



# Do altruists lie less?☆

Rudolf Kerschbamer, Daniel Neururer\*, Alexander Gruber

Department of Economics, Univ. of Innsbruck, Universitätsstraße 15, Innsbruck 6020, Austria

## ARTICLE INFO

### Article history:

Received 19 April 2018  
Revised 17 October 2018  
Accepted 31 October 2018  
Available online 17 November 2018

### JEL classification:

C91  
D63  
D64

### Keywords:

Deception  
Lies  
Social preferences  
Distributional preferences  
Equality equivalence test

## ABSTRACT

Much is known about heterogeneity in social preferences and about heterogeneity in lying aversion – but little is known about the relation between the two at the individual level. Are the altruists simply upright persons who do not only care about the well-being of others but also about honesty? And are the selfish those who lie whenever lying maximizes their material payoff? This paper addresses those questions in experiments that first elicit subjects' social preferences and then let them make decisions in an environment where lying increases the own material payoff and has either consequences for the payoffs of others or no consequences for others. We find that altruists lie less when lying hurts another party but we do not find any evidence in support of the hypothesis that altruists are more (or less) averse to lying than others in environments where lying has no effects on the payoffs of others.

© 2018 The Authors. Published by Elsevier B.V.  
This is an open access article under the CC BY license.  
(<http://creativecommons.org/licenses/by/4.0/>)

## 1. Introduction

Communication is an essential feature of human interaction and in particular of economic activity. Whenever communication is involved, some agents might have material incentives to misrepresent their private information at the expense of other market participants. Examples include markets for experience goods, where sellers have incentives to claim that the quality of the good offered by them is higher than it actually is (Nelson, 1970); markets for label goods, where firms are tempted to assert that their products possess hidden attributes valued by consumers (Bonroy and Constantatos, 2014; Feddersen and Gilligan, 2001); and markets for credence goods, where experts often have incentives to provide a quality that does not fit the customer's needs (Darby and Karni, 1973; Dulleck and Kerschbamer, 2006; Dulleck et al., 2011).

Market and non-market forces often prevent agents from making false claims despite short-run material incentives for misreporting. For instance, sellers in markets for experience goods may resist the temptation to lie because they fear to lose business in the future (Huck et al., 2016); firms in markets for label goods might stay honest because they are afraid that mislabeled products are detected and punished with a high enough probability (Étilé and Teyssier, 2016); and experts in markets for credence goods might act in line with the interests of consumers to avoid a bad reputation (Dulleck et al., 2011; Grosskopf and Sarin, 2010; Mimra et al., 2016; Schneider, 2012).

☆ Financial support from the vice-rectorship of research of the University of Innsbruck through grant number 2013/1/VWL21 and from the Austrian Science Fund (FWF) through special research area grant SFB F63, as well as through grant numbers P22669, P26901 and P27912 is gratefully acknowledged.

\* Corresponding author.

E-mail address: [daniel.neururer@uibk.ac.at](mailto:daniel.neururer@uibk.ac.at) (D. Neururer).

Recent experimental evidence – presented by Gneezy (2005), Dreber and Johannesson (2008), Hurkens and Kartik (2009), Erat and Gneezy (2012), Fischbacher and Föllmi-Heusi (2013), Lopez-Perez and Spiegelman (2013) and Gibson et al. (2013), among others – suggests that some agents avoid making false claims even in environments where lying is difficult or even impossible to detect – or is not punished when detected. By contrast, others seem to cheat whenever cheating is in their material self-interest. A possible explanation for this finding is that it is the result of lying costs which differ across agents (see Erat and Gneezy, 2012; Gibson et al., 2013; Gneezy, 2005 and Gneezy et al., 2013, among others).

Agents differ not only in their willingness to tell a lie, they also differ in their propensity to take the material consequences for others into account when making economic decisions. Indeed, experimental work by Andreoni and Miller (2002), Engelmann and Strobel (2004), Fisman et al. (2007), Cox and Sadiraj (2012) and Kerschbamer (2015) documents large heterogeneity in the distributional – or 'social' – preferences of subjects, with some agents deciding as if they were only interested in their own material payoffs while others seem to care to different degrees for the payoffs of others or for fairness more generally.<sup>1</sup>

Although heterogeneity in lying aversion and heterogeneity in social preferences are by now well documented, relative little is known about the relationship between the two at the individual level. Do altruistic persons lie less than others – for instance, because they are simply more "moral" persons? And do selfish persons lie whenever lying maximizes their material payoff – simply because they do not care for moral values at all? This paper addresses those questions with the help of controlled laboratory experiments. Specifically, we present evidence from experiments that first elicit subjects' social preferences and then let them make decisions in an environment where lying increases the own material payoff and has either consequences for the payoffs of others or no consequences for the payoffs of others.

Eliciting social preferences is a thorny task. In our experiments we use a simple and intuitive approach – the Equality Equivalence Test (Kerschbamer, 2015). This test elicits benevolence in two domains of income allocations – the domain of advantageous inequality where the decision maker is ahead of another person, and the domain of disadvantageous inequality where the decision maker is behind. According to the revealed benevolence, neutrality or malevolence of the decision maker in the two domains, she or he is classified into a distributional preference type.

Different methods have been applied to identify lying aversion in the aggregate or on the individual level. A simple and elegant design has been introduced to the literature by Fischbacher and Föllmi-Heusi (2013): Subjects roll a dice in privacy, report the outcome and get paid based exclusively on their report. Since the experimenter knows the underlying distribution of outcomes, cheating behavior in the aggregate can readily be quantified with this design. The approach is less well suited for the identification of lying aversion at the individual level.

A second strand of literature, pioneered by Gneezy (2005) and slightly modified by Erat and Gneezy (2012), builds on the cheap talk game by Crawford and Sobel (1982) to identify lying behavior on the individual level. In this class of games, a sender learns the true state of the world and then sends a costless message about the state to the receiver. The receiver observes the message of the sender and then chooses an action. After this choice the game ends with payoffs for the two players that depend only on the state of the world and the choice of the receiver, but not on the message sent by the sender. Only the sender is aware about the monetary consequences for both players associated with each choice of the receiver (even after the end of the game). As a consequence, the receiver can not anticipate if the sender has an incentive to transmit the information about the state honestly or not, and the receiver will never learn if the sender said the truth or not.<sup>2</sup>

This paper investigates in experiments involving three treatments whether there is a relationship between the distributional preference type of a subject and his aversion against lying. In all treatments we first elicit subjects' social preferences via the Equality Equivalence Test. Depending on the treatment we then expose them either to a dice-rolling task à la Fischbacher and Föllmi-Heusi (2013) or to one of two versions of a deception game à la Erat and Gneezy (2012).

In the original dice-rolling task as introduced by Fischbacher and Föllmi-Heusi (2013) subjects are asked to roll a dice just once. By comparing the distribution of reports to the predicted distribution of outcomes the researcher receives information about the frequency and extend of lying in the aggregate. On the individual level, however, a reported outcome is only a noisy signal for whether the subject told the truth. To improve the information content of the signal we modify the Fischbacher and Föllmi-Heusi (2013) design slightly by asking the subjects to roll the dice ten times and to self-report the outcomes. The payoff of the subject then corresponds to the sum of self-reported outcomes. We term this the *DICE* treatment.

Our two versions of the deception game à la Erat and Gneezy (2012) have the properties that (i) the sender and the receiver face the same action spaces in both versions and (ii) the sender has the same material incentive to lie in both versions. The two versions – which correspond to the treatments *NEG* and *NEU* in our experiment – differ in the material consequences of lying for the receiver: In *NEG* a lie increases the material payoff of the sender and decreases the material payoff of the receiver and in *NEU* a lie only increases the material payoff of the sender and does not affect the material payoff of the receiver.<sup>3</sup>

<sup>1</sup> The terms distributional preferences and social preferences are used interchangeably in this paper for cases where an agent potentially cares not only for the own material payoff, but also for the consequences of her decisions for the payoffs of others.

<sup>2</sup> This game is often called the 'deception game' – by Gneezy (2005) and Angelova and Regner (2013), for instance.

<sup>3</sup> Erat and Gneezy (2012) present a taxonomy of lies based on the payoff consequences of the lie for the liar and the recipient of the lie: Selfish black lies involve acts that help the liar at the expense of the recipient, Pareto white lies benefit both parties, altruistic white lies harm the liar but benefit the

The main result of our experiments involving 356 subjects is that altruists lie less when lying hurts the other party (as in the *NEG* version of the deception game), but do not lie less than others in environments where lying has no effects on the payoffs of others (as in *DICE* and in the *NEU* version of the deception game).

The rest of the paper is organized as follows. [Section 2](#) reviews the related literature. [Section 3](#) introduces the experimental design and the procedure. The hypotheses are formulated in [Section 4](#). [Section 5](#) presents the experimental results and [Section 6](#) concludes. An appendix contains robustness checks, a detailed explanation about the identification of distributional preferences with the Equality Equivalence Test and the experimental instructions.

## 2. Literature review

In terms of research question and experimental design the paper closest to ours is [Maggian and Villeval \(2016\)](#). These authors investigate in an experiment involving three treatments the correlation between social preferences and lying aversion in children. Each of the three treatments has two stages. The choice in the first stage is used to identify the social preference type of the subject and the choice in the second stage is used to identify lying at the individual level. Specifically, in stage 1 subjects are asked to decide between two allocations, each involving an own material payoff and a material payoff for another subject. In all treatments, one of the two allocations involves an egalitarian distribution, the second allocation creates advantageous or disadvantageous inequality, depending on the treatment. In the Selfishness treatment, by choosing the asymmetric allocation instead of the symmetric one, the subject increases its own material payoff but decreases that of the second subject, while in the Altruism and the Efficiency treatment choosing the asymmetric allocation increases the other subject's payoff either at an own cost (in the Altruism treatment) or at no own cost (in the Efficiency treatment). Depending on her choice, the subject is then assigned one of two distributional preference types. For instance, in the Selfishness treatment, a subject choosing the asymmetric allocation instead of the symmetric one, is classified as selfish while a subject deciding for the symmetric allocation is classified as inequality averse. In all treatments the second stage involves the same two allocations as the first stage but now the computer randomly "proposes" one of the two allocations. Subjects are then asked to report the allocation proposed by the computer and the reported allocation is then implemented.<sup>4</sup> For the identification of lying aversion in the second stage the authors restrict their attention to subjects where the computer-proposed allocation did not match the allocation chosen by the subject in stage 1. If such a subject reports the allocation she or he had chosen in stage 1, the choice of the subject is classified as lying behavior. This means, that in the Selfishness treatment subjects classified as selfish can only tell a black lie (that is, a lie that harms the other side) while subjects classified as inequality averse can only tell a white lie (i.e., a lie that benefits the other party). The authors find that black lies by selfish children are more frequent than white lies by pro-social children.

