Economics Division

University of Southampton

Southampton SO17 1BJ, UK

Discussion Papers in
Economics and Econometrics

Title: Concern for Relative Standing and Deception

By : Spyros Galanis (University of Southampton) and Michael Vlassopoulos (University of Southampton)

No. 1403

Concern for Relative Standing and Deception[*]

Spyros Galanis[†]
University of Southampton

Michael Vlassopoulos[‡]
University of Southampton

January 14, 2014

Abstract

We report results from a sender-receiver cheap talk game, which explores whether an individual's decision to deceive is influenced by a concern for relative standing in a reference group. We show theoretically that positively biased senders, who think they are higher in the deception distribution than they actually are, will correct their beliefs and increase their cheating, when presented with information on the actual deception distribution. Hence, a predominantly positively biased group of senders will increase its average deception. Moreover, within a group, being more positively biased implies cheating less. The experimental data confirm both of these hypotheses.

Keywords: Deception, Lying, Sender-receiver game, Concern for Rank
JEL: C91, D03, D83

1   Introduction

The act of deceiving others in order to promote one's interests is both common and well documented. Some familiar examples are tax payers cheating on their tax returns, politicians lying about their future actions in order to get elected, insurees lying about their characteristics so that they obtain a better contract, job applicants lying about their qualifications in order to get hired, or salespeople lying about the quality of the product in order to secure a better deal. Standard economic theory treats deception as the by-product of a cost-benefit analysis that each agent performs, given his preferences. For instance, in contract and mechanism design theory, agents have a private information characteristic that they can lie about, if this will make them better off.

However, many other factors, unrelated to a direct own cost-benefit analysis, may influence the decision to deceive. For instance, Gneezy (2005) shows that people do not only care about their own gain from cheating but also take into account the losses that cheating may impose on others. Recently, several experimental studies of deception have attempted to identify the determinants of lying and to estimate individuals' intrinsic cost of lying (e.g. Gneezy et al., 2013; Gibson et al., 2013).[1]

The purpose of the present paper is to contribute to this literature, by examining whether individuals have a cost of cheating that is a function of one's perception of his relative standing in the cheating distribution of his peer group. In other words, we examine whether an individual's decision to deceive is influenced by how much he thinks his peers, who are closest to him in the distribution, will cheat. This perspective has implications for cheating behavior when someone receives information regarding the deceptive behavior of his peers that we test and confirm in an experimental setting. Our most important finding is that a basic determinant of whether deception will increase or decrease is whether subjects are positively or negatively biased, as measured by the difference between their actual and perceived decile in the distribution of deception. If a group is positively biased then information will increase average cheating, whereas if it is negatively biased, information will decrease average cheating.

The idea that many economic decisions (e.g. consumption, saving, labor supply, education, charitable giving) are influenced by concerns for relative standing compared to a reference group of peers is an old one in economics (Veblen, 1899; Duesenberry, 1949), but is also one that has received more recent attention and empirical support (Clark et al., 2008; Postlewaite, 2011a,b; Heffetz and Frank, 2011). In this paper we combine this literature with the literature that studies the determinants of deception mentioned above: we consider whether a connection exists between one's perceptions about the degree of cheating that prevails in the peer group and his own decision to cheat. A relevant example is tax compliance. Assuming that the amount of tax evaded is proportional to earned income, then a middle-class individual's decision to evade his taxes may be more influenced by his perception of how much other middle-class individuals cheat, rather than how much the very rich or the very poor cheat.

However, inferring whether individuals have a concern for relative standing by looking at whether their cheating responds to information about the deception of others is an exercise fraught with identification problems, as information may also induce other types of learning. To explain

---

[1]See also Ellingsen et al. (2009), Hurkens and Kartik (2009), Lundquist et al. (2009), Erat and Gneezy (2012), Abeler et al. (2012), Fischbacher and Follmi-Heusi (2013).

more concretely, suppose that an individual has beliefs p which describe the uncertainty of his own payoff (e.g. probability of auditing in a tax evasion environment) and beliefs q about how much his peers will cheat (e.g. amount of tax evasion among individuals with similar income). Hence, he faces a trade off between maximizing his expected payoff given p and incurring the psychological cost of cheating more than his peers, according to his beliefs q. Having concern for relative ranking means that information which changes q, will change deception. However, information about how much others cheat may also change p. For example, information about tax evasion may also update one's probability p of auditing, hence confounding the two effects. To overcome this problem, we develop a more elaborate version of the sender-receiver cheap talk game that has been used extensively in economics to study deception (Gneezy, 2005), with the added property that information about cheating of others does not influence p.[2] We achieve this by having a more finely gradated measure of deception (as opposed to the binary deception or no deception measure in Gneezy (2005)), so that learning how much others have deceived cannot help you infer their beliefs p.[3]

To address our research question we develop an experimental design consisting of two stages. In the first stage, the participants play a single-shot, cheap-talk, deception game. Player 1, the sender, ranks six possible actions, each specifying a payoff for him and player 2, the receiver. The payoffs always sum to £10. He then sends a message to the receiver with a ranking of the actions. The receiver, without knowing which action is associated with which payoff, picks a number between 1 and 6, which determines the action and therefore the payoffs of both players. The message specifies that the sender has ranked the actions according to the receiver's payoff. Hence, the sender has an incentive to deceive the receiver, in order to obtain a higher payoff. We measure cheating by the distance between the honest and the reported rankings, using a cheating score function, which maps rankings into a number between 0-9. We communicate this measure to the senders only. In the second stage, we elicit the beliefs of the senders regarding how their cheating ranks relative to other senders in their group.

We consider our measure of deception to be plausible because it has the following two properties. First, a higher deception score represents a larger departure from the honest ranking. Second, as the deception score increases, one can always find rankings with that score, which weakly increase the sender's payoff and therefore weakly decrease the receiver's payoff, given the most plausible beliefs of the senders regarding the behavior of the receivers.[4] In other words, under these beliefs our measure has the property that a higher score "cheats" the receiver more.

We define biasedness as the difference between actual and perceived (median) decile. We say that a sender is positively biased if his actual decile is higher than his perceived (median) one. Such

---

[2]Other studies that have used an experimental setting similar to that of Gneezy (2005) to study deception are Sanchez-Pages and Vorsatz (2007), Dreber and Johannesson (2008), Hurkens and Kartik (2009), Sutter (2009), Rode (2010), and Innes and Mitra (2012).

[3]We are aware of two other ways of cancelling the effect of information on updating p. The first is making p objective, as in Gibson et al. (2013). The shortcoming with this approach for our purposes is that lying is costly for the experimenter and not another subject, and this may affect the decision to lie. The second is the method suggested by Gneezy et al. (2013), where deception may harm another subject, but the extent of deception does not translate into a lower payoff for the receiver. So in terms of payoff consequences to the receiver, the sender's deception action is binary, whereas for our purposes we need to have a more finely gradated measure of deception.

[4]As in Gneezy (2005), we find that the position chosen more often by the receivers is the first one. In other words, they follow the recommendation of the sender.

a sender has underestimated the cheating of others, thinking that his score puts him in the top (e.g. decile 3), whereas in reality he is in the bottom (e.g. decile 8). If relative standing matters to him, then due to his wrong beliefs he has incurred a higher psychological cost. Hence, if we provide information that reveals to him that his initial score puts him at decile 8, then the cost he incurs drops and he re-optimizes by increasing his score.[5] A similar argument suggests that a negatively biased sender would decrease his score when presented with information about the distribution of scores. Therefore, for a group that is predominantly positively biased, providing information about the distribution of scores leads to a higher average score. Using a similar reasoning we can argue that, within the same group, being more negatively biased implies incurring a lower cost and therefore a higher cheating score. We formalize these hypotheses in section 2.3.

