: gender, age, race, education, time spent on MTurk

Table 5: Probit Regressions for OL treatment with fraction of high SM reports as the dependent variable. Dataset restricted high FM reports.

	(1)	(2)	(3)	(4)
Share of high FM reports	0.0170	0.0478	0.107**	0.0611***
(aka upward lie)	(0.0746)	(0.0364)	(0.0441)	(0.00164)
Belief prop. lying SM		0.106***	0.117***	0.0986***
		(0.0322)	(0.0297)	(0.0370)
Fairness			0.107**	
			(0.0510)	
Sincerity			−0.349***	
			(0.0181)	
Greed avoidance			0.0879	
			(0.0604)	
Pos. reciprocity				0.0269
				(0.0194)
Neg. reciprocity				0.0827
				(0.0696)
Belief reciprocity				0.0810
				(0.0569)
Controls	*no*	*yes*	*yes*	*yes*
Constant	0.216***	−0.109	0.449***	−0.203
	(0.0146)	(0.0863)	(0.121)	(0.195)
N	269	269	269	269
Pseudo R^2	0.0000	0.0435	0.0788	0.0499

Robust standard errors clustered on sessions in parentheses.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$
Controls: gender, age, race, education, time spent on MTurk

Table 6: Probit Regressions for OL treatment with fraction of high SM reports as the dependent variable. Reaction to true low report vs. upward lying. Dataset restricted to FM with a true low result of the die roll.

	(1)	(2)	(3)	(4)
Share of true high FM reports	0.222	0.236	0.228	0.229
compared to true low reports	(0.210)	(0.197)	(0.231)	(0.181)
Belief prop. lying SM		0.0858**	0.0871***	0.0844***
		(0.0337)	(0.0244)	(0.0317)
Fairness			−0.0224	
			(0.0350)	
Sincerity			−0.149**	
			(0.0585)	
Greed avoidance			0.0690	
			(0.164)	
Pos. reciprocity				0.0275
				(0.0355)
Neg. reciprocity				0.0519
				(0.0340)
Belief reciprocity				0.0115
				(0.0626)
Controls	*no*	*yes*	*yes*	*yes*
Constant	0.216***	0.143	0.495***	0.0987
	(0.0146)	(0.527)	(0.0327)	(0.529)
N	320	320	320	320
Pseudo R^2	0.0054	0.0336	0.0427	0.0350

Robust standard errors clustered on sessions in parentheses.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Controls: gender, age, race, education, time spent on MTurk

Table 7: Probit Regressions for OL treatment with fraction of high SM reports as the dependent variable. Reaction to true reports. Dataset restricted to true FM reports.

	(1)	(2)	(3)	(4)	(5)	(6)
OL treatment (true low FM)	−0.0468	−0.0252	−0.0214			
	(0.181)	(0.184)	(0.188)			
OL treatment (true high FM)				0.155	0.144	0.145
				(0.183)	(0.206)	(0.194)
Fairness		−0.0759			-0.241^{***}	
		(0.0574)			(0.0559)	
Sincerity		-0.172^{***}			0.115	
		(0.0647)			(0.135)	
Greed avoidance		0.0502			0.0939	
		(0.0664)			(0.105)	
Pos. reciprocity			−0.0113			0.0673
			(0.0782)			(0.0721)
Neg. reciprocity			0.118^{***}			0.0772
			(0.0440)			(0.0791)
Belief reciprocity			−0.0368			−0.0391
			(0.0579)			(0.0855)
Controls	*no*	*yes*	*yes*	*no*	*yes*	*yes*
Constant	0.263	1.201^{***}	0.447^{***}	0.283^{***}	0.553^{**}	0.546
	(0.180)	(0.174)	(0.160)	(0.0586)	(0.239)	(0.378)
N	468	468	468	248	248	248
Pseudo R^2	0.0002	0.0334	0.0217	0.0028	0.0355	0.0168

Robust standard errors clustered on sessions in parentheses.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Controls: gender, age, race, education, time spent on MTurk

Table 8: Probit Regressions for *SL* and OL treatment comparison with fraction of high SM reports as the dependent variable. Specification 1-3 dataset restricted to true low FM reports. Specification 4-6 dataset restricted to true high FM reports.

