Journal of Behavioral and Experimental Economics 100 (2022) 101920

Contents lists available at ScienceDirect

Journal of Behavioral and Experimental Economics

journal homepage: www.elsevier.com/locate/jbee

Will I tell you that you are smart (dumb)? Deceiving Others about their IQ or about a Random Draw

Giovanni Burro [a,*], Alessandro Castagnetti [b]

[a] Bocconi University Department of Finance IGIER BIDSA, Italy
[b] University of Warwick, United Kingdom

ARTICLE INFO

JEL classification:
C92
D83
D90
D91

Keywords:
Cheating
Deception
Lying costs
Ability versus luck

ABSTRACT

We investigate whether individuals' deception rates differ by whether the messages they send to others are about the latter's relative ability or not. We also study whether they are more likely to deceive when it is in their interest to make others believe that they are either of high ability or low ability. In the experiment, participants play a sender–receiver game. The experiment features a 2 × 2 factorial design. First, we vary whether the state is determined by the receiver's relative performance in an IQ test or it is about a randomly drawn number. Second, monetary incentives, which are common knowledge, are such that the sender is better off (worse off) when the receiver's action is about him being of high (low) relative ability, while the receiver benefits from selecting the action that matches his true ability. We do not find systematic differences in the deception rates due to the ego-relevance of the condition. However, we find that senders believe they are more likely to move the beliefs of the receivers when they send a deceiving message saying that the receiver is of high ability.

1. Introduction

Communication is at the heart of social interactions. People with private information may be willing to misreport this information if this guarantees them an economic advantage. However, as opposed to what neoclassical economics would predict – that individuals would communicate whatever increases their own payoff – people lie less than expected (Gerlach et al., 2019). Recent literature in social sciences has underlined that individuals are constrained by psychological costs related to lying, therefore limiting the extent of how much they misreport (Abeler et al., 2019). Importantly, these costs may heavily depend on what type of information it carries to the recipient of the messages (Gneezy et al., 2017). We hypothesize that the psychological costs of lying may be significantly reduced if the (deceiving) message provides a psychological benefit to the counterpart. For instance, if the message flatters the recipient. If this is true, we can better understand the amount of information that ultimately will be transmitted in various strategic economic interactions.

In this paper, we focus on how deceiving behavior is shaped in social interactions, depending on the ego-relevance of the content of the communication. We look at whether individuals find it easier (more difficult) to deceive others when the message they send enhances (diminishes) others' personal characteristics. Differences in these psychological costs may play a key role in producing overconfident beliefs (Bolger et al., 2019; Moore & Healy, 2008) and may influence the outcome of crucial relevant economic interactions. Take the case of a car dealer (or even a private citizen) trying to sell a sports car. He might be interested in making the customer believe that he has good driving skills, which will enable him to enjoy the vehicle entirely. Think about a sales agent in a fashion store. He will have an incentive to make the customer believe he looks shining in the suit he is thinking of buying.

In our experiment, subjects play a sender–receiver game. In the game, the sender sends a message about the state of the world to the receiver, who then takes an action. The experiment features a 2 × 2 factorial design. We vary whether the state is determined by the receivers' performance in an IQ test (ego relevant condition) or through a random number (not ego-relevant condition). We then vary whether the sender's incentive is for the receiver to take an action corresponding to him being of a better ranking (positive condition) or a worse ranking (negative condition). Before and after the game, we elicit the beliefs of the sender about the beliefs of the receiver's relative ranking. The factorial design allows us to have exogenous variation in the ego-relevance of the messages and, therefore, to study the implications of these psychological costs on deception rates across conditions.

* Corresponding author.
E-mail addresses: giovanni.burro@unibocconi.it (G. Burro), s.castagnetti@warwick.ac.uk (A. Castagnetti).

https://doi.org/10.1016/j.socec.2022.101920
Received 9 August 2021; Received in revised form 2 June 2022; Accepted 19 July 2022
Available online 25 July 2022

We do not find systematic differences in the deception rates due to the ego-relevance of the condition. Moreover, this null effect is not driven by differences in senders' prior beliefs nor by differences in receivers' relative ability. We do not find strong evidence that the psychological costs of lying are reduced when deception provides flattering information to the counterpart. The result holds if we look at both the intensive and extensive margin of deception. Overall, the experimental results strongly suggest that these psychological costs are not strong enough to influence the outcome of interactions with misalignment of incentives.

We also explore senders' beliefs about the beliefs that the receivers have about their rank. First, we find that the senders believe that the receivers are overconfident about their abilities in the ego-relevant conditions. Second, we find that the senders believe in being able to influence the beliefs of the receivers about their own rank. In all conditions, the senders' prior and posterior beliefs differ significantly. Third, we find that when the senders engage in deceiving flattering behavior, they believe in being even more effective at moving the receivers' posterior beliefs. This suggests that senders believe that individuals are likely to incorporate flattering information into their posterior beliefs more readily.

2. Relation to the literature

This paper contributes to two main streams of literature. First, it contributes to the literature on communication experiments. In particular, we refer to the literature on cheap talk (sender–receiver) games, where the sender is informed about the state of the world, and there is misalignment of interests between senders and receivers. Indeed, the experimental design that we implement here is based on a cheap talk game that borrows features from Cai and Wang (2006) and Wang et al. (2010). Blume et al. (2020) offer a comprehensive literature review on cheap talk games. In brief, the experimental findings are that, with respect to theoretical predictions, senders reveal more information, and receivers react more to the messages (Crawford & Sobel, 1982). We add to this literature the study of senders' behavior in communication games where the state is ego-relevant and how it affects behavior compared to the standard case where the state is not ego-relevant.

The study of senders' behavior has started a literature on deception and lying aversion. This is the second and main literature to which we make our contribution. Seminal works on deceiving behavior in the sender–receiver game are Gneezy (2005) and Erat and Gneezy (2012).[1] Gneezy (2005) finds that senders are sensitive to their gain when deciding to lie. Moreover, they not only care about how much they gain from a lie but also how much the other side loses. However, this unselfish motive diminishes with the size of the sender's gains. Erat and Gneezy (2012) study two types of white lies: those that help others at the expense of the person telling the lie (altruistic white lies) and those that help both others and the liar (Pareto white lies).[2] They find that many senders are reluctant to tell even a Pareto white lie, demonstrating a pure lying aversion. In contrast, some senders are willing to tell an altruistic white lie that hurts them a bit but significantly helps others. Gerlach et al. (2019) and Rosenbaum et al. (2014) carry out meta-analyses of dishonest behavior in experiments, including the sender–receiver game. They show that dishonest behavior depends on situational factors, such as reward magnitude and externalities, and personal factors, such as the participant's gender and age. For example, deception rates depend on social norms (Mitra & Shahriar, 2020), on one's mood (Siniver & Yaniv, 2019), and the reporting system (Behnk et al., 2019). Interestingly, Gawn and Innes (2019) find that delegation reduces lying aversion: subjects are more willing to lie through a delegate than to lie directly.

From this recent and active literature, we draw our main experimental research question that the psychological costs of lying may also vary depending on the ego-relevance content of the information transmitted to the receiver. However, our results suggest that the ego-relevance of the messages does not significantly influence the psychological costs of lying.