The experimental data confirm these two hypotheses. We find that the control group is predominantly positively biased and that the average score in the treatment group is 11% higher, a statistically significant difference. Moreover, we find that negatively biased senders have a higher average score than the positively biased ones. Finally, it is not the case that subjects react to any type of information, or that they blindly imitate the behavior of subjects in the reference group. In particular, when we only release information about the average deception score or about the scores of the highest and lowest deciles, we find no statistically significant difference in average cheating.[6]

The paper is organized as follows. In section 2 we describe the experimental setting and we formalize our hypotheses. In section 3 we present our results and in section 4 we conclude.

2   Experimental Design

Our experiment consists of two stages. In the first stage participants play a single-shot, cheap-talk, deception game in which senders transmit a message to receivers, who then take an action which determines the payoffs of both players. In this setting the sender may be tempted to deceive the receiver in order to ensure a higher own payoff. In the second stage we elicit senders' beliefs regarding how their cheating in the sender-receiver game ranks relative to other subjects.

2.1   The Sender-Receiver Deception Game

The experiment consists of two stages. In the first stage, there are two players, a sender and a receiver, and each participant is randomly and anonymously assigned one of the two roles. There are 6 outcomes, $A = \{a_1, \ldots, a_6\}$, that the receiver can choose from, and each outcome specifies an allocation of £10 among the two players. The payoffs associated with each of the 6 outcomes are (0,10), (2,8), (4,6), (6,4), (8,2) and (10,0), where the first number specifies the sender's payoff and the second number specifies the receiver's payoff, in GBP. Therefore, there is a clear ranking of the 6 outcomes for the sender if own monetary payoff is the only relevant dimension he cares about, whereas the ranking for a receiver who is only concerned about own payoff is the exact reverse.

---

[5]A crucial step in this argument is that information about the cheating distribution does not alter his beliefs p about what the receiver will do, only his beliefs q about the cheating of other senders. We elaborate on this argument in section 2.3, where we formally explain our theory.

[6]In a tax evasion experiment with objective probability of auditing, Fortin et al. (2007) also find that mean group behavior has no impact on tax evasion decision under self-consistent expectations.

The sender's ranking is communicated to the receiver through an onscreen message, displayed below in Figure 1. Effectively, the sender's message is a ranking of the monetary payoffs, made by the sender from the perspective of the receiver.

Message to player B

Player A has ranked the 6 outcomes, in terms of your payoff, from highest to lowest.

Please choose a number between 1 (representing the outcome that player A has ranked highest according to your payoff) and 6 (representing the outcome that player A has ranked lowest according to your payoff).

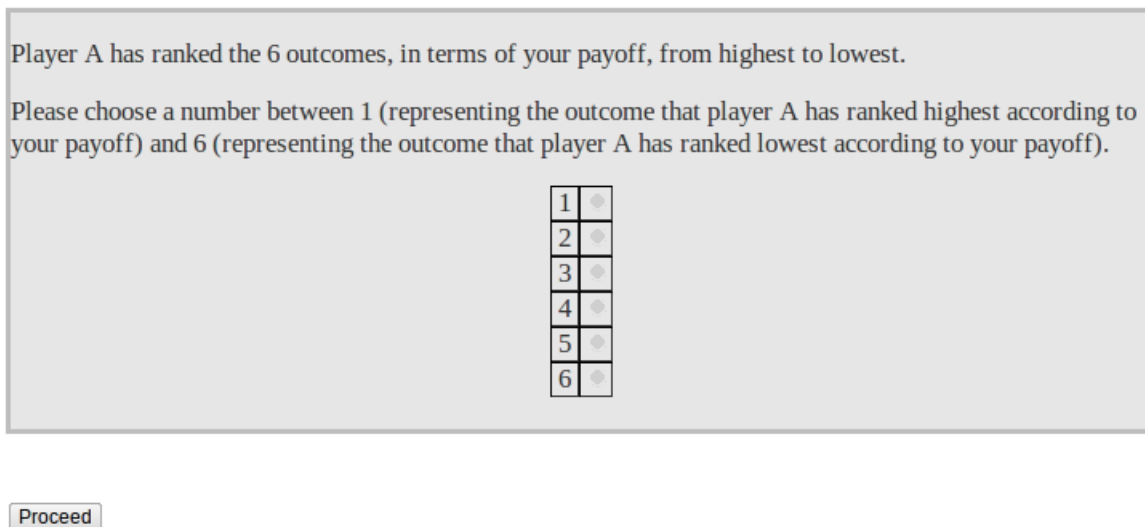| 1 | ◇ |
2	◇
3	◇
4	◇
5	◇
6	◇

Proceed

Figure 1: Screenshot of Message from Sender to Receiver

Note that the outcomes are defined such that the incentives of the sender are not aligned with those of the receiver, which means that the sender has an incentive to deceive the receiver. However, the receiver does not know this as he is never informed about the payoffs associated with each outcome. Although he is told that the sender has specified a ranking of the outcomes according to the receiver's interests, he has no way of knowing what are the available payoffs for him or for the sender. Therefore, he has no way of knowing whether his interests oppose those of the sender, or not, and the sender is also aware of this. The only payoff he learns is the one he receives after he makes his choice, without being informed about the sender's payoff.

To capture the level of deception associated with each ranking of the outcomes, we specify a score function that calculates the difference between the true ranking according to the receiver's payoff and any other ranking. We communicate this to the sender but not to the receiver. The score is calculated as follows: for each action $a \in A$, let $t(a) \in \{1, 2, \ldots 6\}$ be the true ranking of a from the perspective of the receiver. Let $r(a)$ be the reported ranking of a from the perspective of the receiver. We define the deception score of report r to be

$$s(r) = \frac{1}{2}\sum_{a \in A}|r(a) - t(a)|.$$

The honest report has a score of 0, whereas the highest deception can be achieved by reporting the sender's ranking, instead of the receiver's, which yields a score of $1/2 * (5 + 5 + 3 + 3 + 1 + 1) = 9$.

The possible scores are all integers between 0 and 9. Note that if report r has a higher score than report r', then r is further away from the true report than r' is, according to this metric.

This is because the sum of the absolute differences between the true and reported rankings for each action is higher under r than under r'. Hence, we provide senders with a set of available strategies (the reports) and a score function, which is a way of ranking these strategies in terms of how deceitful they are.

To facilitate computation and allow subjects to familiarize themselves with how rankings of the outcomes map into deception scores we supply senders with a calculator that can compute, for each reported ranking that the subject may enter, the associated deception score and we allow subjects to experiment by entering different rankings before submitting their final choice. Moreover, to ensure that subjects have understood how the deception score is computed we ask senders several comprehension questions, before we allow them to proceed to submitting their final choice.

Finally, we communicate the sender's message to the receiver (without revealing the payoffs) and ask him to choose an outcome, that will determine the payoff for both agents for this stage of the experiment. Since we are interested in the behavior of senders, to economize on the number of subjects, each receiver made a single choice, which was used to determine the payoff of multiple senders. We randomly chose one of these senders to determine the payoff of each receiver. Hence, both sender and receiver play a single-shot game and this is common knowledge.