We view our work as complementary to this interesting research. An important difference between studies is the social preference elicitation procedure employed: In [Maggian and Villeval \(2016\)](#) the elicitation task and the set of social preference types tested for vary across treatments while we use the same procedure and test for the same set of types in all our treatments. A second difference between our design and theirs is that whether lying has negative or neutral consequences for another subject depends only on the treatment in our design, while it depends also on the social preference type of the subject in their design.

Other papers related to ours are [Cappelen et al. \(2013\)](#) and [Sheremeta and Shields \(2013\)](#). [Cappelen et al. \(2013\)](#) conduct an experiment based on a sender-receiver game and examine how the decision to lie depends on factors like the content of the lie, the market context and personal characteristics of the sender. An important feature of their experiments is that lying is beneficial for both the sender and the receiver in all of their treatments. An other important feature (in relation to our work) is that all subjects are asked to make a standard dictator decision before being exposed to the sender-receiver game. With respect to this decision the authors find that subjects who give higher shares in the dictator game show also a higher degree of lying aversion in the sender-receiver game. This finding is remarkable because pro-social subjects have an additional incentive for lying compared to selfish subjects in their sender-receiver games where all lies are Pareto white lies.

[Sheremeta and Shields \(2013\)](#) show that other-regarding preferences and lying aversion are also correlated in the context of selfish black lies. In their experiments subjects first play both roles in a sender-receiver game where senders observe with equal probability one of two signals and where they have a material incentive to lie when they see one of the two signals. If the lie is believed by the receiver then the sender is better off and the receiver is worse off compared to the truthful report of the sender. After having made the two decisions in the sender-receiver game, subjects face subsequently first a risk preference elicitation task and then a distributional preferences elicitation procedure. With respect to social preferences the authors find that ahead averse subjects tend to lie less in the sender-receiver game.<sup>5</sup> One possible interpretation of this

---

recipient, and spiteful black lies harm both parties. In this taxonomy, *NEG* corresponds to the context of selfish black lies, while *NEU* sits at the border between selfish black lies and Pareto white lies.

<sup>4</sup> Actually, the authors use a more indirect procedure where the computer randomly chooses a shape – a sun or a star – and the child is asked to report the shape on the screen. However, since each shape is associated with one of the two allocations used in the first stage this indirect procedure seems to be strategically equivalent to the procedure described in the main text.

<sup>5</sup> In the distributional preference elicitation procedure subjects are asked to make four choices between two allocations, each involving an own material payoff and a material payoff for another subject. As in the Equality Equivalence Test one of the two allocations involves an egalitarian distribution while

**Table 1**  
Choices in the Equality Equivalence Test.

Disadvantageous Inequality Block*					
LEFT		Your Choice (please mark)		RIGHT	
you get	passive person gets			you get	passive person gets
1.5 (0.8) Euro	3 (1.5) Euro	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	2 (1) Euro	2 (1) Euro
1.9 (0.9) Euro	3 (1.5) Euro	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	2 (1) Euro	2 (1) Euro
2 (1) Euro	3 (1.5) Euro	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	2 (1) Euro	2 (1) Euro
2.1 (1.1) Euro	3 (1.5) Euro	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	2 (1) Euro	2 (1) Euro
2.5 (1.2) Euro	3 (1.5) Euro	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	2 (1) Euro	2 (1) Euro
Advantageous Inequality Block*					
LEFT		Your Choice (please mark)		RIGHT	
you get	passive person gets			you get	passive person gets
1.5 (0.8) Euro	1 (0.5) Euro	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	2 (1) Euro	2 (1) Euro
1.9 (0.9) Euro	1 (0.5) Euro	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	2 (1) Euro	2 (1) Euro
2 (1) Euro	1 (0.5) Euro	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	2 (1) Euro	2 (1) Euro
2.1 (1.1) Euro	1 (0.5) Euro	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	2 (1) Euro	2 (1) Euro
2.5 (1.2) Euro	1 (0.5) Euro	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	2 (1) Euro	2 (1) Euro

\*The labels "Disadvantageous Inequality Block" and "Advantageous Inequality Block" were not shown to the subjects.

finding is that more pro-social subjects are more averse to lying. However, given that lying in their sender-receiver game is harmful for the receiver, their finding is also consistent with the hypothesis that more pro-social subjects have the same lying aversion as others but lie less because lying hurts the other party. In this respect our experiments are more informative because in addition to a sender-receiver game where lying hurts the receiver, we also have treatments where lying has no effect on the payoffs of others. This allows to discriminate between the two explanations for the findings by [Sheremeta and Shields \(2013\)](#).

### 3. Experimental design and procedure

Our experimental design involves three treatments implemented between subjects. Each treatment consists of two parts. In Part 1 of each treatment we elicit the distributional preferences of subjects by exposing them to the Equality Equivalence Test introduced by [Kerschbamer \(2015\)](#). In Part 2 subjects are either exposed to a dice-rolling task à la [Fischbacher and Föllmi-Heusi \(2013\)](#) (treatment *DICE*) or they play one of two versions of a deception game à la [Erat and Gneezy \(2012\)](#) (treatments *NEG* and *NEU*).

#### 3.1. Part 1: The equality equivalence test

In the Equality Equivalence Test (called *EET* henceforth) each subject is exposed to a series of binary choices between allocations that both involve an own payoff for the decision maker and a payoff for a randomly matched anonymous second subject, the passive person.<sup>6</sup> In the version of the test implemented in our experiments subjects are exposed to ten binary choices. In each of the ten binary decision problems one of the two allocations is symmetric (i.e., egalitarian – giving the same material payoff to both agents) while the other one is asymmetric (involving unequal payoffs for the two subjects). In half of the choice tasks the asymmetric allocation involves disadvantageous inequality (that is, the decision maker receives a lower material payoff than the second subject), while in the other half it involves advantageous inequality (that is, the decision maker receives a higher payoff than the second subject). In both domains the *EET* systematically varies the price of giving (or taking) by increasing the material payoff of the decision maker in the asymmetric allocation while keeping all other payoffs constant. The ten binary choices were presented to subjects in two blocks – as shown in [Table 1](#).<sup>7</sup>

Given the design of the *EET*, a rational decision maker switches – in each block – at most once from the symmetric to the asymmetric allocation (and never in the other direction).<sup>8</sup> As shown by [Kerschbamer \(2015\)](#) the switching points in the two blocks can be used to construct a two-dimensional index representing the archetype of distributional concerns

the second allocation creates advantageous or disadvantageous inequality. Subjects are then classified as either ahead-averse, or behind-averse or pro-social according to the choices in one or two of the four tasks. This means that a subject might end up in more than one of these classes or in none of them.

<sup>6</sup> We employed a double role assignment protocol similar to the one used by [Andreoni and Miller \(2002\)](#) in their dictator games. This means that in our protocol each subject makes distributional choices, and each subject gets two payoffs, one as an active decision maker and one as a passive person.

<sup>7</sup> The payoffs of the Equality Equivalence Test differed slightly across treatments. In the two versions of the deception game (that is, in treatments *NEG* and *NEU*), we used the test version given in [Table 1](#) while in the dice-rolling task (treatment *DICE*) the payoffs were the ones shown in parentheses.

<sup>8</sup> Around 3% of the subjects failed this basic consistency check and we excluded them from our analysis.

**Table 2**  
Payoffs (in Euros) in the two different treatments.

Action Choice of Receiver	NEG		NEU	
	Sender	Receiver	Sender	Receiver
HEADS (= actual outcome)	2	3	2	3
TAILS	3	2	3	3

and preference intensity: The *x*-score (ranging from  $-2.5$  to  $+2.5$  in integer steps) measures pro-sociality in the domain of disadvantageous inequality, and the *y*-score (again ranging from  $-2.5$  to  $+2.5$  in integer steps) measures pro-sociality in the domain of advantageous inequality. In both domains a positive (negative) score means benevolence (malevolence) and a higher score means ‘more benevolent’ (‘less malevolent’). Also, in both domains scores in  $\{-0.5, 0.5\}$  are consistent with neutrality.

Combining the information about benevolence, neutrality or malevolence of a decision maker in the two domains allows classifying subjects into archetypes of distributional preferences. Specifically, we define the following types:

- ALT: a decision maker who has both scores positive and at least one of them strictly larger than 0.5 is classified as altruist;
- SEL: a decision maker who has both scores in  $\{-0.5, 0.5\}$  is classified as selfish;
- OTHERS: subjects that do not fall in any of the categories defined above are classified as ‘others’.<sup>9,10</sup>

### 3.2. Part 2: The three treatments

#### 3.2.1. Treatment DICE: The dice-rolling game

For the *DICE* treatment we used a modified version of the dice-rolling game by Fischbacher and Föllmi-Heusi (2013): Subjects were asked to throw a die in privacy ten times and to report the outcomes on the computer screen. Subjects were informed that they would receive the sum of the reported outcomes in ECUs as payment for this part of the experiment, with 5 ECUs = 1 Euro. The experimental environment was such that neither the experimenter nor any other subject could possibly observe the outcome of the die roll. This assures complete anonymity and gives subjects the possibility to cheat with virtually no costs of lying. The payment was carried out by a second person in absence of the experimenter to further reduce any possible demand effects and this was common knowledge at the beginning of the experiment.