Our closest predecessors are Ho and Yeung (2014) and Gneezy et al. (2017). The experiment of Ho and Yeung (2014) also features two roles: agents and clients. The agent is informed about the absolute performance of the client and sends him performance feedback. The client then reports a level of happiness that determines the agent's payoff. The paper finds that agents inflate the feedback, and clients report higher happiness levels. While there are many common features with our experiment, there are also significant differences, of which two stand out. First, the experimental conditions are not designed to tightly study deception rates and associated psychological costs of the senders, depending on the transmitted ego-relevant information. Second, the incentives are deliberately misaligned across the two roles in our experiment. This allows us to study lying behavior in a strategic environment.

Gneezy et al. (2017) investigate how individuals send feedback to others regarding their attractiveness. They show that people avoid giving accurate face-to-face feedback to less attractive individuals, even if lying in this context comes at a monetary cost to both the person who gives the feedback and the receiver. However, when feedback is provided anonymously, the aversion to giving negative feedback is reduced. While in our experiment we do not vary the level of anonymity (as we preserve it throughout), it enhances the scope of this paper in at least two dimensions. First, we look at the provision of feedback and deception rates when there are incentives to both enhance and diminish the receiver's IQ ability, and we can compare this to the provision of information when the underlying state is not ego-relevant. Second, in our setup, the players have misaligned incentives, so we can better learn deception rates in strategic settings.

Overall, we contribute a single important finding to the literature. While extensive literature emphasizes how psychological costs of lying differ across different settings and characteristics of the situation at hand, we do not find that in a relevant strategic interaction they significantly shape the truthfulness of the information being sent. In particular, in the case of this experiment, senders are not more or less likely to send opportunistic messages when they are about the receivers' IQ ability. They are as prone to send overly exaggerated positive (negative) messages in the positive (negative) ego-relevant condition as they are in the positive (negative) not ego-relevant conditions. In other words, their messages in the game are not shaped by the ego-relevance of the setting.

3. Experimental design

To causally investigate whether subjects are more (less) likely to deceive others when the message they send enhances (diminishes) others' personal characteristics, an experiment with the following features is required. First, a game with at least two players that allows for communication (e.g., the transmission of messages). Second, an action in the game that measures ones' propensity to deceive and, on top, an incentive compatible mechanism to study the impact of the messages on beliefs. Third, exogenous variation in the ego relevance of the task (i.e., whether messages are about one's personal characteristics or not). The laboratory provides a suitable setting to create an environment with these specifics. The experiment consists of two main parts, which are carefully explained below. First, subjects are asked to complete an IQ test. Subjects then play a sender–receiver game. Senders are

[1] Here, we focus on strategic interaction settings. For the literature on cheating that focuses on settings without strategic interactions, see Abeler et al. (2019).

[2] Capraro et al. (2019) studied, borrowing methods from physics, the evolutionary behavior of senders and receivers for white and black lies. Here, we are focused on one-shot interactions.

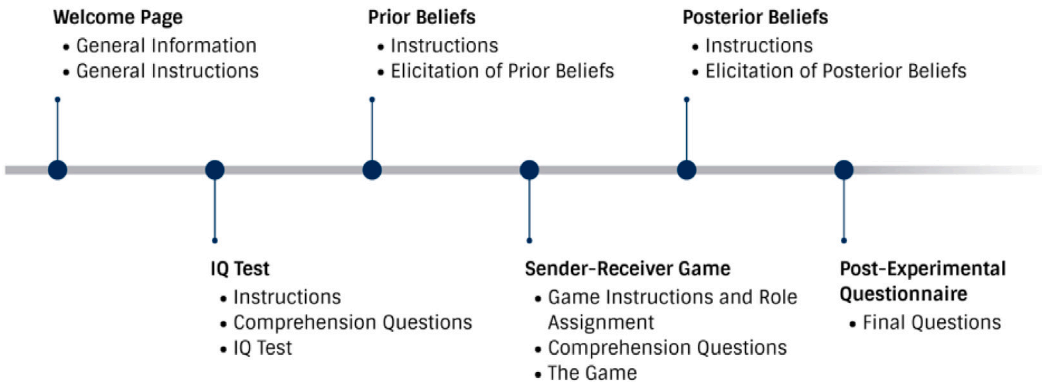

Fig. 1. Timeline of the experiment.

informed about the state of the world and send a message about the state to receivers, who then take an action. Payoffs are thus realized. The incentives in the game are such that the receivers' best interest is that their actions match the state, while senders profit the most from receivers taking a specific action.

The experiment features a 2 × 2 between-subject design. First, we vary the ego-relevance of the state. In the ego-relevant (not ego-relevant) condition, the state is determined by the receiver's performance in the IQ test (by a random draw). The state is just the ranking of the receiver in a group formed by the receiver and 9 other subjects who took part in a pilot session of the experiment. Second, we vary the payoffs in the game. In the positive (negative) condition, senders profit the most when the receivers take a low (high) ranked action in the game, corresponding to them being closer to the top (bottom) of the ranking. Before and after the game, we elicit receivers' beliefs about their ranking, and senders' beliefs about the receivers' beliefs. Fig. 1 shows the timeline of the experiment.

3.1. The IQ test

At the outset of the experiment, we asked participants to complete an IQ test: the Raven Advanced Progressive Matrices (APM) test. They had 10 min to complete the test, which consisted of 20 Raven matrices. The 20 Raven matrices used in the experiment are reported in Appendix E. Subjects earned Soles 5.00 per correct answer out of three randomly chosen questions.[3] Importantly, we also truthfully informed subjects that this test is often used to measure fluid intelligence (i.e., reasoning ability) and general intelligence; and that high scores in this test highly correlate with economic variables like income and occupation and health variables like health quality and longevity (Sternberg et al., 2001). This was accomplished to increase the ego-relevance of the task.

3.2. The sender-receiver game

After the test, subjects were randomly and evenly split into one of two roles: senders and receivers. Thus, the computer program randomly created sender–receiver pairs, and each pair played the sender–receiver game.[4] The game works as follows. First, the sender is informed about the state of the world (the "realized state"), which can take any value in the state space, $S_s = \{1, 2, 3, \dots, 10\}$, and that depends on the receiver's rank. Both players are aware of how the state of the world is determined. The sender then decides which message, m_s, to send to the receiver. The message space corresponds to the state space: $M_s = S_s = \{1, 2, 3, \dots, 10\}$. The receiver then reads the message sent by the sender and chooses an action: $a_r \in A_c = \{1, 2, 3, \dots, 10\}$. Payoffs are determined by both the receiver's action in the game and the state of the world. Fig. 2 gives a visual representation of the game. To make sure that subjects understood the main features of the game, they were asked to complete a comprehension questionnaire. They could not play the game until they answered these questions correctly.

3.2.1. Treatment variations

To causally study whether individuals are more likely to fool others depending on the psychological costs connected to the type of information transmitted, the experiment features a 2 × 2 factorial design. The factors correspond to variations in the ego-relevance of the state and how payoffs are determined in the game. We explain them in detail below.