2.2   Eliciting Perceptions of Relative Deception

In the second stage, only applicable to senders, we ask subjects a series of incentivized questions aiming to reveal their perceptions as to how deceptively they have acted relative to other subjects. In particular, we ask participants a series of choices between a bet on their relative standing in terms of their deception score and a lottery of the type displayed below in Figure 2. These questions allow us to deduce in which of the 10 deciles of the distribution of scores (top 10%, between 10% and 20%, ..., bottom 10%) each participant places his median belief.[7] Our approach of eliciting subjects' beliefs of their relative standing follows an incentive compatible mechanism suggested in the recent overconfidence literature (e.g. Benoît and Dubra, 2011).[8] We told participants that one of questions would be randomly selected at the end of the experiment to determine their payoff for this stage of the experiment.

2.3   Hypotheses

In order to formalize our hypotheses, we build a simple model where concern about one's relative standing influences the chosen score. Let A be the set of actions, with typical element $a \in A$. Let R be the set of rankings that the sender can choose from, with typical element $r \in R$. A ranking $r \in R$ is a function that maps each action $a \in A$ to a number between 1 and 6. The receiver picks a position between 1 and 6, without knowing which action corresponds to which position. Each action $a \in A$ specifies a payoff for both the sender and the receiver. For example, if the sender picks ranking r and the receiver picks position 3, then the payoffs are determined by action a such

---

[7]We deduce the exact decile for subjects who place themselves at the tails (deciles 1-3 and 7-10). The residual category is for subjects who place themselves in the middle deciles 4-7.

[8]We also asked participants what probability they assign on being in the top 10% and being in the bottom 10% of the distribution using a series of similar choices between a bet and a lottery.

Recall that your score is 6, there are 9 possible scores: 0,1,2,3,4,5,6,7,8 and 9, and a higher score indicates a larger departure from the true ranking.

Please choose your preferred option.

○ You will receive £4 if your score is in the top 30%. This means that your score is HIGHER than at least 70% of the scores of the other participants.

○ The computer will run a lottery that will award you £4 with probability 50% and £0 with probability 50%.

[Submit]

_____

If you want to read again instructions regarding how the score is calculated please [click here](#).

Figure 2: Screenshot of Second Stage Question Eliciting Perceptions of Relative Deception

that $r(a) = 3$. Let $u^i : A \to \mathbb{R}$ be sender i's instantaneous utility from action a. Note that we do not make the assumption that the sender's instantaneous utility is necessarily monotonic with respect to his own payoff.

We postulate that sender i's utility function, U^i, has two components. The first is a function of his chosen ranking, r, and of his probabilistic beliefs about which position the receiver will pick. Beliefs are modelled by having a probability distribution over the set of all positions, $p \in \Delta\{1, \ldots, 6\}$, where ΔK denotes the set of all probability distributions over a set K. We write this first component as $\sum_{a \in A} p(r(a))u^i(a)$.

The second component depends on the sender's perceived relative standing. As explained in section 2.1, each ranking r generates a score $s(r)$. If we order the scores by all senders in the reference group, from highest to lowest (breaking ties by randomizing), we can create 10 deciles, where decile 1 contains the scores that are in the top 10% of the distribution, decile 2 contains the scores that are between the top 20% and top 10% of the distribution, and so on. Hence, the scores in decile 1 are those exhibiting the highest cheating, whereas those in decile 10 exhibit the lowest cheating, relative to the reference group.

Let $D = \{1, \ldots, 10\}$ be the set of all deciles. We postulate that each sender i incurs an instantaneous cost $C^i : D \to \mathbb{R}$ by choosing a score that places him in decile $d \in D$. As d increases (hence moving closer to the "low cheating" decile 10), cost decreases. We allow senders to have different cost functions C^i, so that some care more about being closer to the low cheating decile 10 than others.

The sender does not know the scores of the other senders when choosing a ranking r, which generates score $s(r)$. He therefore has probabilistic beliefs over the set of all deciles. To simplify the model, we assume that the sender cares only about his median belief decile, $d^i_m(s(r))$, which is the decile $d \in D$, such that he assigns at least 50% probability that his actual decile is at least d and at least 50% probability that his actual decile is at most d, if he chooses score $s(r)$. A higher score, holding the scores of the other senders constant, implies a lower (i.e. closer to 1) median decile. Hence, d^i_m is a decreasing function of score. Summarizing, sender i's utility function is the

6

following:

$$U^i(r,p) = \sum_{a \in A} p(r(a))u^i(a) - C^i(d_m^i(s(r))).$$

The sender chooses a ranking r that maximises U^i, given his beliefs p about what the receiver will do and his median beliefs function, d_m^i, about what the other senders will do. Let $R(s) = \{r \in R : s(r) = s\}$ be the set of rankings that achieve a score of s. We can rewrite U^i as a function of the score s and beliefs p: $U^i(s,p) = V^i(s,p) - C^i(s)$, where $V^i(s,p) = \max_{r \in R(s)} \sum_{a \in A} p(r(a))u^i(a)$. That is, for each score, which determines i's cost, the sender picks ranking r that maximizes the value of the first component, among all rankings generating the same score.[9] For ease of exposition, we assume that V^i and C^i are defined on the compact interval of scores $[0, 9]$. The two components of $U^i(s,p)$ are depicted in Figure 3, for fixed p.

Note that C^i depends on i's beliefs q about what the other senders will choose *and* on the intrinsic cost of being on decile d. We separate these two components by postulating the following functional form: $C^i(s) = a^i C_0^i(s)$, where $a^i > 0$ and $C_0^i(s)$ represents the intrinsic cost, with $C_0^i(0) = 0$ and $\frac{\partial C_0^i}{\partial s} \geq 0$.

It is worthwhile to note that we do not make any assumptions about the relationship between V^i and score. In that way, we can accommodate many different types of senders. For instance, a sender who thinks that the receiver will trust him, so that he believes that it is more probable that the receiver will choose position x over $x + 1$, for each $x = 1, \ldots, 5$, can be modelled by having an increasing V^i with respect to s. This is because higher scores can be generated by rankings that place the high sender's payoffs to the low positions, such as $x = 1, 2, 3$. Alternatively, a sender who thinks that the receiver does not trust him, so that he believes that it is more probable that the receiver will choose $x + 1$ over x, can be modelled by a decreasing V^i with respect to s. More general beliefs p can be depicted by a non-monotonic V^i, as in Figure 3.

The general form of V^i also allows for senders who care about types of deceptions different than the one we propose. For example, consider a sender who believes with probability one that the receiver will pick the action he ranked in position 6 and thinks that rankings which differ only with respect to positions 1 to 5 are equally deceptive, as they "hurt" the receiver equally. Such a sender is indifferent (according to the first component) between two rankings that assign the same action to position 6, and therefore $V^i(s,p)$ depends only on what rankings $r \in R(s)$ specify in position 6.

In the experiment we extract $d_m^i(s)$ using an incentivised method. We also compute the sender's actual decile, $d_a^i(s)$. We say that sender i is *positively biased* if he has chosen a ranking that generates score s and $d_a^i(s) - d_m^i(s) > 0$, *negatively biased* if $d_a^i(s) - d_m^i(s) < 0$ and *unbiased* if $d_a^i(s) - d_m^i(s) = 0$. A positively biased sender has overestimated his relative position, thinking that his chosen score places him higher in the distribution (e.g. decile 3), than in reality (e.g. decile 8).