#### 3.2.2. Treatments NEG and NEU: Two versions of the deception game

For the *NEG* and the *NEU* treatment we used a modified version of the deception game by Erat and Gneezy (2012): At the start of each session we randomly assigned one of two roles to subjects, either the role of a sender or the role of a receiver. Then we informed subjects that we would randomly form pairs, each consisting of a sender and a receiver.<sup>11</sup> We then informed the senders that we have tossed a fair coin before the start of the experiment and that the outcome of the coin toss was HEADS.<sup>12</sup> We also informed the sender that the payments to the two subjects in a pair would only depend on the action choice of the receiver. Specifically, the receiver in each pair has two available actions – HEADS and TAILS – and the payments to the two subjects as a function of the action chosen by the receiver were as shown in Table 2. The sender was informed that his or her task in this part of the experiment was to send one of two possible messages to the receiver – either Message 1 or Message 2:

Message 1: “The outcome of the coin flip is HEADS”.

Message 2: “The outcome of the coin flip is TAILS”.

The receiver observes the message sent by the sender and then has to choose which option (HEADS or TAILS) to implement without knowing the monetary consequences of the two options. It is important to note that at the end of the game, the receiver only learns his own payoff from the implemented option. The receiver does not learn the payoff of the sender from the implemented option nor the payoffs of the option not implemented and all of this is common knowledge at the beginning of the game.

Table 2 summarizes the monetary consequences of the two options for each treatment in Euros.<sup>13</sup> Option HEADS implements the same allocation – a material payoff of 2 Euros for the sender and one of 3 Euros for the receiver – in both

<sup>9</sup> Unfortunately, we have too few equality averse, inequality averse and spiteful subjects to make useful inferences for these distributional preference types. Therefore, we include these observations only in the aggregate results section and concentrate on the ALT vs. SEL comparison in the further analysis.

<sup>10</sup> In Appendix B, we show how the classification in distributional preference types works in detail.

<sup>11</sup> In the instructions we use the terms “Group B member” and “Group A member” instead of the terms sender and receiver.

<sup>12</sup> We informed the subjects on request that we performed a single coin toss for the whole experiment and that the payoffs in the instructions were chosen according to the outcome of the coin flip.

<sup>13</sup> The experimental currency was ECUs, with 10 ECUs = 1 Euro.

treatments. Option TAILS implements the same payoff (of 3 Euros) for the sender in both treatments. So, the treatments differ only in the material consequences of option TAILS for the receiver. In *NEG*, the material consequence of option TAILS for the receiver is 2 Euros and in *NEU*, option TAILS implies a monetary payoff of 3 Euros for the receiver. Under the assumption that subjects are exclusively interested in their own material income, senders have an incentive to induce the receivers to choose TAILS in both treatments.

The experiment was conducted with paper-and-pen (and several other design features reported below were applied) to convince subjects that neither other subjects nor the experimenters could identify the person who has made any particular decision. This was done in an attempt to minimize the impact of experimenter demand and audience effects.<sup>14</sup>

### 3.3. Experimental procedures

Sixteen experimental sessions were conducted at the University of Innsbruck from November 2011 to June 2014. 356 subjects from various academic backgrounds participated in total and each subject participated in one session only. The invitations were sent out using either the ORSEE recruiting system (Greiner, 2015) or the Hroot recruiting system (Bock et al., 2014).<sup>15</sup>

After arrival subjects assembled in the laboratory and then instructions for Part 1 of the experiment – the *EET* – were distributed and read aloud. Instructions for the *EET* informed subjects that (i) their earnings for this part of the experiment would be determined at the end of the experiment; (ii) they would receive two cash payments for this task, one as a decision maker and one as a passive person; (iii) for their earnings as a decision maker one of the 10 decision problems would be selected randomly and the alternative chosen in that decision problem would be paid out; and (iv) their earnings as a passive person would come from another participant (that is, not from the passive person of the subject under consideration). After reading the instructions and answering questions from the subjects in private, subjects were asked to make their decisions for the *EET*.

Then the instructions for Part 2 of the experiment – either the *DICE*, the *NEG* or the *NEU* part, depending on the treatment – were distributed and read aloud. Instructions for *DICE* informed subjects that their payment for this part of the experiment would only depend on their own choice and that their own choice would not have any consequences for the payments of other subjects. Subjects were also informed that (i) their task in this part of the experiment was to roll a six-sided dice ten times and to self-report the outcomes on the computer screen; (ii) the sum of the reported outcomes would determine their payment for this part of the experiment – with each point yielding an ECU and 5 ECUs translating in 1 Euro; and (iii) that cash payments will be carried out after the experiment by a third person, who is not informed about the experiment's procedure and content. After reading the instructions and answering questions from the subjects in private, subjects were asked to start with the experiment.

Instructions for the *NEG* and the *NEU* treatment informed subjects that (i) there are two roles in this part of the experiment, the role of a 'Group A member' and the role of a 'Group B member'; (ii) each Group A member is matched with exactly one Group B member and vice versa, and that at no point in time will a participant discover the identity of the person she/he is matched with; (iii) Group B members are informed about the outcome of a coin toss and Group A members are not; (iv) Group B members are called to send a message about the outcome of the coin toss to Group A members and that the choice of the message has no implications on the payoffs of any of the two players; (v) Group A members are called to choose either 'HEADS' or 'TAILS' after observing the message from Group B members and that this choice determines the payoffs for both players; (vi) Group A members know nothing about the payoff consequences for both players associated with each choice and that Group B members are aware about the payoff consequences for both players related to each choice of Group A members; (vii) Group A members will learn only their own payoff out of the chosen option and nothing else (even after the experiment ended); and (viii) cash payments could be collected a few days after the experiment from one of the secretaries together with the payments from Part 1 of the experiment. After reading the instructions and answering questions from the subjects in private, Group B members were handed out a decision sheet and an empty envelope and they were asked to fill out the decision sheet in private. After the Group B members made their decisions, they put the decision sheets into the unmarked envelopes. Envelopes were then collected with a letterbox by an experimenter. Additionally, Group B members were asked to answer the following question:

Q1: "Out of 100 Group A members, how many do you think follow the message of the paired Group B member on average?"<sup>16</sup>

In the meantime another experimenter opened the envelopes of the Group B members (this was done in a third room without subjects) and transferred the messages of the Group B members on the decision sheets of the Group A members. Then, Group A members were handed out the decision sheets in an unmarked envelope randomly and they were asked to fill out the decision sheet in private. After the Group A members made their decisions, they put the decision sheets into the unmarked envelopes and the envelopes were then collected with a letterbox by an experimenter.

<sup>14</sup> Our experimental design is inspired by the (almost) double blind procedures employed by Hoffman et al. (1994), Cox (2004) and Cox and Sadiraj (2012). See List (2007) and Zizzo (2010) for a discussion on experimenter demand effects and Hoffman et al. (1994), Andreoni and Petrie (2004), and Andreoni and Bernheim (2009) for experimental evidence indicating that audience effects might have a large impact on subjects' behavior in dictator-game like situations.

<sup>15</sup> The University of Innsbruck changed from the ORSEE to the Hroot system in February 2014.

<sup>16</sup> This question was incentivized in the following way: in each session we paid the subject whose answer was closest to the true value 5 Euros.

## 4. Hypotheses

All our predictions are based on the null hypothesis that all distributional preference types have the same aversion against lying. This is only a working hypothesis, of course – ultimately, our main research question is exactly whether distributional archetypes differ systematically in their lying aversion. We are agnostic about the nature of the lying aversion and simply assume – in line with much of the rest of the literature – that at least some subjects have such an aversion.<sup>17</sup>

### 4.1. Prediction for DICE

Since in the dice-rolling task lying has only consequences for the own material payoff and not for the payoffs of others and since distributional types are assumed to differ only in their attitude towards the material payoff of others our null hypothesis directly leads to the following prediction:

**Hypothesis 1.** *There is no systematic difference between selfish and altruistic subjects in the frequency and extend of lying in DICE.*

Let  $\sigma^x$  stand for the average sum of reported outcomes in DICE among subjects classified as being of distributional type  $x \in \{SEL, ALT\}$ . Using this notation we test H1 by comparing  $\sigma^x$  across distributional preference types.

### 4.2. Prediction for NEU

Since in the NEU version of the deception game lying has only consequences for the own material payoff and not for the payoffs of others and since distributional preference types are assumed to differ only in their attitude towards the material payoff of others our null hypothesis directly leads to the following prediction:

**Hypothesis 2.** *There is no systematic difference between selfish and altruistic subjects in the frequency of lying in NEU.*

Let  $f_{NEU}^x$  stand for the frequency of message 'Tails' among subjects classified as being of distributional type  $x \in \{SEL, ALT\}$  in treatment NEU. Using this notation we test H2 by comparing  $f_{NEU}^x$  across types.

### 4.3. Prediction for NEG

Selfish subjects do not care for the material payoff of others, they simply maximize their own material payoff. Since the payoff of the sender increases to the same extend by lying in NEG and in NEU, our null hypothesis leads directly to the following prediction:

**Hypothesis 3.** *There is no systematic difference in the frequency of lying of selfish subjects between NEG and NEU.*

Using the notation introduced above we test H3 by comparing  $f_{NEG}^{SEL}$  to  $f_{NEU}^{SEL}$ .

Altruistic subjects are willing to give up own material payoff to increase the material payoffs of others. Since the payoff of the receiver is negatively affected by lying in NEG but independent of the message in NEU our null hypothesis directly leads to the following prediction:

**Hypothesis 4.** *The frequency of lying among altruistic subjects is lower in NEG than in NEU.*

We will test this hypothesis by comparing  $f_{NEG}^{ALT}$  to  $f_{NEU}^{ALT}$ .