Ego relevance variation. In the ego-relevant condition, receivers were ranked according to their IQ scores. In particular, their scores were compared to the scores of 9 other subjects who took part in a pilot session of the experiment. They were informed of this ranking procedure and that the scores elicited a strict ordering. If two or more subjects had the same score, then it was randomly determined whose ranking was higher. In the not ego-relevant condition, the ranking was determined by the draw of a random number. In particular, subjects were informed that their ranking would be determined by a random number drawn from a given distribution. In a between-subject design, these distributions could take one of the following forms: (1) uniform distribution where each rank was drawn with equal probability; (2) a positively skewed distribution where higher rankings were drawn with higher probability; (3) a negatively skewed distribution where lower rankings were drawn with higher probability. We varied the distributions to have exogenous variation in prior beliefs in the not ego-relevant condition. This is done to have a distribution of prior beliefs that more closely resembles that in the ego-relevant condition. In the Appendix, we provide a detailed description of the distributions.[5]

This experimental variation allows us to study whether senders are more likely to send overly positive messages when the state is about the receivers' relative performance in the IQ test, compared to the case in which the state has been randomly determined.

[3] At the time of the experiment (June 2020), the exchange rate was: \$ 1.00 = Peruvian Soles 3.54.

[4] To prevent framing effects, in the experiment senders (receivers) were called Player 1 (Player 2).

[5] The analysis of prior beliefs, after the experiment was carried out, shows that prior beliefs in the not-ego relevant treatment are indeed close to those in the ego-relevant treatment. A Wilcoxon rank test, comparing prior beliefs in the two conditions, cannot reject the null hypothesis of equality of the distributions (p=0.59).

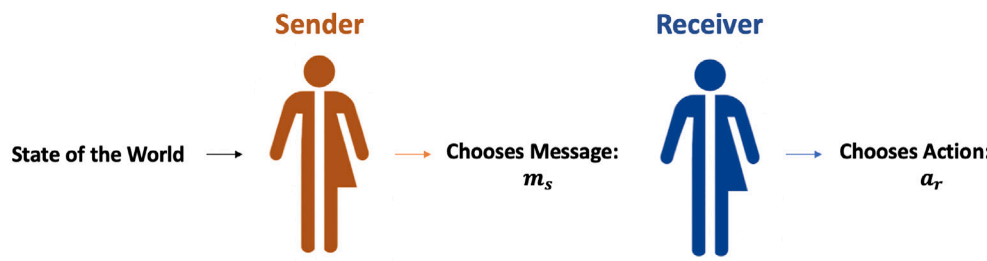

Fig. 2. The sender–receiver game.

Payoff variation. The sender's payoff was determined by the receiver's action in the game. In the positive condition, her payoffs increased monotonically as the receiver played lower numbers in the game. While in the negative condition, payoffs were reversed: the sender's payoff monotonically increased as the receiver played higher numbers. Therefore, in both conditions, the sender's payoff was not dependent on the receiver's rank. The receiver's payoff in the game (and irrespective of the condition) was determined by his action and the realized state. In particular, the receiver's payoff was maximum when his action matched the state and monotonically decreased as his action deviated (in absolute terms) from the realized state. Fig. 3 shows the payoff structure for both players and by condition. The payoffs for both senders and receivers are similar to the ones in Jin et al. (2021) and Wang et al. (2010). Both players knew the game's payoff structure and were explicitly made aware of the misalignment of interests in the game incentives across roles.

This variation allows us to further investigate the link between the ego-relevance of the state and the propensity to fool others. In particular, it lets us also study whether there are asymmetric responses to negative news by ego-relevance of the state. These analyses are crucial as they allow to exclude other confounding effects that may be affecting differences in actions due to the ego-relevance of the state. For instance, senders might believe that the receivers respond more to messages about their personal characteristics. We can assess this possibility by having the 2 × 2 factorial design while also cleanly answering our research question. Fig. 4 provides a summary of the resulting experimental treatments.

3.3. Prior and posterior beliefs

Before and after the sender–receiver game, we asked participants the following set of beliefs.

Receivers' beliefs. We asked receivers their prior (posterior) beliefs about their relative ranking before (after) the sender–receiver game.[6] This corresponded to their relative ranking in the IQ test in the ego-relevant condition. While in the not ego-relevant condition, it was determined through the random draw. In particular, we elicited the entire distribution of these prior beliefs. That is, receivers had to write down the probability with which they thought they occupied each of the 10 ranks.

Senders' beliefs. We asked senders to report their beliefs about what their matched receivers thought their mean ranking was before and after the game. Again, we elicited the entire distribution of the prior (posterior) beliefs.[7] Importantly, before the belief elicitation stage, senders were informed about their matched receiver's rank, how the ranking was determined, and that the receivers did not know their own true rank.

The elicitation of both prior and posterior beliefs is crucial to studying senders' actions in the game controlling for prior beliefs. Conversely, one could acknowledge differences in gameplay that ultimately are not driven by ego-related psychological costs but by differences in prior beliefs. Second, posterior beliefs allow us to analyze whether messages sent in the game are thought to impact receivers' beliefs.

The elicitation procedure was monetarily incentivized with the Binarized Scoring Rule proposed by Hossain and Okui (2013) and with a fixed price of Soles 20.00. Under this method, truthful reporting is orthogonal to subjects' risk preferences, and it does not rely on expected utility theory.[8] We explicitly and truthfully told participants that the elicitation mechanism guaranteed that it was in their best interest to report their true beliefs. We, instead, did not explain to subjects how the procedure worked as withholding the description of the mechanism increases truthful reporting[9] (Danz et al., 2020).

3.4. Post-experimental questionnaire

At the end of the experiment, subjects were asked a set of unincentivized questions. First, senders were asked to report the probability with which they thought that their matched receivers followed the message they sent. Similarly, receivers were asked to report the probability with which they believed that their matched senders sent a truthful message. We then asked them a general willingness to take risks question (Dohmen et al., 2011). Finally, participants completed a demographic questionnaire that included questions about their age, gender, and student status.

3.5. Implementation

The experimental sessions were conducted in July 2020. We recruited subjects through the Orsee recruitment system, and we used the pool of participants registered at the economics laboratory of Universidad Catolica del Perú in Lima, Perú. After removing those participants who did not complete the experiment (mainly due to internet connection issues) and their matched partners, we were left with a total of 314 senders: 141 senders took part in the ego-relevant and positive treatment, 53 to the not ego-relevant and positive treatment, 68 to the ego-relevant and negative treatment, and 52 to the not ego-relevant and negative treatment. On average, sessions lasted 45 min. Participants earned an average payment of Soles 12.00, including the show-up fee of Soles 5.00. We programmed and conducted the experiment in oTree (Chen et al., 2016). The Appendix includes the experimental instructions (translated from Spanish) and descriptive statistics of the sample (Table B.1).

The sessions were conducted online. Each participant registered in advance for an online session (and only once) that took place on a

[6] At the time we asked participants their prior beliefs, they were not aware of what they would be doing in the following task of the experiment.

[7] For both belief questions and roles, we imposed the natural constraint that these probabilities needed to sum up to 100%.

[8] For a detailed explanation of this elicitation procedure, see also Schotter and Trevino (2014).

[9] The interested participants, however, could click on a button to read a detailed description of the elicitation method.