Suppose that all senders in a reference group are presented with the actual distribution of scores of another, similar group of senders. This information can be used by each sender as a signal, in

---

[9]V^i is well defined if R is finite and, more generally, if $R(s)$ is compact for all s and $\sum_{a \in A} p(r(a))u^i(a)$ is continuous in R.
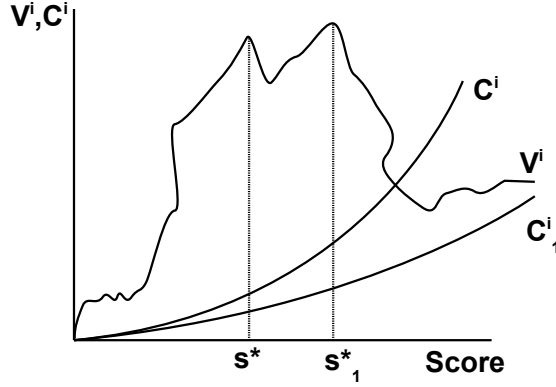
Figure 3: Optimal Score

order to update his median beliefs about the mapping between scores and deciles. Recall that the sender cares about his relative standing with respect to members of his own reference group, who receive the same information as him. Therefore, when forming his median beliefs about the mapping between scores and deciles, he also needs to incorporate the response to the same information of the other senders in his own reference group.

Suppose that sender i is positively biased and his chosen score is s^* (Figure 3). After receiving the information about another reference group, he realizes that he has overestimated his relative position. For example, if he thought that s^* placed him in decile 4, now he realizes that it places him in decile 7. In order to update his beliefs, the parameter of biasedness decreases to a_1^i, resulting in a downward rotation of the C^i function. Let C_1^i be the new cost function. Moreover, $h^i(s) = C^i(s) - C_1^i(s) = (a^i - a_1^i)C_0^i(s)$ is an increasing function of s. Because V^i is unchanged, this implies that sender i will weakly increase his score.[10] A similar argument shows that if the sender is negatively biased, he will weakly decrease his score.

The change in score described above does not take into account that other senders will also change their score, after receiving the same information. If sender i can formulate second order beliefs, he will realize that each positively biased sender will weakly increase his score, whereas each negatively biased sender will weakly decrease his score. If sender i believes that the positively biased senders are more than the negatively biased ones (as we find in the control group), then a_1^i will decrease further, implying that the cost function will rotate downwards again and i will choose an even higher score. Increasingly sophisticated senders can compute higher order beliefs. We reach an equilibrium if all agents are unbiased when presented with the distribution of scores of all other senders.

If the senders are not sophisticated enough, they will not be able to compute their higher order beliefs, and as a result an equilibrium may not be reached. In fact, as shown by Nagel (1995) and other papers, subjects usually cannot compute higher than third order beliefs, so we expect that

---

[10]Recall that at s^* we have the highest difference between V^i and C^i. Because the difference between C^i and C_1^i is increasing with s, it must be that the highest difference between V^i and C_1^i must be at $s_1^* \geq s^*$.

in an experimental setting, not all senders will be unbiased after receiving the same information. We assume that the ability to compute higher order beliefs and therefore to correctly change one's score is independent of whether a sender is positively or negatively biased.

In the argument above we implicitly assume that information about the scores does not change the beliefs p or the functional form $V^i(s,p)$. Hence, the first component of U^i does not change. There are two reasons why we make such an assumption. The first is that the senders choose their score before learning the response of the receivers, and they only play the game once. Hence, the score distribution of other senders does not convey any information about how to better play the game.

The second reason is that even if a sender wanted to use the distribution of scores in order to extract the beliefs of the other senders and update his own beliefs, this is almost impossible. The reason is that each score corresponds to many rankings, consistent with very diverse beliefs p. For example, a score of 5 or 6 can be generated by rankings which are optimal for a self-interested sender who does not care about his relative position and attaches probability 1 to the receiver picking position x, where $x = 1, \ldots, 6$. That is, degenerate beliefs about all possible positions are consistent with a score of 5 or 6. A score of 7 is consistent with position $x = 1, \ldots, 5$, a score of 8 is consistent with position $x = 1, \ldots, 4$, and so on. If we allow for mixed beliefs, risk aversion or concern for relative standing, then each belief p is consistent with many more scores. Hence, knowing someone's score cannot help discovering his beliefs p.

This can be contrasted with the setting of Gneezy (2005), where knowing a sender's action (lie or not lie) with the above characteristics implies knowing his beliefs p. That is, providing information about the behavior of other senders may transmit information about what their beliefs are, and hence influence the sender's decision, even if he does not care of the deception of others. In the current setting this way of transmitting information about beliefs is excluded.

We say that a group of senders is positively biased if the majority of senders are positively biased and the average absolute value of biasedness is the same across the negative and the positively biased senders.[11] The definition of a negatively biased group is similar. We treat positively and negatively biased senders symmetrically. In particular, we assume that two senders who are identical in everything except that one is positively and the other is negatively biased (but the absolute value of their biasedness is the same) will correct their score by the same amount in absolute values. We are now ready to formulate our hypotheses.

Hypothesis 1. *Suppose that we present the distribution of scores of a group of senders to another group whose members have not made their choice yet. If the group is positively biased, then average cheating will increase, when compared to an otherwise identical group, where this information is not presented. If the group is negatively biased, then average cheating will decrease.*

For our next hypothesis, we assume that the difference between actual and perceived decile is monotonic in score. That is, if $s > s'$ then $d_a^i(s) - d_m^i(s) \leq d_a^i(s') - d_m^i(s')$. We say that sender i is more positively biased than sender j if i has chosen a ranking with score s^i, j has chosen a ranking with score s^j and $d_a^i(s^i) - d_m^i(s^i) > d_a^j(s^j) - d_m^j(s^j)$.

Suppose that two senders have the same beliefs p, instantaneous utility u and intrinsic cost

---

[11]This is in fact what we find in the experimental data reported below.

of being in a decile, C_0, but i is more positively biased than j. We will show that $s^i \leq s^j$. Suppose by contradiction that $s^i > s^j$. Because the difference between actual and perceived decile is monotonic in score, we have $d_a^i(s^j) - d_m^i(s^j) \geq d_a^i(s^i) - d_m^i(s^i)$, which implies that $d_a^i(s^j) - d_m^i(s^j) > d_a^j(s^j) - d_m^j(s^j)$. Ignoring tie breakings, having the same score will generate the same actual decile, hence $d_a^i(s^j) = d_a^j(s^j)$.[12] This implies that $d_m^i(s^j) < d_m^j(s^j)$, therefore $a^i > a^j$, as i overestimates the decile generated by score s^j, relative to j. But this means that C^i is obtained by rotating C^j to the left and, using the same argument as in Hypothesis 1, we have $s^i \leq s^j$, a contradiction. Therefore, if i is more positively biased than j then $s^i \leq s^j$. We formulate our hypothesis below.

Hypothesis 2. *Being more positively biased implies a lower score. Moreover, positively biased senders have a lower average score than negatively biased senders.*

2.4 Treatments

In addition to the control, our experimental design involves three treatments. In all treatments, subjects play the same sender-receiver game, with the only difference that we show some extra information to the senders, before they submit their ranking of the actions. Moreover, we tell them that we have communicated this information to all other senders taking part in this treatment. Therefore, this information is common knowledge.

In the first treatment, which we refer to as the *score* treatment, we show them the distribution of the scores from the control group. In particular, we tell subjects what proportion of the control group chose a score of 0, what proportion chose a score of 1 and so on. In a second treatment, that we refer to as the *mean* treatment, we only show them the mean of the distribution of scores of the control group. Finally, in the third treatment that we term *tails* treatment, we show them the percentage of participants choosing score 0 and the percentage of participants choosing score 9, in the control group. The aim of these last two treatments is to decompose the information about the distribution of cheating score, in order to understand which pieces of information are more relevant to the subjects.