Hypotheses 2, 3 and 4 together imply:

**Hypothesis 5.** *The frequency of lying in NEG is higher for selfish subjects than for altruistic subjects.*


Fig. 1 illustrates hypotheses 2 - 5 in a graphical way:

## 5. Results

### 5.1. Results of the equality equivalence test

Table 3 reports the results of the EET. It shows the distributional preference types of subjects who participated in our treatments, except for those who failed the consistency check of the test (at most one switch from Right to Left and no switch in the other direction in each of the two blocks of choices displayed in Table 1). In total 12 subjects failed the consistency check and those subjects are excluded from our further analysis.

<sup>17</sup> Dufwenberg and Gneezy (2000) assume that some people are averse against lying because they feel guilty if they disappoint others' payoff expectations and Battigalli and Dufwenberg (2007) formulate the related concept of 'guilt from blame' where a player cares about others' inferences regarding their willingness to disappoint others' payoff expectations. Inferences are also important in several recent explanations for lying aversion: While in Gneezy et al. (2016) and in Abeler et al. (2016) a decision maker derives disutility proportional to the probability others assign to the fact that the decision maker lies, in Dufwenberg and Dufwenberg (2017) the decision maker gets disutility proportional to the degree in which he is perceived to cheat.



**Fig. 1.** Frequency of message Tails for altruistic and selfish subjects in the NEG and NEU treatment.

**Table 3**

Proportion of subjects according to their distributional preferences.

	NEG	NEU	DICE
	<i>N</i> = 104	<i>N</i> = 108	<i>N</i> = 132
ALT	33.65% (35)	46.30% (50)	34.85% (46)
SEL	42.31% (44)	37.04% (40)	56.06% (74)
OTHERS	25.04% (25) <sup>a</sup>	16.66% (18) <sup>b</sup>	9.09% (12) <sup>c</sup>

<sup>a</sup> Values in parentheses correspond to the absolute number of observations. In treatment NEG, the group OTHERS consists of 19 inequality averse and 6 spiteful subjects.

<sup>b</sup> In treatment NEU, the group OTHERS consists of 3 equality averse, 12 inequality averse and 3 spiteful subjects.

<sup>c</sup> In treatment DICE, the group OTHERS consists of 1 equality averse, 6 inequality averse and 5 spiteful subjects.

The type distributions are similar across treatments with some notable exceptions: SEL are less and ALT are more frequent in NEU than in the other two treatments, and the OTHERS category is overrepresented in NEG. While the differences in the type distributions across treatments are remarkable, they are not important for our further analysis since our main comparisons are comparisons across types within a given treatment and comparisons across treatments for a given type.

### 5.2. Aggregate results: DICE

The average outcome of the throw of a fair die is 3.5. If cheating was not an issue the average outcome in DICE should therefore be 35. However, cheating occurred on a broad basis and the mean outcome in DICE differs significantly from 35 (mean = 43.11;  $p = .000$ , single sample T-test). On the one extreme is a fraction of 3.8% of subjects who report 6 for each of the 10 die throws. On the other extreme there is a fraction of 37.1% of subjects who cannot be identified as cheaters on a level of 85%. These findings are well in line with similar results in the literature (see Mazar et al. (2008), for instance). Fig. 2 shows the distribution of the sum of the reported throws and Fig. 3 shows the relative frequencies of all reported die throws.

### 5.3. Aggregate results: NEG and NEU

The results of the two versions of the deception games aggregated over all distributional preference types are presented in Table 4. The figures in row [1] indicate the relative frequency of lying, i.e. sending the incorrect message TAILS. Assuming that senders believe that receivers will follow their message (with probability one) the gain for the sender from lying is 1 Euro in both versions of the deception game. In NEG, where the loss for the receiver is also 1 Euro, 44% of the senders lie while in NEU, where the receivers gain nothing and loose nothing compared to HEADS, the fraction of senders who lie is 75%. These figures suggest that the fraction of subjects sending message TAILS is significantly different across treatments which is indeed the case ( $p = .000$ , Fisher's exact).



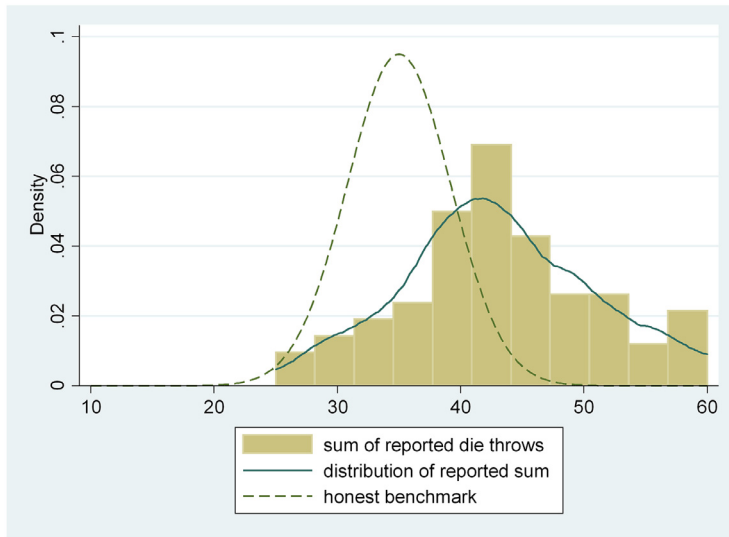


Fig. 2. Histogram of the sum of the reported die throws.

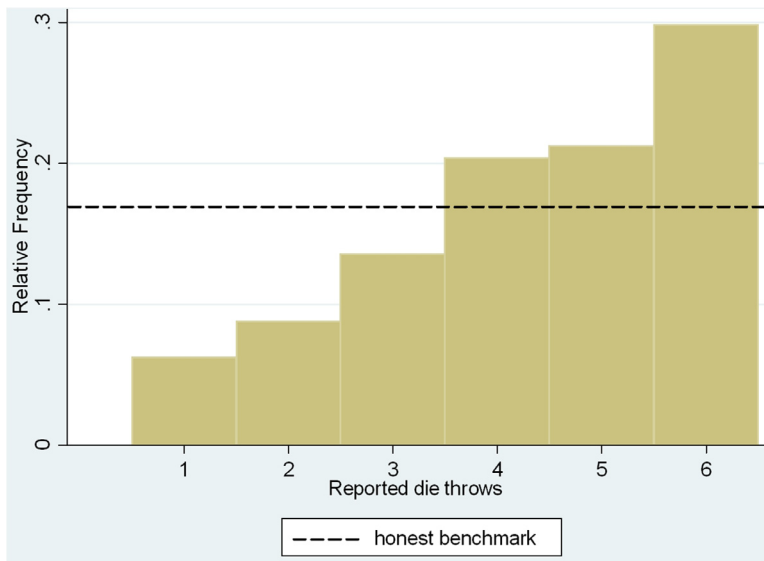


Fig. 3. Relative frequencies of all reported die throws.

**Table 4**  
Results for NEG and NEU.

	NEG	NEU
[1] sender chooses message TAILS	0.44	0.75
	(46)	(81)
[2] sender chooses message TAILS contingent on answers to Q1 > 50	0.56	0.82
	(34)	(69)
[3] receiver implements sender's message	0.71	0.76
	(80)	(82)
[4] mean of senders' answer to Q1 (divided by 100)	0.57	0.69
	(104)	(106) <sup>a</sup>
[5] proportion of senders who answer Q1 with values > 50	0.59	0.79
	(61)	(84)

<sup>a</sup> Values in parentheses correspond to the actual number of observations. Two subjects did not answer Q1.

**Table 5**  
Means of the Reported Sums in DICE.

ALT	43.13 (46)
SEL	43.39 (74)

Values in parentheses correspond to the actual number of observations.

The assumption that all senders believe that receivers will follow their message (with probability one) is unrealistic, of course. It is therefore more meaningful to condition the chosen message on the belief of the sender. To do so, we classify the senders into two groups, depending on their answer to Q1: In the first group we include all senders who answered Q1 with values between 0 and 49, and in the second group we put all senders who answered Q1 with values between 51 and 100. In line with [Sutter \(2009\)](#), we argue that sending the incorrect message TAILS within the first group has a different interpretation than sending the incorrect message TAILS within the second group.<sup>18</sup> Row [2] of [Table 4](#) presents the relative frequencies of sending the incorrect message TAILS contingent on the answer to Q1 having a value larger than 50. In *NEG*, 56% of the senders lie according to this modified lying definition while in *NEU* 82% of the senders send the incorrect message TAILS. The frequency of lie-telling according to this modified lying definition is again significantly different across treatments ( $p = .001$ , Fisher's exact). These results are in line with the findings in the previous literature on cheap-talk sender-receiver games indicating that people are sensitive to the material consequences for themselves and for the other party when deciding to lie. See [Gneezy \(2005\)](#), [Erat and Gneezy \(2012\)](#), [Gibson et al. \(2013\)](#) and [Gneezy et al. \(2013\)](#), for instance.

Row [3] of [Table 4](#) reports the relative frequencies of receivers actually following the sender's message, and there is no significant difference across the two treatments ( $p = .495$ , Fisher's exact). This is not surprising because the receivers have no information about the payoffs of the two players associated with the different outcomes in any of the two treatments.

Row [4] of [Table 4](#) shows the means of the answers to Q1. In *NEG* senders think that around 57 receivers (out of 100) and in *NEU* senders believe that 69 (out of 100) follow their message. The distributions of the answers to Q1 are significantly different across treatments ( $p = .000$ , Mann-Whitney). Furthermore, row [5] of [Table 4](#) shows the proportion of senders who answer Q1 with values between 51 and 100 for each treatment. The differences between the proportions are statistically significant ( $p = .002$ , Fisher's exact). This is in line with [Camerer et al. \(1989\)](#), who argue that it is difficult for an informed party to neglect own information (i.e., the non-aligned payoff structure in *NEG* and the aligned payoff structure in *NEU*) when forming expectations about how an uninformed party will behave.

#### 5.4. Test of hypotheses

[Table 5](#) reports the means of the reported sums in *DICE* grouped by distributional preference type. Selfish subjects report on average 43.39 and altruists report on average 43.13. These figures suggest that there is no significant difference across distributional preference types in the frequency and extend of lying. This is confirmed by statistical tests ( $p = .8624$ , T-Test).