comparison	base	delta
$\tilde{a}_2^l(SL) = \tilde{a}_2^h(SL)$	0.6036	0.1436
$\tilde{a}_2^l(BL) < \tilde{a}_2^h(BL)$	0.6	0.1359
$\tilde{a}_2^l(BL) < \tilde{a}_2^l(SL)$	0.6	0.1359
$\tilde{a}_2^h(SL) < \tilde{a}_2^h(BL)$	0.6116	0.1420
$\tilde{a}_2^{l,l}(OL) < \tilde{a}_2^{h,h}(OL)$	0.5855	0.1515
$\tilde{a}_2^{l,l}(OL) < \tilde{a}_2^{h,l}(OL)$	0.5855	0.1784
$\tilde{a}_2^{h,h}(OL) < \tilde{a}_2^{h,l}(OL)$	0.6693	0.1751
$\tilde{a}_2^{l,l}(OL) < \tilde{a}_2^l(SL)$	0.5849	0.1250
$\tilde{a}_2^{h,h}(OL) = \tilde{a}_2^h(SL)$	0.6116	0.1632

Table 9: Overview MDE for two proportions test.

Instructions

HIT on Mturk

Scientific study, survey ($1 + possible bonus, ˜10 min)

This is a scientific study conducted by researchers from Frankfurt School of Finance & Management. It consists of a short survey on personal attitudes for which you will be paid $1 and a simple game in which you might earn extra money.

Your Worker ID will be retrieved automatically when you click to link to start the project. It will only be used for assigning the payment to the right account and to control that you have not participated in this HIT before. On the last page of the survey, you will receive a personalized completion code. Please copy and paste this completion code in the box below so that we can verify that you have completed the survey.

When you have finished with the project, come back here and submit the HIT. We will approve payments within two days.

Please click on the link below in order to start.

Make sure to leave this window open as you complete the project.

When you are finished, you will return to this page to paste the code in the box below.

Provide the participation code here: _ _ _ _ _ _ _ _ _

Survey on unipark

Instructions for First Mover

Before the survey starts, you will play a simple game where you can earn additional money that will be added to the $1 we pay you for the survey. In this game you will be matched with another worker who will take part in the study within the next two days. This worker will not learn your ID.

SL and OL treatment

In this game the computer will simulate the roll of a six-sided die.

BL treatment

In this game you will have to roll a six-sided die. Note that you will roll the die in private, so that the outcome cannot be observed by anyone else.

After observing the outcome of the die roll, we will ask you to report the result of your die roll.

All

An outcome of...

means the result is ”Low”

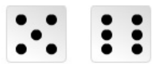

means the result is ”High”

SL treatment

After the simulation of the die roll, the result will be filled in automatically.

The worker you are matched with will be informed about your result. After having observed your result, this worker will roll a die in private and report either “High” or “Low.”

OL treatment

The worker you are matched with will be informed about the actual outcome of the simulated die roll as well as about your reported result. After receiving this information, this worker will roll a die in private and report either “High” or “Low.” Note that the additional payment is based solely on the reported results and not on the outcomes of the simulated and actual die roll.

BL treatment

The worker you are matched with will be informed about your reported result. After having observed your reported result, this worker will also roll a die in private and report either “High” or “Low.”

All

The additional payment is calculated as follows

case	report	payment
1	You: High Other: Low	You: $1.10 Other: $0.10
2	You: Low Other: High	You: $0.10 Other: $1.10
3	You and the other worker report both ”High” or both ”Low.”	Each of you has a 50% chance of getting the $1.10. getting the $1.10. A random draw decides if you or the other worker get the $1.10. you or the other worker get the $1.10.

The additional payment will be sent to you as a bonus within the next two days. The bonus payment will include a message stating the reports of both workers and the resulting payment.

Do you wish to participate?

If you do not wish to participate, please close this window now.

If you do wish to participate, please press ”Continue.”

Next Page

SL treatment

[Here result of the die roll is shown]

Your result is ”Low.”

OL treatment

The simulation of the die roll is:

[Here result of the die roll is shown]

An outcome of...

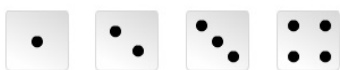

means the result is ”Low”

means the result is "High"

Please report either "High" or "Low" by checking one of the two boxes below.

_ _ Low   _ _High

BL treatment

Getting started:

Please get a die or use a virtual one.

You can virtually roll a die for instance on https://www.freeonlinedice.com/#dice.

Now please roll the die exactly once.

An outcome of...

means the result is "Low"

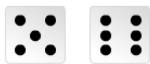

means the result is "High"

Please report the result by checking one of the two boxes below.