2.5 Procedures

The experiment was conducted online with University of Southampton students in the fall of 2012 and winter of 2013. Participants were recruited through email announcements. Participants who expressed interest in the experiment received login information (username, password and url). All further experimental instructions, to be found in the appendix, were provided to participants upon login into the experimental website. The experiment lasted 20-25 minutes and participants were paid a participation fee of £3 and had the opportunity to earn additional money. Total average payment was £9.25. Payments were carried out in cash by a research assistant.

A total of 432 participants (403 senders and 29 receivers) of diverse academic backgrounds participated,[13] with the sample being balanced in the gender dimension (51% males and 49%

---

[12]Note that in this general model we have assumed a compact interval of scores $[0, 9]$, so the probability of two agents have exactly the same score is zero.

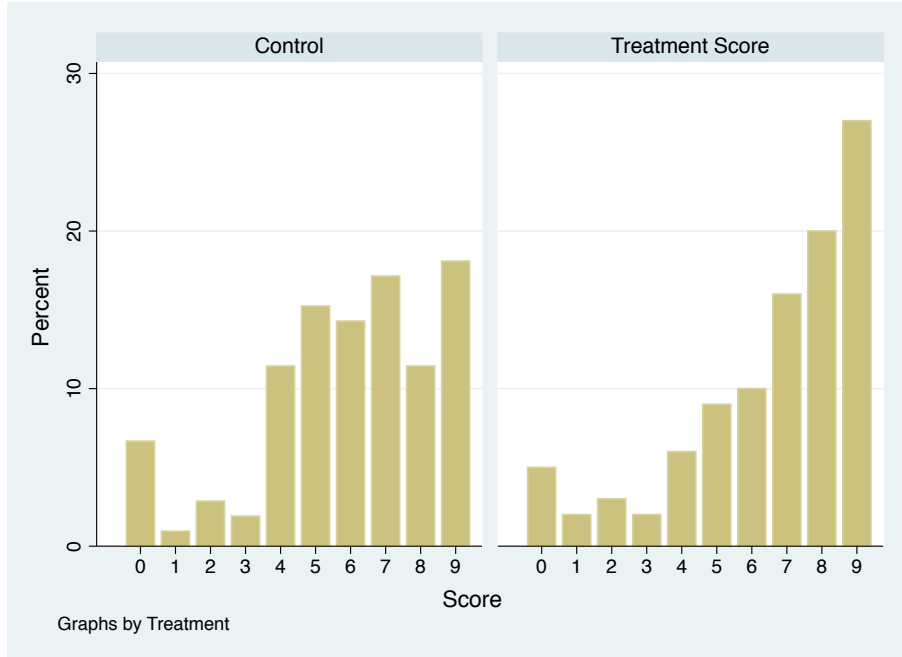[13]34% study Economics or Management.

Figure 4: Probability Distribution Function

females). The distribution of participants across treatments is as follows: 105 in the control, 100 in the *score* treatment, 99 in the *mean* treatment and 99 in the *tails* treatment.

3   Results

3.1   Main Results

We begin the presentation of our experimental results by discussing the behavior of senders in the first stage of the experiment. Figure 4 illustrates the probability distribution of the deception score separately for the control and the score treatment. We find that in the control condition the highest score (9) is the most popular choice as it is chosen by roughly 20% of subjects. Each of the numbers between 5 and 8 are chosen by at least 10% of subjects, while it is perhaps remarkable that almost 7% of subjects chose the truthful ranking which yields a score of 0.[14] When we turn attention to the score treatment, we see a noticeable shift of weight toward the highest scores 8 and 9 which combined are now chosen by almost half of the subjects (as opposed to 30% in the control).

We next consider the degree of biasedness of the senders. The results are summarized in Table 1. Recall from section 2.3 that we define a sender's bias to be the difference between the actual decile of the chosen score and the perceived decile, which we deduce in stage 2 of the experiment.[15] Note

---

[14]Almost 70% of receivers chose an action ranked in top 3 by a sender. In particular, 34.5% chose the highest ranked action, 24.1% the second highest, 10.3% the third highest, 17.2% the fourth highest, 3.5% the fifth and finally 10.3% the sixth highest.

[15]In order to determine the actual deciles, we ranked all senders according to their score, breaking ties by random-

that as we do not deduce where exactly in deciles 4-7 a subject belongs to we apply the following assignment protocol for subjects whose perceived decile is between 4 and 7: we assign zero bias if their actual decile also lies in the interval 4-7; when the actual decile is less than 4 we assign the difference between actual and 4; when the actual is greater than 7 we assign the difference between actual and 7.[16]

As can be seen in Table 1 for the control group, there are more positively than negatively biased subjects. In particular, 46 subjects (43.8%) are positively biased, whereas 33 are negatively biased (31.4%). The average bias, across all subjects in the control group, is 0.2. Moreover, in absolute value the biasedness of the positively biased is statistically indistinguishable from those that are negatively biased (Mann-Whitney two-sided test; p-value=0.26). Therefore, according to our theory, the control group fits the definition of a positively biased group.

	Control			Score Treatment		
	Freq.	Percent	Mean Score	Freq.	Percent	Mean Score
Negatively Biased	33	31.4	7.5	25	25	8.3
Unbiased	26	24.8	5.9	23	23	5.9
Positively Biased	46	43.8	4.8	52	52	6.1
All	105		5.9	100		6.6

Table 1: Summary Statistics: Bias and Deception Score

We are now ready to examine whether *Hypothesis 1* holds. According to it we expect that since the control is positively biased the average score in the score treatment will be higher than in the control group. A Mann-Whitney test indicates significant difference in the distribution of scores between the control and the score treatment (p=0.024, two-sided test). This is confirmed by regression analysis reported in Table 2. In particular, in columns (1) and (2) we report OLS regressions of the deception score on a dummy for being in the score treatment with and without a control for gender. We find that the mean deception score in the control group is 5.94, whereas in the score treatment the mean deception score is significantly higher by 11%. Columns (4) and (5) report OLS regressions of the probability of selecting a score greater than 6. These regressions also indicate that subjects in the score treatment are more likely to choose a deception score that is higher than 6.

---

izing. We then assigned decile 1 to the top 10%, and similarly for the rest.

[16]We also tried an alternative assignment protocol, whereby we assigned the midpoint 5.5 to those whose perceived decile is between 4 and 7. We then took differences between their actual decile and 5.5. This protocol produces a more positive bias.

	Deception Score			Score ≥ 6	
	(1)	(2)	(3)	(4)	(5)
Constant	5.943***	5.656***	6.014***	0.61***	0.576***
	(0.24)	(0.30)	(0.21)	(0.05)	(0.05)
Score	0.637*	0.686*	0.711**	0.12*	0.126*
	(0.35)	(0.35)	(0.32)	(0.065)	(0.045)
Female		0.579			0.067
		(0.36)			(0.065)
Bias			-0.354***		
			(0.03)		
Obs.			205		

Notes: *,** and *** denote, respectively, significance at the 10%, 5% and 1% levels. Heteroskedasticity-robust standard errors are in parentheses.

Table 2: OLS regressions.