**Result 1.** *In line with Hypothesis 1, there is no systematic difference between selfish and altruistic subjects in the frequency and extend of lying in the dice-rolling task.*

[Fig. 4](#) shows the proportions of senders who sent the incorrect message TAILS in the *NEG* and the *NEU* version of the deception game, grouped by distributional preference type and contingent on the answer to Q1 being larger than 50.<sup>19,20</sup> In *NEU*, 77% of the selfish subjects and 87% of the altruist subjects lie. These figures suggest that – in line with the theoretical prediction – there is no significant difference in lying behavior in *NEU* across distributional preference types. The pairwise comparison across types confirms this inference: For SEL vs. ALT Fisher's exact test yields  $p = .363$ . We therefore conclude:

**Result 2.** *In line with Hypothesis 2, there is no systematic difference between selfish and altruistic subjects in the frequency of lying in the *NEU* version of the deception game.*

The proportion of lies within the group of selfish subjects is 75% in *NEG* and 77% in *NEU*. These figures suggest that – in line with the theoretical prediction – there is no significant difference in lying behavior of selfish subjects between *NEG* and *NEU*. Again, this is confirmed by statistical tests ( $p = 1.000$ , Fisher's exact).

<sup>18</sup> See [Appendix A](#) for alternative classifications according to the answers to Q1.

<sup>19</sup> Overall, 25 ALT and 20 SEL subjects answer Q1 with values larger than 50 in *NEG* and 38 ALT and 34 SEL subjects do so in *NEU*. By contrast, we have only 12 IAV subjects in *NEG* and 7 IAV subjects in *NEU* who answer Q1 with values larger than 50. This low number of IAV subjects was the reason to concentrate on the comparison between ALT and SEL in our main analysis.

<sup>20</sup> Our results also survive an alternative classification where we count only those subjects as liars who send the incorrect message expecting that it is quite likely that the other person will follow ( $Q1 \geq 70$ ). They are also robust to applying the definition of deception introduced by [Sutter \(2009\)](#). See [Appendix A](#) for details.

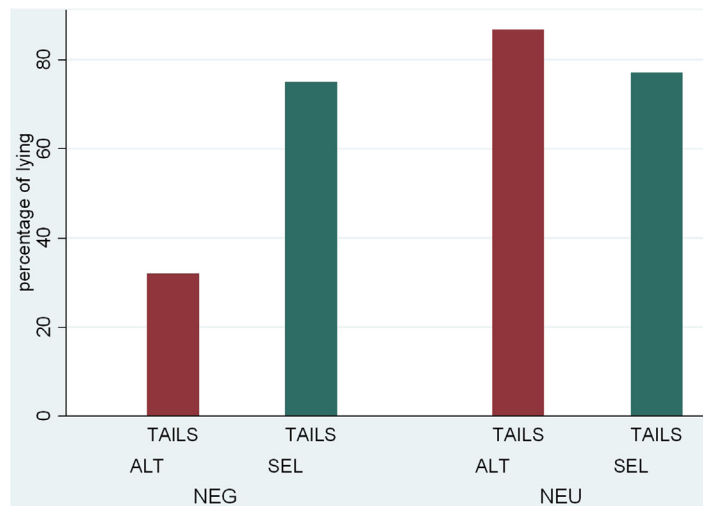


Fig. 4. Proportions of senders who lie contingent on answers to Q1 > 50 grouped by their distributional preference type.

**Result 3.** In line with Hypothesis 3, there is no systematic difference in the frequency of lying of selfish subjects between the NEG and NEU version of the deception game.

The proportion of lies within the group of altruistic subjects is 32% in NEG and 87% in NEU. These figures suggest that – in line with the theoretical prediction – there is a significant difference in the frequency of lie-telling of altruistic subjects between NEG and NEU. The pairwise comparison across treatments confirms this conjecture: For altruistic subjects the difference between NEG and NEU in the frequency of dishonest messages is highly significant ( $p = .000$ , Fisher's exact).

**Result 4.** In line with Hypothesis 4, the frequency of lying among altruistic subjects is lower in the NEG than in the NEU version of the deception game.

In NEG, 32% of the altruistic subjects and 75% of the selfish subjects lie. These figures suggest that – in line with the theoretical prediction – there is a significant difference in lying behavior in NEG between selfish and altruistic subjects. The pairwise comparison across types confirms that the difference in the frequency of lying between SEL and ALT subjects is highly significant ( $p = .007$ , Fisher's exact).

**Result 5.** In line with Hypothesis 5, the frequency of lying in the NEG version of the deception game is higher for selfish subjects than for altruistic subjects.

## 6. Conclusion

The results presented in this paper indicate that altruistic subjects lie less when lying hurts another party but do not lie less than selfish subjects in environments where lying helps the liar but has no effects on the payoffs of others. Specifically, our data suggests the following two main conclusions: First, our results indicate a relatively consistent relationship between the behavior in the EET and the behavior in the treatments DICE, NEG and NEU. This finding is not obvious because previous studies often find low predictive power of measures of distributional preferences for related games (see Blanco et al., 2011). Second, our findings suggests that the norm of distributive justice is uncorrelated with the truth-telling norm at the individual level. Indeed, with our data we were not able to reject our null hypothesis that lying costs – whatever their origin – are distributed similarly across agents of different distributional preference types.

If the findings of this paper are confirmed in future experimental research, they have positive and negative implications for the selection of agents for jobs in markets plagued by asymmetric information. The negative side of our findings is that screening agents along a single dimension of morality does not necessarily yield moral popes: An agent who has a long history of altruistic behavior (documented, for instance, by an impressive track record of volunteer work in his or her CV) might still be inclined to be dishonest when dishonesty helps the liar but does not have negative consequences for relevant others. The positive side of our findings is that agents who have been selected because of their altruism are, on average, more likely to be inclined to take the consequences of their behavior for others into account when deciding whether to stay honest in a business relation.

## Declarations of interest

None.




Fig. 5. Frequencies of answers to Q1 in the treatments NEG and NEU.

**Table 6**  
Results for NEG and NEU.

	NEG	NEU
[1] sender chooses message TAILS	0.44 (46)	0.75 (81)
[2] sender chooses message TAILS contingent on answers to Q1 > 50	0.56 (34)	0.82 (69)
[3] sender chooses message TAILS contingent on answers to Q1 $\geq$ 70	0.60 (21)	0.90 (60)
[4] liars and sophisticated liars according to the definition proposed by Sutter (2009)	0.62 (56)	0.81 (79)

Values in parentheses correspond to the actual number of observations.

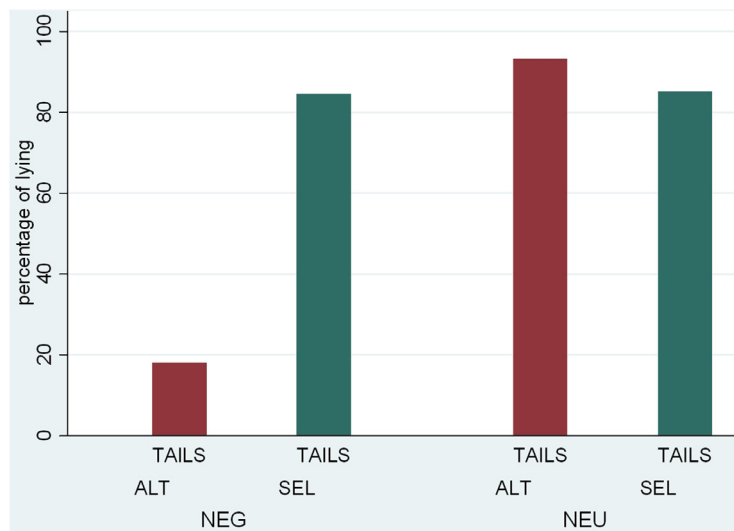
## Appendix A. Robustness of results for NEG and NEU to alternative classifications of lying

Fig. 5 shows the answers to Q1 for the treatments *NEG* and *NEU*. As indicated in Section 5.3 the answers differ between the two treatments and a Kolmogorov–Smirnov Test confirms that the difference in distributions is significant ( $p = .000$ ).

In Section 5.3, we classified the senders into two groups, depending on their answer to Q1: In the first group we included all senders who answered Q1 with values between 0 and 49, and in the second group we put all senders who answered Q1 with values between 51 and 100. In the analysis we then focused on the behavior of subjects in the second group. In this section, we use two alternative classifications of lying based on the answers to Q1 and check if the results based on these alternative classifications are in line with the results reported in Section 5.3.

First, we focus on subjects who are relatively sure that the other person will follow their message by examining only subjects who answered Q1 with values larger or equal than 70. Row [3] of Table 6 presents the relative frequencies of sending the incorrect message TAILS contingent on the answer to Q1 having a value larger or equal than 70 (we show in Table 6 also the results according to the other classifications for purposes of comparison). In *NEG*, 60% of the senders lie according to this modified lying definition while in *NEU* 90% of the senders send the incorrect message TAILS. The frequency of lie-telling according to this modified lying definition is again significantly different across treatments ( $p = .001$ , Fisher's exact). These results are in line with the findings reported in Section 5.3 for the aggregate data.

Second, we analyze the results in the same way as Sutter (2009) by classifying senders who answered Q1 with values larger than 50 and send the incorrect message TAILS as liars, and additionally classifying senders who answered Q1 with values smaller than 50 and send the correct message HEADS as sophisticated liars. Row [4] of Table 6 presents the relative frequencies according to this definition. In *NEG*, 62% of the senders lie according to this lying definition while in *NEU* 80%



**Fig. 6.** Proportions of senders who lie contingent on answers to  $Q1 \geq 70$  grouped by their distributional preference type.

of the senders lie. The frequency of lying according to this definition is once again significantly different across treatments ( $p = .006$ , Fisher's exact) and the results are again in line with the findings reported in Section 5.3 for the aggregate data.

We now perform the same tests as in Section 5.4 based on these alternative definitions of lying. Since the results remain unchanged we copy the text from Section 5.4 and only change the numerical values.