_ _ Low   _ _High

Instructions for Second Mover

Before the survey starts, you will play a simple game where you can earn additional money that will be added to the $1 we pay you for the survey. In this game you will be matched with another worker who has already taken part in the study. This worker will not learn your ID.

In this game you will have to roll a six-sided die. Note that you will roll the die in private, so that the outcome cannot be observed by anyone else.

After observing the outcome of the die roll, we will ask you to report the result of your die roll.

An outcome of...

means the result is ”Low”

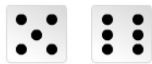

means the result is ”High”

SL treatment

The worker you are matched with did not roll a die. Instead, the roll of a six-sided die has been simulated by a computer and ”Low” was automatically filled in for the outcome 1 to 4 and ”High” for the outcomes 5 or 6.

Before you are asked to roll the die, you will be informed about the result of the worker you are matched with.

OL treatment

The worker you are matched with did not roll a die in private. Instead, the roll of a six-sided die has been simulated by a computer and the outcome has been shown to the worker. After observing the outcome of the die roll, the worker reported either ”High” or ”Low.”

Before you are asked to roll the die, you will be informed about the actual outcome of the simulated die roll as well as about the reported result of the worker you are matched with. Note that the additional payment is based solely on the reported results and not on the outcomes of the simulated and actual die roll.

BL treatment

The worker you are matched with has already rolled a die and reported either ”High” or ”Low.” Note that this player also rolled the die privately, so that the actual outcome of the die roll has not been observed by anyone else. Before you are asked to roll the die, you will be informed about the reported result of the worker you are matched with.

All

The additional payment is calculated as follows

case	report	payment
1	You: High Other: Low	You: $1.10 Other: $0.10
2	You: Low Other: High	You: $0.10 Other: $1.10
3	You and the other worker report both "High" or both "Low."	Each of you has a 50% chance of getting the $1.10. getting the $1.10. A random draw decides if you or the other worker get the $1.10. you or the other worker get the $1.10.

The additional payment will be sent to you as a bonus within the next two days. The bonus payment will include a message stating the reports of both workers and the resulting payment.

Do you wish to participate?

If you do not wish to participate, please close this window now.

If you do wish to participate, please press "Continue."

Next page

Getting started:

Please get a die or use a virtual one.

You can virtually roll a die for instance on https://www.freeonlinedice.com/#dice.

SL and BL treatment

The worker you are matched with has reported the result "Low."

Now please roll the die exactly once

OL treatment

The simulation of the die roll for the worker you are matched with was

[Here result of the die roll is shown.]

The worker you are matched with has reported the result "Low."

Now please roll the die exactly once.

All

An outcome of...

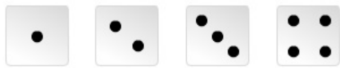

means the result is "Low"

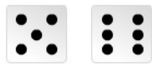

means the result is "High"

Please report the result by checking one of the two boxes below.

_ _ Low _ _High

Instructions for SL_lab and BL_lab

Most of these instructions are identical for the treatments "Single lie" and "Both lie." Parts that refer to only one of the treatments are indicated by SL_lab and BL_lab, respectively.

Welcome to this experiment!

You are participating in an economic experiment on decision making. All decisions are made anonymously, i.e., other participants will not learn the identity of a participant who has made a certain decision. The payment is anonymous as well, i.e., a participant will not learn the amount of payment received by other participants.
All instructions will be displayed on screen. Please read the instructions carefully. If you have trouble understanding the instructions, please take a second look at them. If you still have questions, please give us a hand signal.

General information

- The experiment consists of multiple parts and a short questionnaire.
- You will receive instructions for each part before the respective part starts. Please click on "OK" only after you have read the instructions and have no further questions.
- At the end of the experiment, you will receive an overview of your results.
- Each part can influence your payment. Therefore, please consider each of your decisions carefully.
- We use the unit "tokens" during the experiment. It will be converted into euro at the end. The exchange rate is one euro for 10 token.
- After the experiment, please remain seated in your cubicle until we call you for the payment.

Instructions part 1

Procedure

In this part of the experiment, random groups of two participants will be formed. In each group, there are two roles, participant A and participant B. You will be informed whether the role of participant A or participant B was assigned to you before this part starts.