The difference can also be seen in Figure 5 which shows that the cumulative distribution function of the score treatment first order stochastically dominates that of the control.
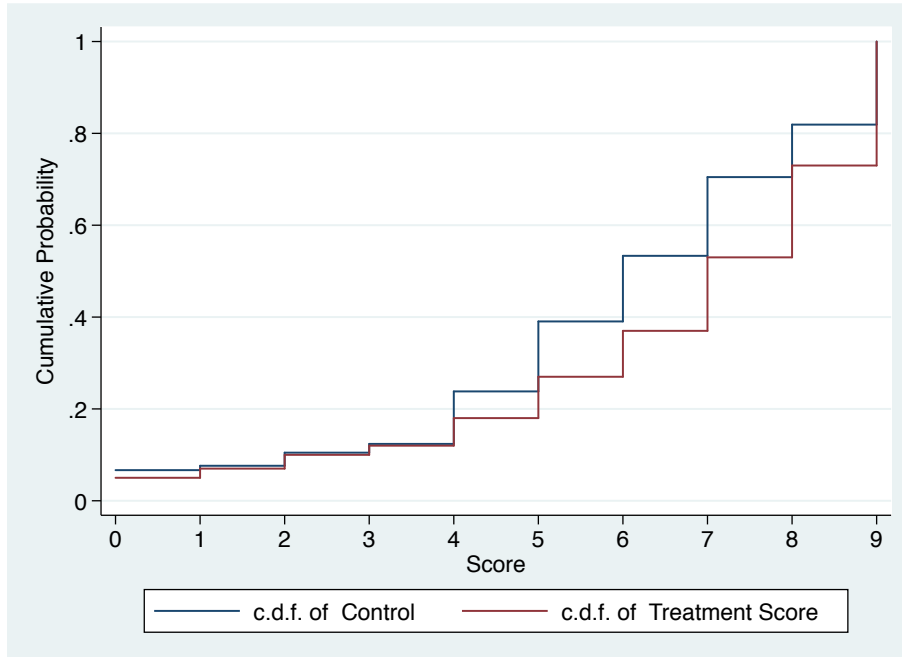

Figure 5: Cumulative Distribution Function

We next examine whether *Hypothesis 2* holds in our data. A first confirmation of this hypothesis

is provided in column (3) of Table 2, which indicates that there is a negative and significant association between the deception score chosen and the "size" of a subject's bias. This can be further seen by comparing the distribution of scores of positively biased subjects to that of negatively biased (averages are displayed in Table 1). A Mann-Whitney test strongly rejects equality for both the control and score treatment.

3.2 Additional Results

In this subsection we present some results from the two auxiliary treatments mean and tails. In Table 3 we present regressions of score in columns (1) and (2) and of the probability of selecting a score greater than 6. In all cases, we see no evidence of an effect of having seen the mean score or the tails of the score distribution of the control. This is corroborated by a Mann-Whitney test (p-value=0.11 and 0.71, respectively, for pairwise comparison between control and mean and control and tails treatment).

	Deception Score		Score ≥ 6	
	(1)	(2)	(3)	(4)
Constant	5.943***	5.869***	0.61***	0.624***
	(0.24)	(0.28)	(0.05)	(0.06)
Mean	0.491	0.480	0.087	0.085
	(0.34)	(0.35)	(0.07)	(0.07)
Tails	0.047	0.040	0.047	0.049
	(0.36)	(0.36)	(0.07)	(0.07)
Female		0.149		−0.030
		(0.29)		(0.06)
Obs.	303			

Notes: *,** and *** denote, respectively, significance at the 10%, 5% and 1% levels. Heteroskedasticity-robust standard errors are in parentheses.

Table 3: OLS regressions: Treatments Mean and Tails.

It is important to note that these auxiliary results are consistent with our theory of concern for relative standing influencing deception. Observing the mean of the control distribution does not alter one's perception of his relative standing. Hence, we do not expect that the mean deception score of those in the mean treatment will be different from that of the control. The same is true if one observes the percentage of participants choosing score 0 or score 9. This information can be used to calculate the score of those in the top decile 1 and in the bottom decile 10. However, if a subject's median decile is not one of these two deciles, then this information will not change his perception and therefore his score will stay the same. Therefore, on the aggregate, we should

not expect a significant difference in the average mean score of those in the tails treatment and the control.

3.3 Discussion

In this section, we make three points regarding our experimental findings. First, the fact that there are differences in deception when subjects receive information regarding the whole distribution as opposed to only the average or the tails suggests that we do not just find that informing people that there is more deception than they thought leads them to deceive more. The experimental results indicate that what matters is information regarding their relative deception, which is only revealed in the score treatment.

Second, a few recent experimental studies in economics and social psychology (Innes and Mitra, 2012; Gino et al., 2009; Fosgaard et al., 2013) have demonstrated that exposure to other people's deception can influence one's own deceptive behavior, due to concerns for social conformity (Cialdini and Goldstein, 2004; Ayal and Gino, 2011). Our evidence also suggests that social conformity is present. Our contribution is to illustrate the relevance of a specific type of social conformity. In particular, our mechanism suggests that each subject conforms to the cheating exhibited by the peers closest to him in the distribution, and not necessarily to the prevalent behavior in the group. In other words, social conformity occurs with respect to a sub-group, defined by one's preferred position in the deception distribution.

Finally, as discussed in the introduction and in section 2.3, providing information about the cheating of others may not only change one's perception of his relative standing, but also update his beliefs p about other parameters of the game. In that case, change in the cheating behavior cannot be solely attributed to a concern for relative standing, as the two effects are confounded. For example, information about the deception scores of others may reveal what are their beliefs regarding what the receivers will do, which could be used by the senders of the treatment group in order to update their own beliefs. We believe this is unlikely to be driving behavior in our experiment for two reasons. First, the senders of the control group play the game only once and we record their cheating score before they learn their payoff. Hence, their choice is not the outcome of experience play. Second, we specifically designed the game in a way such that each cheating score is associated with many different rankings of the possible distributions of the £10. Therefore, each score is consistent with many different beliefs of the sender about what the receiver will do. This means that receiving information about the cheating score of a sender does not reveal what his beliefs are about what the receiver will do.[17] This is to be contrasted with the setting in Gneezy (2005) and subsequent papers, where there are only two actions (cheat or no cheat), so providing information about the choice of a sender can reveal what his beliefs are as well.

4 Conclusion

We postulate that individuals have a cost of lying that is a function of their relative standing in the deception distribution of their peer group. We confirm such a relationship between deception

---

[17]We elaborated on this argument in section 2.3.

and relative standing experimentally. In particular, we postulate that a basic determinant of whether deception will increase or decrease is whether subjects are positively or negatively biased, as measured by the difference between their actual and perceived decile in the distribution of deception. Our theory suggests that being positively biased means overestimating the cost of deception. Therefore, receiving information that corrects this overestimation will lead to more cheating. Alternatively, ceteris paribus, being more positively biased means that you cheat less. We confirm both of these hypotheses. Indeed, when we released information to a positively biased treatment group about the deception scores of the control group, this led to a significant increase in cheating. Moreover, we found that, within a group, being more positively biased implied cheating less.

An area where our findings may have interesting policy implications is tax compliance. For instance, according to Leicester et al. (2012), in the UK, the gap between collected tax revenues and those that were due based on tax law was almost 8 per cent, or around £35 billion, in 2009-10. In other countries the tax gap can be higher. Leicester et al. (2012) cites several behavioral elements that are relevant for tax compliance, like overestimation of the likelihood of detection, tax complexity, or social factors, such as information about whether others pay their taxes. The message of this paper is that information about tax compliance of others can either increase or decrease tax compliance, depending on whether the target group is positively or negatively biased. Therefore, before releasing such information, one should first estimate the biasedness of the group.