(a) Payoff Table in the Positive Condition

	Receiver's Action in the Game									
	1	2	3	4	5	6	7	8	9	10
Receiver's Ranking = 1	15,15	14,14	13,13	12,12	11,11	8,8	7,7	5,5	3,3	1,1
Receiver's Ranking = 2	15,14	14,15	13,14	12,13	11,12	8,11	7,8	5,7	3,5	1,3
Receiver's Ranking = 3	15,13	14,14	13,15	12,14	11,13	8,12	7,11	5,8	3,7	1,5
Receiver's Ranking = 4	15,12	14,13	13,14	12,15	11,14	8,13	7,12	5,11	3,8	1,7
Receiver's Ranking = 5	15,11	14,12	13,13	12,14	11,15	8,14	7,13	5,12	3,11	1,8
Receiver's Ranking = 6	15,8	14,11	13,12	12,13	11,14	8,15	7,14	5,13	3,12	1,11
Receiver's Ranking = 7	15,7	14,8	13,11	12,12	11,13	8,14	7,15	5,14	3,13	1,12
Receiver's Ranking = 8	15,5	14,7	13,8	12,11	11,12	8,13	7,14	5,15	3,14	1,13
Receiver's Ranking = 9	15,3	14,5	13,7	12,8	11,11	8,12	7,13	5,14	3,15	1,14
Receiver's Ranking = 10	15,1	14,3	13,5	12,7	11,8	8,11	7,12	5,13	3,14	1,15

(b) Payoff Table in the Negative Condition

	Receiver's Action in the Game									
	1	2	3	4	5	6	7	8	9	10
Receiver's Ranking = 1	1,15	3,14	5,13	7,12	8,11	11,8	12,7	13,5	14,3	15,1
Receiver's Ranking = 2	1,14	3,15	5,14	7,13	8,12	11,11	12,8	13,7	14,5	15,3
Receiver's Ranking = 3	1,13	3,14	5,15	7,14	8,13	11,12	12,11	13,8	14,7	15,5
Receiver's Ranking = 4	1,12	3,13	5,14	7,15	8,14	11,13	12,12	13,11	14,8	15,7
Receiver's Ranking = 5	1,11	3,12	5,13	7,14	8,15	11,14	12,13	13,12	14,11	15,8
Receiver's Ranking = 6	1,8	3,11	5,12	7,13	8,14	11,15	12,14	13,13	14,12	15,11
Receiver's Ranking = 7	1,7	3,8	5,11	7,12	8,13	11,14	12,15	13,14	14,13	15,12
Receiver's Ranking = 8	1,5	3,7	5,8	7,11	8,12	11,13	12,14	13,15	14,14	15,13
Receiver's Ranking = 9	1,3	3,5	5,7	7,8	8,11	11,12	12,13	13,14	14,15	15,14
Receiver's Ranking = 10	1,1	3,3	5,5	7,7	8,8	11,11	12,12	13,13	14,14	15,15

Fig. 3. Payoff table by condition.
Notes: The tables show the payoff structure by payoff condition. In the top panel (a) the table displays the payoff matrix for the positive payoff condition, while in the bottom panel (b) the table displays the payoff matrix for the negative payoff condition. The columns indicate the receiver's action in the sender–receiver game, while the rows indicate the realized state of the world, which corresponds to the receiver's actual ranking. In each cell, the left entry (in red) shows the sender's payoff, while the right entry (in blue) shows the receiver's payoff. Payoffs are in Peruvian Soles.

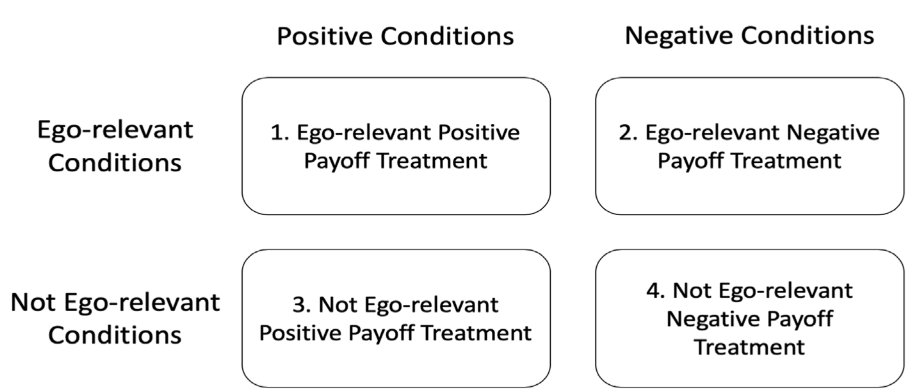

Fig. 4. Summary of Experimental Conditions and Corresponding Treatments.

particular day and time. Registered participants received a reminder the day of the session. Two research assistants supervised the sessions, and participants could contact them (via email or text message) in real-time if needed.

3.6. Power analysis and discussion

In Section 4.1 we will soon discuss that the ego-relevance of the message does not influence the rate of deception. Given our sample size and that our main finding is a null result, we would like to briefly discuss the issue of power to allow the reader to read the remainder of our work with a full understanding of its statistical properties. To mitigate power concerns, we conducted an ex-post power analysis. We fixed the probability of false positive to the standard threshold, $\alpha = 0.05$. In the positive condition, we reach the standard threshold of 80% power to detect a 19pp difference between the not ego and the ego-relevant condition.[10] In the negative condition, we reach 80% power to detect a 26pp effect.[11] We believe that these are meaningful effects.[12] At least, we can conclude that we put an upper bound on the impact of ego-relevance on deception. It is possible that the ego-relevance of the information leads to an increase in the probability of deception. However, our work indicates that this effect is, if existing, relatively small.

4. Results

In this section, we will focus on two sets of results. First, in Section 4.1 we will focus on the possibility that the senders misreport the ranking of the receiver in order to lead the receiver to play a move that would increase the senders' payoffs. Second, in Section 4.2 we will analyze whether and in which case the senders believe the receivers to be overconfident and whether and in which cases they believe that they are able to change the receivers' beliefs about their own relative rank.

4.1. Deceiving behavior

We analyze whether the senders engage in deceiving behavior when they send a message to the receivers and whether this differs depending on the ego-relevance of the task. With "deception" we will refer to the sender sending a message strictly smaller (greater) than the receiver's ranking in the positive (negative) condition. Hence, a sender is said to deceive if she tried to bias the receiver's response towards the action which would maximize the sender's own payoff.[13] In theory, we expect the ego relevance of the task to have a different impact in the positive and in the negative condition. We expect the deception rate to be higher (lower) when the message is ego-relevant, in the positive (negative) condition, with respect to the not ego-relevant one. The idea is that in the positive condition, the psychological cost of deception (Abeler et al., 2019) will be lower if the message is ego-relevant since deception might boost the ego of the respondent. The opposite is true in the negative condition. We do not expect reputation or social norms to play a role since anonymity is preserved throughout. Hence, defining d_e the deception rate in the ego-relevant condition and d_{ne} the deception rate in the not ego-relevant one we test:

$H_0 : d_e \leq d_{ne}$ $H_1 : d_e > d_{ne}$ in positive condition

$H_0 : d_e \geq d_{ne}$ $H_1 : d_e < d_{ne}$ in negative condition

Overall, in 64% of the cases a sender deceived. Deception rates in the four experimental conditions are reported in Fig. 5. They range from around 51% in the negative ego-relevant condition to around 70% in the positive ego-relevant condition, with the other two measures in-between (62% in the negative not ego-relevant condition and 66% in the positive not ego-relevant condition). From this figure, we can already see that the ego-relevance of the condition does not have a relevant impact on the willingness to deceive the respondent, both in the positive and in the negative condition.[14] A one-sided t-test could not reject the null hypothesis that the deception rate was lower (higher) or equal in the ego-relevant task in the positive (negative) condition with p=0.33 (p=0.14) in the positive (negative) condition. To control for possible confounding factors, we perform the following logistic regression:

$$Deception_i = \Lambda(\alpha + \beta Ego_i + \gamma Rank_i + \boldsymbol{\delta}\mathbf{x_i} + \epsilon_i) \quad (1)$$

We perform the analysis separately for the positive and the negative conditions. $Deception_i$ is a dummy equal to 1 if the sender i sent a message strictly smaller (greater) than the receiver's ranking in the positive (negative) condition. Ego_i is a dummy equal to 1 if the sender plays in the ego-relevant condition. $Rank_i$ is the ranking of the receiver to whom the sender i is matched. $\mathbf{x_i}$ is a vector containing controls about sender i. Controls include demographics, IQ, risk preferences, and the prior of the sender. To give a sense of them, past research showed that males tend to be more dishonest than females (Abeler et al., 2019; Dreber & Johannesson, 2008) and that children with higher IQ have a higher likelihood of cheating (Alan et al., 2020).