*Test of the hypotheses based on an alternative classification where only those subjects are counted as liars who send the incorrect message expecting that it is quite likely that the other will follow ( $Q1 \geq 70$ )*

Fig. 6 shows the proportions of senders who sent the incorrect message TAILS in the NEG and the NEU version of the deception game, grouped by distributional preference type and contingent on the answer to Q1 being equal or larger than 70.<sup>21</sup> In NEU, 85% of the selfish subjects and 93% of the altruists subjects lie. These figures suggest that – in line with the theoretical prediction – there is no significant difference in lying behavior in NEU across distributional preference types. The pairwise comparison across types confirms this suspicion: For SEL vs. ALT Fisher's exact test yields  $p = .408$ . We therefore conclude:

**Result 2.1.** *In line with Hypothesis 2, there is no systematic difference across distributional preference types in the frequency of lying in the NEU version of the deception game.*

The proportion of lies within the group of selfish subjects is 85% in NEG and 85% in NEU. These figures suggest that – in line with the theoretical prediction – there is no significant difference in lying behavior of selfish subjects between NEG and NEU. Again, this is confirmed by statistical tests ( $p = 1.000$ , Fisher's exact).

**Result 3.1.** *In line with Hypothesis 3, there is no systematic difference in the frequency of lying of selfish subjects between the NEG and NEU version of the deception game.*

The proportion of lies within the group of altruistic subjects is 18% in NEG and 93% in NEU. These figures suggest that – in line with the theoretical prediction – there is a significant difference in the frequency of lie-telling of altruistic subjects between NEG and NEU. The pairwise comparison across treatments confirms this conjecture: For altruistic subjects the difference between NEG and NEU in the frequency of dishonest messages is highly significant ( $p = .000$ , Fisher's exact).

**Result 4.1.** *In line with Hypothesis 4, the frequency of lying among altruistic subjects is lower in the NEG than in the NEU version of the deception game.*

In NEG, 18% of the altruistic subjects and 85% of the selfish subjects lie. These figures suggest that – in line with the theoretical prediction – there is a significant difference in lying behavior in NEG between selfish and altruistic subjects. The pairwise comparison across types confirms that the difference in the frequency of lying between SEL and ALT subjects is highly significant ( $p = .003$ , Fisher's exact).

<sup>21</sup> Overall, 11 ALT and 13 SEL subjects answer Q1 with values equal or larger than 70 in NEG and 30 ALT and 26 SEL subjects do so in NEU. According to this classification we have only 7 IAV subjects in NEG and 7 IAV subjects in NEU and this low number of IAV types was the reason to concentrate on the comparison of ALT vs. SEL types.




Fig. 7. Proportions of senders who lie according to the definition of Sutter (2009) grouped by their distributional preference type.

**Result 5.1.** *In line with Hypothesis 5, the frequency of lying in the NEG version of the deception game is higher for selfish subjects than for altruistic subjects.*

*Test of the hypotheses based on the definition of lying proposed by Sutter (2009)*

Fig. 7 shows the proportions of liars in the NEG and the NEU version of the deception game, grouped by distributional preference type and contingent on the answer to Q1 (senders who answered Q1 with values larger than 50 and send the incorrect message TAILS and senders who answered Q1 with values smaller than 50 and send the correct message HEADS are classified as liars in the following analysis).<sup>22</sup> In NEU, 78% of the selfish subjects and 84% of the altruists lie. These figures suggest that – in line with the theoretical prediction – there is no significant difference in lying behavior in NEU across distributional preference types. The pairwise comparison across types confirms this suspicion: For SEL vs. ALT Fisher's exact test yields  $p = .567$ . We therefore conclude:

**Result 2.2.** *In line with Hypothesis 2, there is no systematic difference across distributional preference types in the frequency of lying in the NEU version of the deception game.*

The proportion of lies within the group of selfish subjects is 74% in NEG and 78% in NEU. These figures suggest that – in line with the theoretical prediction – there is no significant difference in lying behavior of selfish subjects between NEG and NEU. Again, this is confirmed by statistical tests ( $p = .789$ , Fisher's exact).

**Result 3.2.** *In line with Hypothesis 3, there is no systematic difference in the frequency of lying of selfish subjects between the NEG and NEU version of the deception game.*

The proportion of lies within the group of altruistic subjects is 46% in NEG and 84% in NEU. These figures suggest that – in line with the theoretical prediction – there is a significant difference in the frequency of lie-telling of altruistic subjects between NEG and NEU. The pairwise comparison across treatments confirms this conjecture: For altruistic subjects the difference between NEG and NEU in lying is highly significant ( $p = .000$ , Fisher's exact).

**Result 4.2.** *In line with Hypothesis 4, the frequency of lying among altruistic subjects is lower in the NEG than in the NEU version of the deception game.*

In NEG, 46% of the altruistic subjects and 74% of the selfish subjects lie. These figures suggest that – in line with the theoretical prediction – there is a significant difference in lying behavior in NEG between selfish and altruistic subjects. The pairwise comparison across types confirms that the difference in the frequency of lying between SEL and ALT subjects is highly significant ( $p = .028$ , Fisher's exact test).

**Result 5.2.** *In line with Hypothesis 5, the frequency of lying in the NEG version of the deception game is higher for selfish subjects than for altruistic subjects.*

<sup>22</sup> Overall, 33 ALT and 38 SEL subjects answer Q1 with values not equal than 50 in NEG and 45 ALT and 36 SEL subjects do so in NEU. According to this classification we have 15 IAV subjects in NEG and 12 IAV subjects in NEU.

**Table 7**  
Determination of the x- and y-scores.

subject chooses LEFT for the 1 <sup>st</sup> time in row	in the X-list (Decision Nr. 1 - 5)	in the Y-list (Decision Nr. 6 - 10)
	x-score	y-score
1	2.5	-2.5
2	1.5	-1.5
3	0.5	-0.5
4	-0.5	0.5
5	-1.5	1.5
never	-2.5	2.5

## Appendix B. Identification of distributional preferences with the equality equivalence test

In this part of the appendix we demonstrate the determination of the x- and the y-score for the *EET* version employed in our experiment. The employed test version is displayed in Table 1 of Section 3.1. The upper part of the table shows the 5 binary choices in the disadvantageous inequality block – where in the asymmetric allocation (option LEFT) the decision maker receives a lower payoff than the passive person. The lower part shows the corresponding choices for the advantageous inequality block. In both parts the *EET* systematically varies the price of giving (or taking) by increasing the material payoff of the decision maker in the asymmetric allocation while keeping all the other payoffs constant. Given this design feature, in each of the two domains a rational decision maker switches at most once from the symmetric to the asymmetric allocation (and never in the other direction).<sup>23</sup> Furthermore, the switching points of the decision maker in the two domains can be used to construct the two scores – the x-score measuring the pro-sociality of the decision maker in the domain of disadvantageous inequality, and the y-score measuring pro-sociality in the domain of advantageous inequality.

Table 7 shows how the scores are determined: The x-score is defined as 3.5 minus the row number in which the subject decides the first time for the asymmetric allocation (with the convention that the subject who always decides for the symmetric allocation gets assigned a score of -2.5), while the y-score is defined as the row number in which the subject decides for the first time for the asymmetric allocation minus 3.5 (with the convention that the subject who always decides for the symmetric allocation gets assigned a score of +2.5). This way of defining the scores guarantees that in both domains a positive (negative) score means benevolence (malevolence) and a higher score means 'more benevolent' ('less malevolent').

Given the two scores of a subject we assign types of distributional preferences as follows:

- ALT: a decision maker who has both scores positive and at least one of them strictly larger than 0.5 is classified as altruist;
- EAV: a decision maker who has the x-score positive and the y-score negative and has in addition either the x-score strictly above 0.5 or the y-score strictly below -0.5 is classified as equality averse;
- IAV: a decision maker who has the x-score negative and the y-score positive and has in addition either the x-score strictly below -0.5 or the y-score strictly above 0.5 is classified as inequality averse;
- SEL: a decision maker who has both scores in  $\{-0.5, 0.5\}$  is classified as selfish;
- SPI: a decision maker who has both scores negative and at least one of them strictly below -0.5 is classified as spiteful.

## Appendix C. Experimental instructions

### Experimental Instructions for DICE

*In the following we provide an English translation of our originally German instructions. German instructions are available on request.*

*[General Instructions, at start of session]*

### Welcome to today's experiment.

The experiment consists of two parts and the two parts are completely independent from each other. First, part 1 will be explained and carried out, followed up by part 2. Your total payoff from today's experiment is calculated as the sum of your payoffs from part 1 and part 2.

The experiment is conducted on the computer screen, you make your decisions on the screen. All your choices and answers will remain anonymous.

Please do not talk to each other during the experiment. Please switch off your mobile phones. Please be aware that all tasks not related to the experiment such as surfing the internet, playing games on the computer or reading literature are not allowed and yield to exclusion from the experiment. Please raise your hand if you have any questions at any time.

<sup>23</sup> Rationality here means that the subject has well-behaved (that is, complete and transitive) preferences and that the preferences are strictly monotonic in the own material payoff.

### Payoff information

After finishing the experiment, the experimenter will leave the room and payment will be carried out by a third person, who is not informed about the experiment's procedure and content. You will then be called by your seat number individually. This procedure further guarantees your complete anonymity.

[Instructions for The Equality Equivalence Test]

Part 1 consists of 10 choices. In each of these choices you are anonymously matched with another participant. We call this participant "your passive person". You will later see why we do this. **Your passive person will be chosen randomly. At no time you will receive information about your passive person's identity, nor will your passive person get information about your identity.**

Talking about payment we are speaking of Tokens.

Conversion rate is:

**5 Tokens = 1 Euro**

Each of your decisions is a **choice between alternative "left" and alternative "right"**. Each of your decisions is a choice over allocations between you and your passive person.