Finally, having established that biased beliefs about one's relative standing influences cheating, an interesting question that arises is how to model the formation of biased beliefs and, more importantly, how these beliefs are updated with the arrival of new information. Since all peers in the group will change their beliefs and their actions in response to the new information, a proper model has to take into account how higher order beliefs are formulated. We leave these questions for future research.

References

Abeler, J., **Becker, A.**, and **Falk, A.** (2012). Truth-telling: A representative assessment. *IZA Discussion Paper*

Ayal, S. and **Gino, F.** (2011). Honest rationales for dishonest behavior. In M. Mikulincer and P.R. Shaver, editors, *The Social Psychology of Morality: Exploring the Causes of Good and Evil*, chapter 8, 149–166. American Psychological Association, Washington, DC

Benoît, J.P. and **Dubra, J.** (2011). Apparent overconfidence. *Econometrica*, 79(5): 1591–1625

Cialdini, R.B. and **Goldstein, N.J.** (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, 5: 591–621

Clark, A.E., **Frijters, P.**, and **Shields, M.A.** (2008). Relative income, happiness, and utility: An explanation for the Easterlin paradox and other puzzles. *Journal of Economic Literature*, 46(1): 95–144

Dreber, A. and **Johannesson, M.** (2008). Gender differences in deception. *Economics Letters*, 99(1): 197–199

Duesenberry, J. (1949). *Income, saving, and the theory of consumer behavior.* Harvard economic studies. Harvard University Press

Ellingsen, T., **Johannesson, M.**, **Lilja, J.**, and **Zetterqvist, H.** (2009). Trust and truth. *Economic Journal*, 119(534): 252–276

Erat, S. and **Gneezy, U.** (2012). White lies. *Management Science*, 58(4): 723–733

Fischbacher, U. and **Follmi-Heusi, F.** (2013). Lies in disguise - an experimental study on cheating. *Journal of the European Economic Association*, 11(3): 525–547

Fortin, B., **Lacroix, G.**, and **Villeval, M.C.** (2007). Tax evasion and social interactions. *Journal of Public Economics*, 91(11-12): 2089 – 2112

Fosgaard, T.R., **Hansen, L.G.**, and **Piovesan, M.** (2013). Separating Will from Grace: An experiment on conformity and awareness in cheating. *Journal of Economic Behavior & Organization, forthcoming*

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

Gino, F., **Ayal, S.**, and **Ariely, D.** (2009). Contagion and differentiation in unethical behavior: The effect of one bad apple on the barrel. *Psychological Science*, 20(3)

Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, 95(1): 384–394

Gneezy, U., **Rockenbach, B.**, and **Serra-Garcia, M.** (2013). Measuring lying aversion. *Journal of Economic Behavior & Organization, forthcoming*

Heffetz, O. and **Frank, R.H.** (2011). Preferences for status: Evidence and economic implications. In A.B. J. Benhabib and M. Jackson, editors, *Handbook of Social Economics.* North Holland

Hurkens, S. and **Kartik, N.** (2009). Would i lie to you? On social preferences and lying aversion. *Experimental Economics*, 12(2): 180–192

Innes, R. and **Mitra, A.** (2012). Is dishonesty contagious? *Economic Inquiry*, 51(1): 722–734

Leicester, A., **Levell, P.**, and **Rasul, I.** (2012). Tax and benefit policy: insights from behavioural economics. Technical report, Institute for Fiscal Studies

Lundquist, T., **Ellingsen, T.**, **Gribbe, E.**, and **Johannesson, M.** (2009). The aversion to lying. *Journal of Economic Behavior & Organization*, 70(1-2): 81–92

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85: 1313–1326

Postlewaite, A. (2011a). Social norms and preferences. In A.B. J. Benhabib and M. Jackson, editors, *Handbook of Social Economics.* North Holland

Postlewaite, A. (2011b). Social norms and social assets. *Annual Review of Economics*, 3: 239–259

17

Rode, J. (2010). Truth and trust in communication: Experiments on the effect of a competitive context. *Games and Economic Behavior*, 68(1): 325 – 338

Sanchez-Pages, S. and **Vorsatz, M.** (2007). An experimental study of truth-telling in a sender-receiver game. *Games and Economic Behavior*, 61(1): 86–112

Sutter, M. (2009). Deception through telling the truth? Experimental evidence from individuals and teams. *Economic Journal*, 119(534): 47–60

Veblen, T. (1899). *Theory of the Leisure Class.* MacMillan: New York

Appendix

Experimental Instructions for Senders

Screen 1: Log in

Screen 2: General Information

You are taking part in an economics experiment, the purpose of which is to examine the decisions people make in certain circumstances. You will be paid a participation fee of £3 for completing the experiment. You will also have the opportunity to earn additional money.

The experiment will consist of two stages and will last about 20-25 minutes. In the first stage we will ask you to make a certain decision. Your earning for this stage will depend on your decision and the decision of another participant. In the second stage we will ask you some questions regarding the decision you made in the first stage. Further details and instructions will be provided at the beginning of each stage.

The money that you earn will be paid to you in cash. In particular, after all participants have completed the experiment you will receive an email from us with instructions regarding the specific times and the location at Highfield Campus to pick up your money from. Note that an assistant will pay you privately.

Screen 3: Consent Form

Screen 4: Instructions for Stage 1

In this stage of the experiment, you are randomly matched with another participant. You will be called player A and the participant you are matched with will be called player B. The roles have been randomly allocated. You will not know with whom you were matched, neither now nor after the experiment, and we will not reveal your identity to the participant you are matched with either.

There are 6 outcomes. Each outcome consists of two numbers. The first number indicates a payoff in pounds for player A and the second number indicates a payoff in pounds for player B. These 6 outcomes are as follows:

Player A	Player B
0	10
2	8
4	6
6	4
8	2
10	0

We will ask you to rank these outcomes from 1 to 6. Then, player B will choose a number from 1 to 6, without knowing what the outcomes are. Player B's choice will determine your payments. For example, if you rank outcome (4,6) in position 3 and player B chooses number

3, then you will get £4 and he or she will get £6. If you rank outcome (6,4) in position 4 and player B chooses number 4, then you will get £6 and he or she will get £4.

Before making a choice, player B will receive a message from you, saying that you have ranked the outcomes, from highest to lowest, according to player B's payoff.

Please click below to see what the message will look like.

Screen 5: Message to Player B

Player A has ranked the 6 outcomes, in terms of your payoff, from highest to lowest.

Please choose a number between 1 (representing the outcome that player A has ranked highest according to your payoff) and 6 (representing the outcome that player A has ranked lowest according to your payoff).

Screen 6: Instructions for Stage 1

Note that in the message, player B is told that you have ranked the outcomes, from highest to lowest, according to player B's payoff. However, the message does not display the payoffs associated with each of the 6 outcomes. Therefore, player B can never know whether your ranking is the true one or not.

If you want to see the message sent to player B again, click here.

Screen 7: Payment for Stage 1

Recall that the choice made by player B will determine the payments that you and player B will receive in this stage of the experiment. For example, if you rank outcome (4,6) in position 3 and player B chooses number 3, then you will get £4 and he or she will get £6. If you rank outcome (6,4) in position 4 and player B chooses number 4, then you will get £6 and he or she will get £4. To repeat, when player B is making a choice he or she cannot see any of the payoffs.