The results are reported in Table 1 (full results in Table C.1 in the Appendix). We do not find any effects of the ego-relevance of the condition on deception rates. This holds even after controlling for the ranking of the receiver and controls, including differences in senders' prior beliefs. However, we find that the ranking of the receiver affects the deceiving propensity of the sender. In the positive

Table 1
Deception Extensive Margin. The dependent variable is a dummy equal to 1 if the sender sent a message strictly smaller (greater) than the receiver's ranking in the positive (negative) condition. Ego is a dummy equal to 1 for the ego-relevant condition and Ranking is the ranking of the receiver. Logit average marginal effects. Controls include demographics, IQ, risk preferences, and prior of the sender.

	Dependent variable:					
	Deception					
	Positive condition			Negative condition		
	(1)	(2)	(3)	(4)	(5)	(6)
Ego	0.034	−0.010	−0.021	−0.100	0.069	0.048
	(0.074)	(0.065)	(0.067)	(0.090)	(0.077)	(0.078)
Ranking		0.080***	0.076***		−0.099***	−0.101***
		(0.008)	(0.011)		(0.008)	(0.010)
Controls	No	No	Yes	No	No	Yes
Observations	194	194	194	120	120	120
Pseudo R^2	0.001	0.199	0.218	0.007	0.289	0.321

Note: *p<0.1; **p<0.05; ***p<0.01.

[10] We have power equal to 81.2%, to be more precise, taking for granted a 66.04% deception rate in the positive not ego-relevant condition.

[11] We have power equal to 81.5%, to be more precise, taking for granted a 61.54% deception rate in the negative not ego-relevant condition.

[12] To give a sense of the magnitude of the findings in similar papers, works with similar research questions found an effect that ranged from a minimum of 16pp (Erat & Gneezy, 2012) to a maximum of 59.4pp (Ho & Yeung, 2014), with other effects in the order of 32pp (Erat & Gneezy, 2012), or 19pp to 35pp (Gneezy, 2005). In Section 2 we introduced these papers in detail and we clearly explained that, although close to our research design, some differences remain. Even though none of them find a relatively small effect size, we cannot however exclude that with our design this should be the magnitude that we should expect.

[13] A small percentage of senders (14.6%) deceived in the direction which minimized their own payoff.

[14] The difference in the deception rate in the positive and negative conditions are significantly different, possibly suggesting that there is a higher psychological cost of deception in the negative condition. This difference disappears after controlling for the ranking of the respondent. However, when pooling all data together (to compare the positive and negative conditions), we need to redefine ranking in at least one of the two conditions (e.g. as 11 minus ranking) to make it comparable across conditions. This problem does not arise when analyzing positive and negative conditions separately, as we do in the main text.

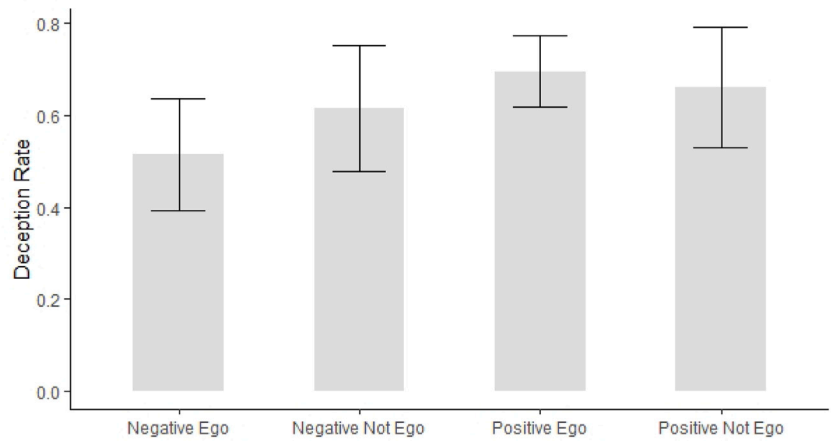

Fig. 5. Deception rate by condition. A sender is said to deceive if she sent a message strictly smaller (greater) than the receiver's ranking in the positive (negative) condition (95% confidence intervals).

condition, a higher ranking induces a higher rate of deceiving behavior, and in the negative condition, a lower ranking induces a higher rate of deceiving behavior. In particular, the probability of engaging in deceiving behavior increases by around 8% when the ranking increases by 1 in the positive condition, and it decreases by around 10% when the ranking increases by 1 in the negative condition. This means the sender is more likely to engage in deceiving behavior if there is more room to deceive and a higher potential reward from deception. We can provide some examples, focusing on the positive condition. First, if the receiver's ranking is 10, there are 9 possible messages that the sender can send to deceive the receiver, while if the receiver's ranking is 2, there is only 1 possible message which allows the sender to deceive the receiver, according to our definition of deception. Second, if the receiver's ranking is 10, deception offers a high potential reward. If the sender sends a message equal to, say 5, and the receiver follows the advice, the sender will get a payoff of 11. If the sender sends a message of 10 and the receiver follows the advice, the sender will get 1. On the contrary, if the receiver's ranking is 2, the sender sends 2, and the receiver plays 2, the payoff for the sender will be 14, while if the sender sends 1 and the receiver plays 1, the sender will get 15. The difference of 1, in this case, is much smaller than the difference of 10 in our previous example, where the receiver's ranking was 10. It suggests that if there are more deceiving options available, the sender is more likely to deceive. We can relate this finding to the literature on the psychological cost of deception (Gerlach et al., 2019).

After looking at the extensive margin of deceiving behavior, we now focus on the intensive margin. We define the intensive margin of deception as the difference between the receiver's ranking and the sender's message in the positive condition. In the negative one, it is defined as the difference between the sender's message and the receiver's ranking. Hence, a bigger value corresponds to higher deception in both cases. Similar to what we said for the extensive margin of deception, we expect the ego relevance of the task to have a different impact in the positive and in the negative condition. In the positive (negative) condition, we expect deception to be higher (lower) in the ego-relevant task. The intensive margin of deception takes values of 0.65 in the negative ego condition, 1.81 in the negative not ego, 2.38 in the positive ego, and 1.72 in the positive not ego. Hence, defining D_e the deception rate in the ego-relevant condition and D_{ne} the deception rate in the not ego-relevant one, we test:

$$H_0 : D_e \leq D_{ne} \;\; H_1 : D_e > D_{ne} \text{ in positive condition}$$

$$H_0 : D_e \geq D_{ne} \;\; H_1 : D_e < D_{ne} \text{ in negative condition}$$

We do find that a one-sided t-test rejects the null hypothesis in both the positive (p<0.1) and the negative (p<0.05) conditions. To investigate this issue in greater detail, we perform the following regression:

Table 2

Deception Intensive Margin. The dependent variable is the difference between the ranking of the receiver and the message sent by the sender, in the positive condition. It is message minus ranking in the negative one. Ego is a dummy equal to 1 for the ego-relevant condition. Prior is senders' prior beliefs about the receivers' beliefs on their own rank, before the message is sent. Controls include demographics, IQ, and risk preferences of the sender.