*Example: You could be asked whether you like to choose alternative left, in which case you receive 5 tokens and your passive person receives 7.5 tokens, or alternative right in which case you receive 5 tokens and your passive person also receives 5 tokens. This decision problem would be presented on the screen as follows:*

LEFT		Your Choice (please mark)		RIGHT	
you get	passive person gets			you get	passive person gets
5 Tokens	7.5 Tokens	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	5 Tokens	5 Tokens

You are asked to make **10 such decisions** in part 1. At the end of the experiment, **one** of these 10 decisions will be randomly chosen as payoff relevant. Your payoff for this part is calculated as follows:

**Payoff as active person:** Suppose the above mentioned decision problem was chosen as payoff relevant and you have chosen alternative "right". You would then receive 5 tokens as active person and your passive person receives 5 tokens.

**Payoff as passive person:** Just like your passive person receives tokens from your decision without doing anything for it, you receive tokens from another participant's decision without doing anything for it, that is you are the passive person for this other participant. It is assured that you will not get the same person as active and passive person. That is, if person X is your passive person, you will certainly not get person X as active person.

Your total payoff for part 1 is calculated as the sum of your payoff as active person and your payoff as passive person. Attention: your decisions only influence your payoff as active person. Your payoff as passive person depends solely on the decisions of another participant.

[Instructions for DICE]

Your task in part 2 is to throw a die and report the outcome on the computer screen. You have 15 seconds time to throw the die and report the outcome on the screen. Please press *next throw* after you reported the outcome. You are asked to throw the die and report the outcome on the screen 10 times. At the end your reported outcomes are summed up to your "summed score" of this part.

### Payoff information

You receive your summed score 1:1 in tokens, paying you with **5 Tokens = 1 Euro** as payoff for part 2.

*Example: You have reported the following outcomes in part 2: 2, 5, 3, 6, 2, 5, 4, 5, 2 and 1, in which case your summed score of part 2 is 35. So, you receive your summed score of part 2 1:1 in tokens, paying you 35 tokens / 5 = 7 Euro as payoff for part 2.*

### Experimental Instructions for NEU and NEG

In the following we provide an English translation of our originally German instructions. We provide the instructions and the decision sheets for the treatment NEU here. The instructions and decision sheets for the treatment NEG are identical to those for treatment NEU (except for the payoffs of the receiver when the receiver chooses TAILS). German instructions are available on request.

[General Instructions, at start of session]

### Welcome to an experiment on decision making. We thank you for your participation!

The experiment consists of two parts and the two parts are completely independent from each other. A research foundation has provided the funds for conducting the experiment. Your total payoff from today's experiment is calculated as the sum of your payoffs from the two parts and the text below will tell you how the amount you earn will be determined.



The experiment consists of two parts and the two parts are completely independent from each other. First, part 1 will be explained and carried out, followed up by part 2. Your total payoff from today's experiment is calculated as the sum of your payoffs from part 1 and part 2.

### No Talking Allowed

Please do not talk to any other participant until the experiment is over. If there is anything that you don't understand, please raise your hand. An experimenter will approach you and clarify your questions in private. In about ten minutes this document will also be read aloud (by an experimenter).

### Anonymity

You will never be asked to reveal your identity to anyone during the experiment. Neither the experimenters nor the other subjects will be able to link you to any of your decisions. In order to keep your decisions private, please do not reveal your choices to any other participant. The following means help to guarantee anonymity:

### Non-Computerized Experiment and Private Code

The task you have to complete during the experiment is conducted in private on a printed form; that is, the experiment is not computerized. You have drawn a small sealed envelope from a box upon entering the room. PLEASE DO NOT OPEN YOUR ENVELOPE BEFORE THE EXPERIMENT STARTS. Your envelope contains your participation number. We will refer to it as "your private code" in the following. Your private code is the only identification used during the experiment and you will also need it to collect your cash payments.

### Cash Payments

During the experiment we shall not speak of Euros but rather of Talers. At the end of the experiment the total amount of Talers you have earned will be converted to Euros at the following rate:

### 10 Taler = 1 Euro

Cash payments can be collected from Monday onwards in room w3.29 on the third floor (West) of this building. You will present your private code to an admin staff person (Mrs. xy) and you will receive your cash payment in exchange. The admin staff person will not know who has done what and why, nor how payments were generated. No experimenter will be present in the room when you collect your money. Also, the private codes of this experiment will be mixed up with the codes of other experiments. This will again help to guarantee that the amount you earn cannot be linked to your decisions. Mrs. xy is available from Monday to Friday between 9 a.m. and 11 a.m., as well as between 2 p.m. and 3 p.m. in room w3.29 on the third floor (West) of this building. Please collect your earnings within a week (you find those details also on the card displaying your private code).

### Two Groups

Before the experiment starts, the participants in this room will be randomly divided into two groups of equal size. The groups are called Group A and Group B. Members of Group A will be seated in this room, members of Group B will be seated in the adjacent room.

### Role Assignment and Start of the Experiment

After the instructions at hand and the detailed instructions have been read aloud and all questions have been answered you (and all other participants in this room) will be asked to open the sealed envelope you drew from the box when entering this room. The envelope contains a card with your private code. The code ends with a number. If this number is even, you are a member of Group A, if it is odd, you are a member of Group B. Members of Group A are asked to take a seat at one of chairs in this room. Members of Group B will be escorted to the adjacent room and asked to take a seat there. Then the decision sheets will be handed out and the experiment starts.

*[Instructions for The Equality Equivalence Test]*

### Matching

Each member of Group A is anonymously paired with a member of Group B. The matching is 1:1; that is, each member of Group A is exactly matched with one member of Group B and vice versa. You will **never learn the identity of the member of the other group you are paired with**. In the same way, the member of the other group you are paired with will not learn your identity. In the following we call the member of the other group you are matched with **the other person**.

### Decision Tasks

Each member of Group A and each member of Group B will be asked to make 10 decisions. Each of the 10 decisions is a choice between the alternatives LEFT and RIGHT. Each alternative has consequences for your own payment and for the payment to the other person you are matched with.

**Example:** You may be asked, if you prefer alternative LEFT, in which you receive 15 Taler and the other person 30 Taler, or alternative RIGHT, in which you receive 20 Taler and the other person receives 20 Taler as well. You then have to decide which of the two alternatives to choose. This decision problem is presented as follows:

LEFT		Your Choice (please mark)		RIGHT	
you get	other person gets			you get	other person gets
15 Taler	30 Taler	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	20 Taler	20 Taler

Each member of Group A and each member of Group B will be asked to make 10 of such decisions in total. The payoffs out of this part of the experiment are determined as follows:

**Payment from your decisions:** For each participant one of the 10 decision situations is selected separately and at random, and the alternative chosen in the respective situation will then be paid out. If for example the situation above would be selected and if you had chosen in the above situation the alternative RIGHT, then you would receive 20 Taler and the other person would receive 20 Taler too.

**Payment as other person:** As the other person you are matched with receives Taler from your decision, without doing anything for it, you also receive Taler from another participant in the experiment, without doing anything for it; that is, you are for another participant the other person. We will ensure that if person X is your other person, then for sure you will not be the other person of person X.

In the case of wrong decisions (no alternative in a row or both alternatives in a row are marked) you get no payment from this part of the experiment and the payment to the other person is determined randomly (by randomly implementing one of the two alternatives).

After you have made your decision on the decision sheet, put the decision sheet into the envelope and wait until an experimenter will collect it. Pay attention that the envelope is not marked in any way.

Decision Sheet

**The decision sheet will look as shown below. Note that this page is NOT the decision sheet. The decision sheet will be handed out to you at the start of the experiment.**

Decision Sheet

**Please enter your private code (you will get no payment out of the experiment in case the private code is missing or wrong):**

.....

LEFT		Your Choice (please mark)		RIGHT	
you get	other person gets			you get	other person gets
15 Taler	30 Taler	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	20 Taler	20 Taler
19 Taler	30 Taler	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	20 Taler	20 Taler
20 Taler	30 Taler	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	20 Taler	20 Taler
21 Taler	30 Taler	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	20 Taler	20 Taler
25 Taler	30 Taler	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	20 Taler	20 Taler

LEFT		Your Choice (please mark)		RIGHT	
you get	other person gets			you get	other person gets
15 Taler	10 Taler	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	20 Taler	20 Taler
19 Taler	10 Taler	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	20 Taler	20 Taler
20 Taler	10 Taler	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	20 Taler	20 Taler
21 Taler	10 Taler	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	20 Taler	20 Taler
25 Taler	10 Taler	LEFT <input type="radio"/>	<input type="radio"/> RIGHT	20 Taler	20 Taler

[Instructions for NEU]

Decision Sheet of a Member of Group B

**If you are a member of Group B your decision sheet will look as shown below. Note that this page is NOT the decision sheet. The decision sheet will be handed out to you at the start of the experiment.**

Decision Sheet – You Are a Member of Group B

**Please enter your private code (you will get no payment out of the experiment in case the private code is missing or wrong):**

.....

**Matching**

Each member of Group B is anonymously paired with a member of Group A. The matching is 1:1; that is, each member of Group B is exactly matched with one member of Group A and vice versa. You will **never learn the identity of the member of the other group you are paired with**. In the same way, the member of the other group you are paired with will not learn your identity. In the following we call the member of the other group you are matched with **the other person**.

Before the start of the experiment we made a coin toss and **the outcome of the coin toss is HEADS**.

The other person you are paired with is **not informed** about the outcome of the coin toss but the other person is aware about the fact that you know the outcome of the coin toss.

We now ask you to send a message to the other person. You have the choice between the following two alternatives:

**Message 1:** “The outcome of the coin toss is **HEADS**.”

**Message 2:** “The outcome of the coin toss is **TAILS**.”

Your message will be shown to the other person, and then he or she will be asked to choose between **HEADS** and **TAILS**. The choice of the other person will determine your payoff and the payoff of the other person in this part of the experiment:

- If the other person chooses the actual outcome of the coin toss, **HEADS**, then **you** will receive **20 Taler** and **he or she** will receive **30 Taler**.
- If the other person chooses **TAILS**, then **you** will receive **30 Taler** and **he or she** will receive **30 Taler**.