If you want to see the message sent to player B again, click here.

Screen 8: Score

We have constructed a formula that calculates the difference between the true ranking according to player B's payoff and any other ranking. In what follows we explain how this formula works.

First, note that the true ranking of outcomes according to player B's payoff is the following:

Player A	Player B	True Ranking
0	10	1
2	8	2
4	6	3
6	4	4
8	2	5
10	0	6

The score for reporting this ranking is 0. Any other ranking places some of the outcomes at a position that does not correspond to their true rank according to player B's payoff. For example, the true rank of outcome (10,0) is 6. If the reported rank of (10,0) is 1, then the absolute difference between true and reported rank is 5. For another example, the true rank of outcome (2,8) is 2. If the reported rank of (2,8) is 4, then the absolute difference between true and reported rank is 2.

The score is calculated by adding the absolute difference between true and reported rank, for all six outcomes, and then dividing by 2.

Below is an example, which shows how to calculate the score for the following hypothetical ranking.

Player A	Player B	True Ranking	Hypothetical Ranking	Absolute Difference
0	10	1	2	$|1 - 2| = 1$
2	8	2	6	$|2 - 6| = 4$
4	6	3	5	$|3 - 5| = 2$
6	4	4	4	$|4 - 4| = 0$
8	2	5	1	$|5 - 1| = 4$
10	0	6	3	$|6 - 3| = 3$
			Total	14
			Score	$14/2 = 7$

Each possible ranking is associated with one of 10 possible scores: 0,1,2,3,4,5,6,7,8 and 9. The higher is the score, the larger is the deviation of the reported ranking from the true one.

Screen 9: Quiz

To ensure that you understand the setup we would like you to answer some questions regarding the following hypothetical scenario. If you want to read again the instructions regarding how the score is calculated please click here.

Suppose that a hypothetical player A has ranked the outcomes in the following way:

Player A	Player B	True Ranking	Hypothetical Ranking
0	10	1	3
2	8	2	4
4	6	3	1
6	4	4	2
8	2	5	5
10	0	6	6

Then the matched hypothetical player B will receive the following message:

"Player A has ranked the 6 outcomes, in terms of your payoff, from highest to lowest.

Please choose a number between 1 (representing the outcome that player A has ranked highest according to your payoff) and 6 (representing the outcome that player A has ranked lowest according to your payoff)."

Please answer the following questions:

1. What is the score of the ranking reported by player A?

2. Suppose that player B chooses option 2 in the screen. What is the payment that player A will receive?

3. Suppose that player B chooses option 5 in the screen. What is the payment that player B will receive?

Screen shown only to Score Treatment

We have previously run this experiment with a group of 106 student participants from the University of Southampton. Before we ask you to provide your ranking we present below some information regarding the scores chosen by this group of participants.

In particular, the table below presents the following piece of information. For each of the 10 possible scores, we report the percentage of participants who have chosen this particular score.

Screen 10: Choice

Please select your ranking by inserting a number from 1 to 6 next to each outcome. Your ranking and Player B's choice will determine your payment in this stage of the experiment.

Note that player B will have no information about the available outcomes and the associated payoffs. Therefore, player B can never know whether the ranking you supply is the true one or not.

Before submitting your ranking you should calculate the associated score by clicking the button (Calculate score) below. You can try different rankings and calculate the associated scores, but note that after you click on submit you will not be able to change your ranking.

Recall that a higher score indicates a larger departure from the true ranking according to player B's payoff.

Screen 11: Instructions for Stage 2 for Control

Recall that the score associated with your reported ranking is XX, and any ranking can obtain one of the 9 possible scores: 0,1,2,3,4,5,6,7,8 and 9, where a higher score indicates a larger departure from the true ranking. If you want to read again the instructions regarding how the score is calculated please click here.

We now ask you to answer a set of questions regarding your belief on how your score compares to the score of other participants.

At the end we will randomly choose one of the questions and use your answer to that question to determine your payment for this stage of the experiment.

Screen 11: Instructions for Stage 2 for Score Treatment

Recall that the score associated with your reported ranking is XX, and any ranking can obtain one of the 9 possible scores: 0,1,2,3,4,5,6,7,8 and 9, where a higher score indicates a larger departure from the true ranking.[18] If you want to read again the instructions regarding how the score is calculated please click here.

The information presented earlier refers to the scores submitted by a previous group of participants. You belong to a new group of participants, and every member of your group has received the same information about the previous group as you.

We now ask you to answer a set of questions regarding your belief on how your score compares to the score of the other participants of the new group.

At the end we will randomly choose one of the questions and use your answer to that question to determine your payment for this stage of the experiment.

Screens 12-15: Instructions for Stage 2[19]

Recall that your score is XX, there are 10 possible scores: 0,1,2,3,4,5,6,7,8 and 9, and a higher score indicates a larger departure from the true ranking.

Please choose your preferred option.

(a) You will receive £4 if your score is in the top Y%. This means that your score is HIGHER than at least (100-Y)% of the scores of the participants other.

(b) You will participate in a computer generated lottery that will award you £4 with probability 50% and £0 with a probability of 50%.

---

[18]Note that although we enumerated all 10 possible scores in the instructions presented to both the control and score treatment, it was erroneously indicated that the possible scores are 9.

[19]We presented successive screens, where Y would take the value of 10, 20 and 30 and we either had top Y% or bottom Y%.

Experimental Instructions for Receivers

Screen 1: Log in

Screen 2: General Information

You are taking part in an economics experiment, the purpose of which is to examine the decisions people make in certain circumstances. You will be paid a participation fee of £3 for completing the experiment. You will also have the opportunity to earn additional money.

The money that you earn will be paid to you in cash. In particular, after all participants have completed the experiment you will receive an email from us with instructions regarding the specific times and the location at Highfield Campus to pick up your money from. Note that an assistant will pay you privately.

Screen 3: Consent Form

Screen 4: Instructions

In this experiment you will play a game with up to 15 other participants. You will be called player B and each other participant you are matched with will be called player A. The roles have been randomly allocated. You will not know with whom you were matched, neither now nor after the experiment, and we will not reveal your identity to the participants you are matched with either.

There are 6 outcomes, and each outcome specifies a payoff for you and one for each player A that you are matched with. Each player A has seen the payoffs associated with each outcome and has ranked the outcomes from 1 to 6.

Then, you will see a message of the following form.

"Player A has ranked the 6 outcomes, in terms of your payoff, from highest to lowest.

Please choose a number between 1 (representing the outcome that player A has ranked highest according to your payoff) and 6 (representing the outcome that player A has ranked lowest according to your payoff)."

Your choice will determine the payoffs for you and for each player A that you are matched with. Your payment will be determined as follows: in the end we will randomly choose one of the player As you are matched with and use their ranking and your choice of outcome to determine your payoff. For example, if you choose number 1, you will get the payment associated with the outcome ranked first by the randomly chosen player A. If you choose number 4, you will get the payment associated with the outcome ranked fourth by the randomly chosen player A.

You will never know what amounts were actually offered in the options not chosen. Furthermore, you will never know the amount that each player A has earned in the chosen option or would have earned in any of the other options not chosen by you.

Because the information conveyed to you by the message of each player A is identical, you will only see one message and your choice will be applied to the rankings of all player

As.

Screen 5: Choice

Player A has ranked the 6 outcomes, in terms of your payoff, from highest to lowest.

Please choose a number between 1 (representing the outcome that player A has ranked highest according to your payoff) and 6 (representing the outcome that player A has ranked lowest according to your payoff).