	Dependent variable:			
	DeceptionI			
	(1)	(2)	(3)	(4)
Ego	0.659	0.804*	−1.161**	−0.941
	(0.454)	(0.427)	(0.516)	(0.586)
Prior		0.678***		−0.518**
		(0.201)		(0.245)
Controls	No	Yes	No	Yes
Observations	194	194	120	120
R^2	0.012	0.126	0.037	0.087
Adjusted R^2	0.007	0.093	0.029	0.030

Note: *p<0.1; **p<0.05; ***p<0.01; Robust s.e. in parenthesis.

$$DeceptionI_i = \alpha + \beta Ego_i + \gamma Prior_i + \boldsymbol{\delta}\mathbf{x_i} + \epsilon_i \quad (2)$$

where $DeceptionI$ is the intensive margin of deception. $Prior_i$ is the sender's i prior beliefs about the receivers' beliefs on their own ranking before the message is sent. We perform the analysis separately for the positive and the negative conditions. The results are reported in Table 2 (full results in Table C.2 in the Appendix). The only noticeable difference is that ego relevance of the task significantly decreases the intensive margin of deception by 1.1 points in the negative condition.[15] This difference is no longer significant once we add controls. In particular, if the sender believes that the respondent thinks of having a higher ranking before the message is sent, the intensive margin of deception significantly increases (decreases) in the positive (negative) condition. This is an expected effect, since a high prior belief means that the sender has more (less) room for deception in the positive (negative) condition.[16]

[15] Notice that in the regression we fit robust standard errors and that the regression naturally tests for a two-sided hypothesis, while we had reported the *p*-value of a one-sided t-test, previously.

[16] Prior beliefs had no impact on the extensive margin of deception, as can be seen in the Appendix. However, the ranking was included as a control when testing the extensive margin of deception. Here, we deliberately excluded it from the list of controls since it is taken into account by the definition of the intensive margin of deception.

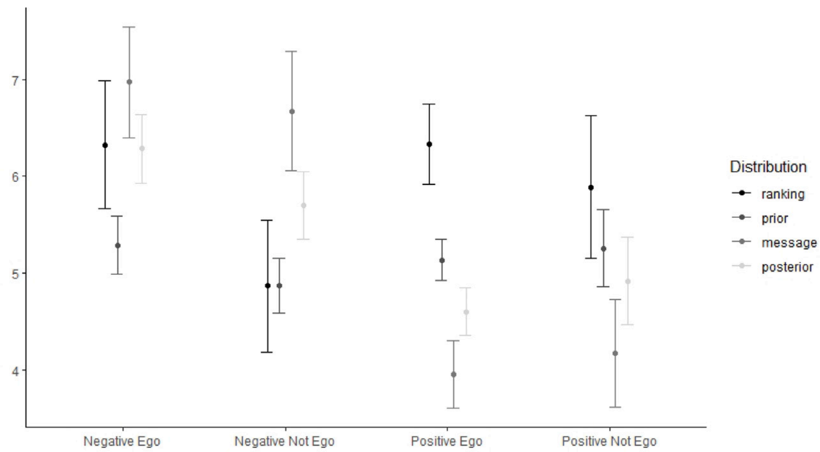

Fig. 6. Rankings, beliefs and messages. The figure reports the mean (with 95% confidence intervals) of the rankings, prior and posterior beliefs of the senders, and the messages sent by the senders for each condition.

4.2. Beliefs analysis

Although we did not identify an effect of ego-relevance on the propensity of the senders to deceive the receivers, we did find that the senders engage in deceiving behavior. We now dig into some psychological mechanisms that might drive senders' behavior. We focus on believed (by senders) overconfidence of the receivers and on the self-perceived ability of the senders to move receivers' beliefs. The first insight into our findings is offered by Fig. 6. Believed overconfidence is measured as the distance between the senders' prior beliefs about the receivers' beliefs on their own ranking and the actual ranking of the receivers. The higher this value, the stronger the sender's belief that the respondent believes to be better than he is. As it is clear from Fig. 6 and statistical tests in the two ego-relevant conditions, this distance is different from 0, being 1.038 in the negative and 1.195 in the positive ego-relevant conditions (two-sided t-test with $p<0.01$.). This means that the senders believe the receivers' beliefs to be around 1 ranking closer to the top with respect to their true rank. Our measure of overconfidence is not statistically different (at the conventional 5% level) from 0 in the not ego-relevant conditions, instead. This finding contributes to the credibility of our manipulation check. The fact that the senders believe that the receivers are overconfident in the ego-relevant settings but not in not ego-relevant ones suggests that the senders believe that our manipulation is successful at influencing the ego-relevance of the task. Moreover, we find that the receivers are overconfident in the ego-relevant conditions while they are not in the not ego-relevant ones. We find that the distance between the actual ranking and the receiver's beliefs about their own ranking before receiving the message is significantly positive in both of the two ego-relevant conditions (two-sided t-test with $p<0.01$), while it is not significantly different from 0 (two-sided t-test with $p>0.05$) in the two not ego-relevant conditions.[17] These findings on beliefs help rule out the hypothesis that a failure of the ego-relevance manipulation might drive our null result.

Moreover, from Fig. 6, we can already see that believed overconfidence is particularly stronger in the negative ego-relevant condition with respect to the negative not ego-relevant one and that the self-perceived ability (of the senders) to move the receivers' beliefs is not related to the ego-relevance of the task. To add more details to this research question we perform some regression analyses (see Table 3):

$$Overconfidence_i = \alpha + \beta Ego_i + \boldsymbol{\delta}\mathbf{x_i} + \epsilon_i \quad (3)$$

[17] $p = 0.67$ in the negative and $p = 0.09$ in the positive not ego-relevant condition.

Table 3
Overconfidence. The dependent variable is the difference between the actual ranking of the receiver and the belief of the sender about the receiver's belief before the message is sent. A bigger value is associated to higher believed (by the sender) overconfidence of the receiver. Ego is a dummy equal to 1 for the ego-relevant condition. Controls include demographics, IQ, and risk preferences of the sender.

	Dependent variable:			
	Overconfidence			
	Positive Condition		Negative Condition	
	(1)	(2)	(3)	(4)
Ego	0.563	0.645*	1.040**	0.963**
	(0.363)	(0.350)	(0.438)	(0.463)
Controls	No	Yes	No	Yes
Observations	194	194	120	120
R^2	0.013	0.039	0.045	0.062
Adjusted R^2	0.008	0.008	0.037	0.012

Note: *p<0.1; **p<0.05; ***p<0.01; Robust s.e. in parenthesis.