Each participant chooses one out of ten envelopes. Each envelope contains either a green or a pink ticket. Overall, there are three envelopes with green and seven with pink tickets. The number of pink and green tickets is identical for each participant.
Subsequently, you and the participant assigned to you will each open one envelope, and then declare the color of the ticket in the envelope. Please open the envelope only if we ask you to do so. Your payment depends on the declared color of your ticket and the declared color of the ticket of the other participant.
We will now explain the exact process

For treatment "SL_lab" only:

- Participant A sees ten envelopes on the screen, of which he selects one.
- Participant A opens the selected envelope by clicking on it, so that he can see the colour of the ticket contained.
- After opening the envelope, the colour will be entered automatically in the designated response field.
- Participant B will be informed about the colour drawn by participant A.
- On the table of participant B, there is a large envelope containing ten small envelopes.
- After participant B was informed about the color drawn by participant A, participant B will be asked to open one envelope. He takes the large envelope and opens it. He then chooses one of the small envelopes. Participant B opens the small envelope to see the color of the ticket contained.
- In contrast to participant A, participant B enters the color of the ticket in the designated response field self-reliant, after opening the small envelope.
- After participant B entered his answer, the participants will be informed about the entered colors of both participants and their results.

For treatment "BL_lab" only:

- On the table of participant A, there is a large envelope containing ten small envelopes.
- Participant A will be asked to open one envelope. He takes the large envelope and opens it. He then chooses one of the small envelopes. Participant A opens the small envelope to see the color of the ticket contained.

- Participant A enters the color of the ticket in the designated response field self-reliant, after opening the small envelope.
- After entering the color, participant B will be informed about the color entered by participant A.
- On the table of participant B, there is a large envelope containing ten small envelopes.
- After participant B was informed about the color drawn by participant A, participant B will be asked to open one envelope. He takes the large envelope and opens it. He then chooses one of the small envelopes. Participant B opens the small envelope to see the color of the ticket contained.
- Participant B enters the color of the ticket in the designated response field self-reliant, after opening the small envelope.
- After participant B entered his answer, the participants will be informed about the entered colors of both participants and their results.

For both:

Calculation of the results

The result depends on the colors entered by both participants.
If the entered color of one participant is green and of the color entered by the other participant is pink, then the participant with the green color receives 80 tokens and the participant with the pink color receives 40 tokens. If both participants entered the same color, it will be randomly chosen who will receive the 80 and who the 40 token.
Now follows a short quiz, which contains questions to the instructions. If there are any questions or problems of understanding, please signal this to us by hand signal.

References

Abeler, J., D. Nosenzo, and C. Raymond (Forthcoming): "Preferences for truth-telling," *Econometrica*.

Arechar, A. A., S. Gächter, and L. Molleman (2018): "Conducting interactive experiments online," *Experimental Economics*, 21, 99–131.

Ashton, M. C. and K. Lee (2009): "The HEXACO–60: A Short Measure of the Major Dimensions of Personality," *Journal of Personality Assessment*, 91, 340–345.

Belot, M. and M. Schröder (2013): "Sloppy work, lies and theft: A novel experimental design to study counterproductive behaviour," *Journal of Economic Behavior & Organization*, 93, 233 – 238.

Benistant, J. and M. C. Villeval (2017): "Unethical Behavior and Group Identity in Contests," *mimeo.*

Berinsky, A. J., G. A. Huber, and G. S. Lenz (2012): "Evaluating online labor markets for experimental research: Amazon. com's Mechanical Turk," *Political Analysis*, 20, 351–368.

Bock, O., I. Baetge, and A. Nicklisch (2014): "hroot: Hamburg Registration and Organization Online Tool," *European Economic Review*, 71, 117 – 120.

Bolton, G. E. and A. Ockenfels (2000): "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, 90, 166–193.

Carpenter, J., P. H. Matthews, and J. Schirm (2010): "Tournaments and Office Politics: Evidence from a Real Effort Experiment," *American Economic Review*, 100, 504–17.

Charness, G. and M. Dufwenberg (2006): "Promises and Partnership," *Econometrica*, 74, 1579–1601.

Charness, G., D. Masclet, and M. C. Villeval (2013): "The Dark Side of Competition for Status," *Management Science*, 60, 38–55.

Cho, I.-K. and D. M. Kreps (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102, 179–222.

Conrads, J., B. Irlenbusch, R. M. Rilke, C. Schielke, and G. Walkowitz (2014): "Honesty in tournaments," *Economics Letters*, 123, 90–93.

Crump, M. J. C., J. V. McDonnell, and T. M. Gureckis (2013): "Evaluating Amazon's Mechanical Turk as a Tool for Experimental Behavioral Research," *PLOS ONE*, 8, 1–18.

DellaVigna, S. and D. Pope (2018): "What Motivates Effort? Evidence and Expert Forecasts," *The Review of Economic Studies*, 85, 1029–1069.