The other person is **not aware of the payoff consequences** for both players when he or she chooses between **HEADS** or **TAILS**. The only information the other person has is the content of the message you send to him or her. Even after the decision of the other person he or she will only learn his or her own payoff from the chosen alternative.

The message I will send to the other person is (please mark a message):

Message 1: "The outcome of the coin toss is HEADS."

Message 2: "The outcome of the coin toss is TAILS."

Question Sheet of a Member of Group B

**If you are a member of Group B your question sheet will look as shown below. Note that this page is NOT the question sheet. The question sheet will be handed out to you at the start of the experiment.**

Question Sheet – You Are a Member of Group B

**Please enter your private code (you will get no payment out of the experiment in case the private code is missing or wrong):**

.....

Please answer the following question. You can earn extra money by answering the following question: The participant, whose answer is closest to the true value gets 50 Taler extra in each session.

**Question: "Out of 100 Group A members, how many do you think follow the message of the paired Group B member on average?"**

.....

After you have made your decision on the decision sheet and answered the question on the question sheet, put the decision sheet and the question sheet into the envelopes and wait until an experimenter will collect them. Pay attention that the envelopes are not marked in any way.

Decision Sheet of a Member of Group A

**If you are a member of Group A your decision sheet will look as shown below. Note that this page is NOT the decision sheet. The decision sheet will be handed out to you at the start of the experiment.**

Decision Sheet – You are a Member of Group A

**Please enter your private code (you will get no payment out of the experiment in case the private code is missing or wrong):**

.....

### Matching

Each member of Group A is anonymously paired with a member of Group B. The matching is 1:1; that is, each member of Group A is exactly matched with one member of Group B and vice versa. You will **never learn the identity of the member of the other group you are paired with**. In the same way, the member of the other group you are paired with will not learn your identity. In the following we call the member of the other group you are matched with **the other person**.

Before starting this experiment, we have made a coin toss, and told the outcome of it to the other person you are paired with, but we are not going to tell it to you. After being informed of the outcome of the coin toss the other person has sent a message to you. The other person had to choose one of the following two messages:

**Message 1:** "The outcome of the coin toss is **HEADS**."

**Message 2:** "The outcome of the coin toss is **TAILS**."

The other person has sent you the following Message: **Message** \_\_\_\_\_

Now we ask you to choose between HEADS and TAILS. Your choice determines your own payoff and the payoff of the other person in the following way:

- If your **decision is identical with the outcome of the coin toss**, you get a Taler and the other person gets b Taler.
- If your **decision is not identical with the outcome of the coin toss**, you get c Taler and the other person gets d Taler.

You do not know the payoff consequences of your decision. This means you do not know the actual value of a, b, c, and d. In contrast, the other person knows exactly the values of a, b, c, and d. The only thing you will learn after the experiment is your own payoff from your choice. You will never learn the payoff of the other person or the payoffs from the alternative not chosen.

I choose (please mark one of the two alternatives):

HEADS

TAILS

After you have made your decision on the decision sheet, put the decision sheet into the envelope and wait until an experimenter will collect it. Pay attention that the envelope is not marked in any way.

## References

- Abeler, J., Nosenzo, D., Raymond, C., 2016. Preferences for truth-telling. IZA Discussion Paper 10188. Bonn.
- Andreoni, J., Bernheim, D., 2009. Social image and the 50-50 norm: a theoretical and experimental analysis of audience effects. *Econometrica* 77, 1607–1636.
- Andreoni, J., Miller, J., 2002. Giving according to GARP: an experimental test of the consistency of preferences for altruism. *Econometrica* 70, 737–753.
- Andreoni, J., Petrie, R., 2004. Public goods experiments without confidentiality: a glimpse into fund-raising. *J. Public Econ.* 88, 1605–1623.
- Angelova, V., Regner, T., 2013. Do voluntary payments to advisors improve the quality of financial advice? an experimental deception game. *J. Econ. Behav. Organ.* 93, 205–218.
- Battigalli, P., Dufwenberg, M., 2007. Guilt in games. *Am. Econ. Rev.* 97, 170–176.
- Blanco, M., Engelmann, D., Normann, H.T., 2011. A within-subject analysis of other-regarding preferences. *Games Econ. Behav.* 72, 321–328.
- Bock, O., Nicklisch, A., Baetge, I., 2014. Hroot: registration and organization online tool. *Eur. Econ. Rev.* 71, 117–120.
- Bonroy, O., Constantatos, C., 2014. On the economics of labels: how their introduction affects the functioning of markets and the welfare of all participants. *Am. J. Agric. Econ.* 97, 239–259.
- Camerer, C., Loewenstein, G., Weber, M., 1989. The curse of knowledge in economic settings: an experimental analysis. *J. Polit. Econ.* 97, 1232–1254.
- Cappelen, A.W., Sorensen, E.O., Tungodden, B., 2013. When do we lie? *J. Econ. Behav. Organ.* 93, 258–265.
- Cox, J., 2004. How to identify trust and reciprocity. *Games Econ. Behav.* 46, 260–281.
- Cox, J., Sadiraj, C.V., 2012. Direct tests of individual preferences for efficiency and equity. *Econ. Inq.* 50, 920–931.
- Crawford, V., Sobel, P.J., 1982. Strategic information transmission. *Econometrica* 50, 1431–1451.
- Darby, M., Karni, R.E., 1973. Free competition and the optimal amount of fraud. *J. Law Econ.* 16, 67–88.
- Dreber, A., Johannesson, M., 2008. Gender differences in deception. *Econ. Letters* 99, 197–199.
- Dufwenberg, M., Dufwenberg, M., 2017. Lies in Disguise – A Theoretical Analysis of Cheating CESifo Working Paper Series, No.6208.
- Dufwenberg, M., Gneezy, U., 2000. Price competition and market concentration: an experimental study. *Int. J. Ind. Organ.* 18, 7–22.
- Dulleck, U., Kerschbamer, R., 2006. On doctors, mechanics, and computer specialists: the economics of credence goods. *J. Econ. Lit.* 44, 5–42.
- Dulleck, U., Kerschbamer, R., Sutter, M., 2011. The economics of credence goods: an experiment on the role of liability, verifiability, reputation, and competition. *Am. Econ. Rev.* 101, 526–555.
- Engelmann, D., Strobel, M., 2004. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *Am. Econ. Rev.* 94, 857–869.
- Erat, S., Gneezy, U., 2012. White lies. *Manage. Sci.* 58, 723–733.
- Etilé, F., Teyssier, S., 2016. Signaling corporate social responsibility: third-party certification versus brands. *Scand. J. Econ.* 118, 397–432.
- Feddersen, T.J., Gilligan, T.W., 2001. Saints and markets: activists and the supply of credence goods. *J. Econ. Manage. Strategy* 10, 149–171.
- Fischbacher, U., Föllmi-Heusi, F., 2013. Lies in disguise – an experimental study on cheating. *J. Eur. Econ. Assoc.* 11, 525–547.
- Fisman, R., Kariv, S., Markovits, D., 2007. Individual preferences for giving. *Am. Econ. Rev.* 97, 1858–1876.
- Gibson, R., Tanner, C., Wagner, A.F., 2013. Preferences for truthfulness: heterogeneity among and within individuals. *Am. Econ. Rev.* 103, 532–548.
- Gneezy, U., 2005. Deception: the role of consequences. *Am. Econ. Rev.* 95, 384–394.
- Gneezy, U., Kajackaite, A., Sobel, J., 2016. Lying Aversion and the Size of the Lie. SSRN Working Paper
- Gneezy, U., Rockenbach, B., Serra-Garcia, M., 2013. Measuring lying aversion. *J. Econ. Behav. Organ.* 93, 293–300.
- Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. *J. Econ. Sci. Assoc.* 1, 114–125.
- Grosskopf, B., Sarin, R., 2010. Is reputation good or bad? *Am. Econ. Rev.* 100, 2187–2204.
- Hoffman, E., McCabe, K., Shachat, K., Smith, V., 1994. Preferences, property rights, and anonymity in bargaining games. *Games Econ. Behav.* 7, 346–380.
- Huck, S., Lünser, G., Tyran, J.R., 2016. Price competition and reputation in markets for experience goods: an experimental study. *RAND J. Econ.* 47, 99–117.
- Hurkens, S., Kartik, N., 2009. Would i lie to you? On social preferences and lying aversion. *Exp. Econ.* 12, 180–192.3.
- Kerschbamer, R., 2015. The geometry of distributional preferences and a non-parametric identification approach: the equality equivalence test. *Eur. Econ. Rev.* 76, 85–10
- List, J., 2007. On the interpretation of giving in dictator games. *J. Polit. Econ.* 115, 482–493.
- Lopez-Perez, R., Spiegelman, E., 2013. Why do people tell the truth? *Exp. Econ.* 16, 233–247.
- Maggian, V., Villeval, M.C., 2016. Social preferences and lying aversion in children. *Exp. Econ.* 19, 663–685.
- Mazar, N., Amir, O., Ariely, D., 2008. The dishonesty of honest people: a theory of self-concept maintenance. *J. Market. Res.* 45, 633–644.
- Mimra, W., Rasch, A., Waibel, C., 2016. Price competition and reputation in credence goods markets: experimental evidence. *Games Econ. Behav.* 100, 337–352.
- Nelson, P., 1970. Information and consumer behavior. *J. Polit. Econ.* 78, 311–329.
- Schneider, H.S., 2012. Agency problems and reputation in expert services: evidence from auto repair. *J. Ind. Econ.* 60, 406–433.
- Sheremeta, R.M., Shields, T., 2013. Do liars believe? beliefs and other-regarding preferences in sender-receiver games. *J. Econ. Behav. Organ.* 94, 268–277.
- Sutter, M., 2009. Deception through telling the truth?! experimental evidence from individuals and teams. *Econ. J.* 119 (534), 47–60.
- Zizzo, D.J., 2010. Experimenter demand effects in economic experiments. *Exp. Econ.* 13, 75–98.