We perform the analysis separately for the positive and the negative condition. The dependent variable $Overconfidence_i$ is the believed (by senders) overconfidence of the receivers. For sender i, it is defined as the difference between the actual ranking of the receiver and the belief of the sender about the receiver's belief before the message is sent. The other variables are defined as in Eqs. (1) and (2). To give a sense of them, it is known for example that males are more confident than females regarding their own investment abilities (Barber & Odean, 2001). Here, we ask if this also translates into believed overconfidence about others. However, they do not have a relevant effect on believed overconfidence, as can be seen from Table C.3 in the Appendix. Although positive, the effect of the ego-relevance of the condition is not significant at the 5% level in the positive condition. In the negative condition, the task's ego-relevance is associated with a significant increase of around 1 point in the overconfidence measure (Table 3). If we pool together these two conditions and run a regression with robust errors and controls (not shown), we find that the ego relevance of the condition leads to higher significant overconfidence, with a pooled effect of 0.8 (ego coefficient significant at $p<0.01$). While there is abundant literature that individuals are overconfident about their personal characteristics and abilities (Moore & Healy, 2008), here we provide evidence that individuals also anticipate that others are overconfident.

Finally, we explore whether the senders believe that they can move the beliefs of the receivers. We once more stress that both prior and posterior beliefs of the senders are the beliefs that the senders have about what the receivers believe of themselves. Hence, what follows

Table 4
Posterior Minus Prior. The dependent variable is the difference between the belief of the sender about the receiver's belief after minus the one before the message is sent. Ego is a dummy equal to 1 for the ego-relevant condition and Ranking is the ranking of the receiver. Message is the message sent by the sender. Deception is a dummy equal to 1 if the sender sent a message strictly smaller (greater) than the receiver's ranking in the positive (negative) condition. Controls include demographics, IQ, and risk preferences of the sender.

	Dependent variable:									
	Posterior minus Prior									
	Positive Condition					Negative Condition				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Ego	−0.194	−0.187	−0.128	0.450	−0.123	0.167	0.157	0.101	0.116	0.184
	(0.209)	(0.209)	(0.205)	(0.393)	(0.212)	(0.242)	(0.232)	(0.226)	(0.378)	(0.218)
Ranking		−0.017			−0.101***		0.007			−0.052
		(0.039)			(0.038)		(0.045)			(0.044)
Message			0.301***		0.339***			0.222***		0.240***
			(0.063)		(0.062)			(0.074)		(0.072)
Deception				−0.262					0.442	
				(0.355)					(0.349)	
Ego ∗ Deception				−0.914**					0.185	
				(0.458)					(0.486)	
Controls	No	No	No	No	Yes	No	No	No	No	Yes
Observations	194	194	194	194	194	120	120	120	120	120
R^2	0.004	0.004	0.189	0.107	0.261	0.004	0.004	0.146	0.046	0.163
Adjusted R^2	−0.002	−0.006	0.180	0.092	0.229	−0.005	−0.013	0.131	0.021	0.103

Note: *p<0.1; **p<0.05; ***p<0.01; Robust s.e. in parenthesis.

has to be interpreted as "the perceived persuasion ability the senders have on the respondents' beliefs". We define the senders' perceived ability to influence the beliefs of the receivers as the difference between the beliefs of the senders about the receivers' beliefs after and before the message is sent. Senders perceived ability to influence respondents' beliefs is statistically different from 0 in all conditions (two-sided t-test with $p<0.01$ in all but the positive not ego-relevant where it is $p<0.05$). Hence, the senders believe in being able to influence the receivers' beliefs irrespective of whether the messages carry information about their IQ or the draw of the random number. To add more details to the picture, we perform the following regressions:

$$Moveprior_i = \alpha + \beta Ego_i + \gamma Rank_i + \nu Message_i + \boldsymbol{\delta} \mathbf{x_i} + \epsilon_i \quad (4)$$

$$Moveprior_i = \alpha + \beta Ego_i + \lambda Deception_i + \psi Ego_i \times Deception_i + \epsilon_i \quad (5)$$

We perform the analysis separately for the positive and the negative condition. The dependent variable $Moveprior_i$ is the difference between the sender's posterior and prior beliefs. A bigger value is associated with a higher believed (by the sender) ability to move the receiver's beliefs. The other variables are defined as in Eqs. (1)–(3). We introduced two separate models, in Eq. (4) and Eq. (5) since the deception variable and the message variable are highly collinear. Bear in mind that in the positive conditions, the sender's objective should be to move the posterior of the receiver downwards (towards lower ranks) and the opposite in the negative conditions. The results are reported in Table 4, while the full results are reported in Table C.4 in the Appendix.

Our first finding is that the senders believe that a higher message is more likely to move the beliefs of the receivers. This holds after controlling for the ranking of the receiver. This suggests that senders strongly believe that messages are not cheap talk but have an effect on the receivers' beliefs.[18] Second, in the positive conditions, we also find that, in case of deceiving behavior, senders believe that they are more likely to push the posterior of the receivers towards a lower ranking when the information is ego-relevant. The coefficient of the interaction between Ego and Deception in column 4 of Table 4 is negative and significant. This suggests that in case of deceiving behavior in the positive ego-relevant condition, the senders believe that they can, holding the prior constant, push the posterior one point lower on average, compared to the not ego-relevant condition. In other words, in the positive ego-relevant condition, the senders believe that their deceiving messages can be more persuasive than in the not ego-relevant condition, possibly thinking that the receivers will be flattered by the deceiving messages. In sum, while senders do not engage in higher deceiving behavior in the ego-relevant task, as we observed in Section 4.1, they believe that they will be more persuasive in case of deception. In the negative conditions, it is true that senders believe that a higher message is more likely to move the beliefs of the receivers. However, we do not detect any noticeable pattern coming from the ego relevance of the task and/or the deceiving behavior of the sender, since the coefficient of Ego, the one of Deception, and that of their interaction are not significant. The senders are not confident of persuading the senders in making them believe they performed poorly.

5. Concluding remarks

In this paper, we studied whether deception is influenced when the communication is about the recipient's personal characteristics in a setting of strategic information transmission. We hypothesized that providing information about others' IQ abilities may influence the psychological costs of lying. The experiment's findings clearly showed that deception is not higher (lower) when they enhance (diminish) the IQ ability of the recipient of the message. We did not find systematic differences on the deception rates due to the ego-relevance of the condition. Thus, we find that deception does not react to whether it provides good or bad vs neutral news about the personal characteristics of the receiver. One possible explanation for the absence of an increase in the deception rate due to the ego-relevance of the condition might be that the senders do not believe that they would be more successful in deceiving the receivers in the ego-relevant conditions. To explore this possibility, senders were asked to report the probability with which they thought that their matched receivers followed the message they sent. If the senders believed that the receivers should be more likely to be deceived in the ego-relevant condition, we would expect this probability to be higher in the ego-relevant case, particularly in case of deception. This is not true, as can be seen in the analyses contained in the Appendix (Table C.5). Senders are not more confident that the receiver will match the message in the ego-relevant conditions and, in particular, when they try to deceive the receivers in the ego-relevant conditions.

Our null result may also be driven by the fact that the incentives are common knowledge. The receivers know that the senders have an

[18] Remember that we elicited beliefs in an incentive-compatible fashion. In other words, senders had an economic incentive to report their beliefs truthfully.

incentive and know exactly how big that incentive is. In some real-life interactions, the situation is more nuanced. People receiving messages can only suspect that the senders have an incentive to deceive them, and possibly they do not know how big that incentive is. With our set-up, we refer to all those situations where it is clear that the agent conveying a message has a clear incentive to deceive. Think about sales agents in general and the examples we gave in Section 1, in particular.