Diekmann, A., W. Przepiorka, and H. Rauhut (2015): "Lifting the veil of ignorance: An experiment on the contagiousness of norm violations," *Rationality and Society*, 27, 309–333.

Doğan, G. and R. A. Roggema (2016): "Lying to be Fair," *mimeo.*

Dohmen, T., A. Falk, D. Huffman, and U. Sunde (2009): "Homo Reciprocans: Survey Evidence on Behavioural Outcomes," *The Economic Journal*, 119, 592–612.

Dufwenberg, M. and G. Kirchsteiger (2004): "A theory of sequential reciprocity," *Games and Economic Behavior*, 47, 268 – 298.

Falk, A. and U. Fischbacher (2006): "A theory of reciprocity," *Games and Economic Behavior*, 54, 293 – 315.

Faravelli, M., L. Friesen, and L. Gangadharan (2015): "Selection, tournaments, and dishonesty," *Journal of Economic Behavior & Organization*, 110, 160 – 175.

Fehr, E. and S. Gächter (1998): "Reciprocity and economics: The economic implications of Homo Reciprocans," *European Economic Review*, 42, 845 – 859.

Fehr, E. and K. M. Schmidt (1999): "A Theory of Fairness, Competition, and Cooperation," *The Quarterly Journal of Economics*, 114, 817–868.

Fischbacher, U. (2007): "z-Tree: Zurich toolbox for ready-made 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.

Gächter, S., L. Gerhards, and D. Nosenzo (2017): "The importance of peers for compliance with norms of fair sharing," *European Economic Review*, 97, 72 – 86.

Gibson, R., C. Tanner, and A. F. Wagner (2013): "Preferences for truthfulness: Heterogeneity among and within individuals," *American Economic Review*, 103, 532–48.

Gneezy, U. (2005): "Deception: The role of consequences," *American Economic Review*, 95, 384–394.

Gneezy, U., A. Kajackaite, and J. Sobel (2018): "Lying Aversion and the Size of the Lie," *American Economic Review*, 108, 419–53.

Harbring, C. and B. Irlenbusch (2008): "How Many Winners are Good to Have?: On Tournaments with Sabotage," *Journal of Economic Behavior & Organization*, 65, 682–702.

Horton, J. J., D. G. Rand, and R. J. Zeckhauser (2011): "The online laboratory: conducting experiments in a real labor market," *Experimental Economics*, 14, 399–425.

Houser, D., S. Vetter, and J. Winter (2012): "Fairness and cheating," *European Economic Review*, 56, 1645 – 1655.

Khalmetski, K. and D. Sliwka (2017): "Disguising Lies-Image Concerns and Partial Lying in Cheating Games," *mimeo.*

Kocher, M. G., S. Schudy, and L. Spantig (2017): "I lie? We lie! Why? Experimental evidence on a dishonesty shift in groups," *Management Science.*

Krawczyk, M. W. (2011): "A model of procedural and distributive fairness," *Theory and Decision*, 70, 111–128.

Levine, D. K. (1998): "Modeling altruism and spitefulness in experiments," *Review of economic dynamics*, 1, 593–622.

Mason, W. and S. Suri (2012): "Conducting behavioral research on Amazon's Mechanical Turk," *Behavior Research Methods*, 44.

Mazar, N., O. Amir, and D. Ariely (2008): "The dishonesty of honest people: A theory of self-concept maintenance," *Journal of marketing research*, 45, 633–644.

Muehlheusser, G., A. Roider, and N. Wallmeier (2015): "Gender differences in honesty: Groups versus individuals," *Economics Letters*, 128, 25–29.

Paolacci, G. and J. Chandler (2014): "Inside the Turk: Understanding Mechanical Turk as a Participant Pool," *Current Directions in Psychological Science*, 23, 184–188.

Perugini, M., M. Gallucci, F. Presaghi, and A. P. Ercolani (2003): "The personal norm of reciprocity," *European Journal of Personality*, 17, 251–283.

Rabin, M. (1993): "Incorporating Fairness into Game Theory and Economics," *The American Economic Review*, 83, 1281–1302.

Trautmann, S. T. (2009): "A tractable model of process fairness under risk," *Journal of Economic Psychology*, 30, 803 – 813.

Wenzhong, Z. and F. Limin (2012): "A Case Study of Siemens' Violation of Business Ethics in Argentine Based On Stakeholder Theory," *Global Journal of Management and Business Research*, 12, 59–65.