We also found that the senders believe the receivers to be overconfident in ego-relevant conditions. The belief in overconfidence is significantly higher in the negative ego-relevant condition than in the negative not ego-relevant condition. Moreover, senders believe that they will be more effective at persuading the receivers of their relative ability in case of deception in the positive ego-relevant condition with respect to the positive not ego-relevant condition. This difference does not arise between the ego and not ego-relevant negative conditions. These pieces of evidence speak to the fact that, although senders do not believe they can fully convince receivers, they believe that self-serving beliefs exist and that they can at least partially influence those beliefs.

As we already discussed in Section 3.6, a possible concern for our work is the potential lack of power. This concerns our main finding that the ego-relevance of the message does not influence the rate of deception. An ex-post power analysis tells us that, in the positive (negative) condition, we reach the standard threshold of 80% power to detect a 19pp (26pp) difference between the not ego and the ego-relevant condition. While these can be deemed meaningful effects, we cannot rule out that a smaller effect exists. However, if such an effect exists, we put a meaningful upper bound on its magnitude.

Finally, the issue of publication bias has raised concerns in both economics (Brodeur et al., 2020; Chopra et al., 2022; Stanley, 2005) and psychology (Francis, 2013; Kühberger et al., 2014). This bias arises if a preference for statistically significant results arises. This might potentially create a gap in the advancement of the literature. Moreover, an increase in the publication of null results informs the wider research community that some effects are not present. We are aware that our paper is not the final word on whether the ego-relevance of the messages lowers/increases the psychological cost of deception. While we provide evidence that this feature does not affect deception in strategic interactions, it is still possible that in some other instances this happens. For example, this might happen when the messages are not anonymous and concern attractiveness (Gneezy et al., 2017), when the stakes are higher, or when the players are not fully aware of the fact that the senders profit from deception. However, we find the case of complete information to be a very interesting lower bound on the universe of possible deception environments. Moreover, we believe that in many real-life interactions, "receivers" can be almost sure that the "senders" have an incentive to deceive them. To conclude, we believe that we have contributed a relevant finding to this stream of literature. When the message on ego-relevance concerns cognitive abilities and the incentives to deceive are common knowledge, the ego-relevance of the state does not lower or increase deception costs. However, when the senders deceive the receivers about their positive cognitive abilities (they communicate a higher IQ than the receivers showed), they believe deception can be more persuasive.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data and code

Available here on Mendeley Data.

Acknowledgments

We thank very useful advice from the editor, two anonymous referees, Despoina Alempaki, Lajos Kossuth, Daniel Sgroi, and Renke Schmacker. This project also benefited from discussions with seminar participants at the PhD Forum (Economics — Warwick University), Brown Bag Seminar (WBS — Warwick University), IDEE meeting (Economics — Bologna University), and the SIN meeting. We thank Marco Gutierrez and Carlos Kong for their excellent work as research assistants. We also thank Jose Arellano, Kristian Lopez Vargas, and Alejandro Lugon for giving us access to the LEEX-PUCP laboratory. Financial support by the Leverhulme Trust (The Bridges Programme), United Kingdom and the Experimental and Behavioral Economics Research Group at Warwick University are gratefully acknowledged. Ethical approval (ID number: ECONPGR04/19) was granted by the Warwick Ethics Committee in May 2020.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.socec.2022.101920.

References

Abeler, J., Nosenzo, D., & Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, *87*(4), 1115–1153.
Alan, S., Ertac, S., & Gumren, M. (2020). Cheating and incentives in a performance context: Evidence from a field experiment on children. *Journal of Economic Behaviour and Organization*, *179*, 681–701.
Barber, B. M., & Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment*. *Quarterly Journal of Economics*, *116*(1), 261–292.
Behnk, S., Barreda-Tarrazona, I., & García-Gallego, A. (2019). Deception and reputation–An experimental test of reporting systems. *Journal of Economic Psychology*, *71*, 37–58.
Blume, A., Lai, E. K., & Lim, W. (2020). Strategic information transmission: A survey of experiments and theoretical foundations. In *Handbook of experimental game theory*. Edward Elgar Publishing.
Bolger, F., Hasker, K., & Murad, Z. (2019). Overconfidence: On its nature, universality, origins and relationship with performance. Working paper.
Brodeur, A., Cook, N., & Heyes, A. (2020). Methods matter: P-hacking and publication bias in causal analysis in economics. *American Economic Review*, *110*(11), 3634–3660.
Cai, H., & Wang, J. T.-Y. (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior*, *56*(1), 7–36.
Capraro, V., Perc, M., & Vilone, D. (2019). The evolution of lying in well-mixed populations. *Journal of the Royal Society Interface*, *16*(156).
Chen, D. L., Schonger, M., & Wickens, C. (2016). OTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, *9*, 88–97.
Chopra, F., Haaland, I., Roth, C., & Stegmann, A. (2022). The null result penalty. Working paper.
Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica*, *50*(6), 1431–1451.
Danz, D., Wilson, A. J., & Vesterlund, L. (2020). Belief elicitation: Limiting truth telling with information or incentives. CESifo Working Paper.
Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., & Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, *9*(3), 522–550.
Dreber, A., & Johannesson, M. (2008). Gender differences in deception. *Economics Letters*, *99*(1), 197–199.
Erat, S., & Gneezy, U. (2012). White lies. *Management Science*, *58*(4), 723–733.
Francis, G. (2013). Replication, statistical consistency, and publication bias. *Journal of Mathematical Psychology*, *57*(5), 153–169.
Gawn, G., & Innes, R. (2019). Lying through others: Does delegation promote deception? *Journal of Economic Psychology*, *71*, 59–73.
Gerlach, P., Teodorescu, K., & Hertwig, R. (2019). The truth about lies: A meta-analysis on dishonest behavior. *Psychological Bulletin*, *145*(1), 1.
Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, *95*(1), 384–394.
Gneezy, U., Gravert, C., Saccardo, S., & Tausch, F. (2017). A must lie situation – Avoiding giving negative feedback. *Games and Economic Behavior*, *102*, 445–454.
Ho, T.-H., & Yeung, C. (2014). Giving feedback to clients. *Management Science*, *60*(8), 1926–1944.
Hossain, T., & Okui, R. (2013). The binarized scoring rule. *Review of Economic Studies*, *80*(3), 984–1001.

Jin, G. Z., Luca, M., & Martin, D. (2021). Is no news (perceived as) bad news? An experimental investigation of information disclosure. *American Economic Journal: Microeconomics*, *13*(2), 141–173.

Kühberger, A., Fritz, A., & Scherndl, T. (2014). Publication bias in psychology: A diagnosis based on the correlation between effect size and sample size. *PLoS One*, *9*(9).

Mitra, A., & Shahriar, Q. (2020). Why is dishonesty difficult to mitigate? The interaction between descriptive norm and monetary incentive. *Journal of Economic Psychology*, *80*.

Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, *115*(2), 502.

Rosenbaum, S. M., Billinger, S., & Stieglitz, N. (2014). Let's be honest: A review of experimental evidence of honesty and truth-telling. *Journal of Economic Psychology*, *45*, 181–196.

Schotter, A., & Trevino, I. (2014). Belief elicitation in the laboratory. *Annual Review of Economics*, *6*(1), 103–128.

Siniver, E., & Yaniv, G. (2019). Optimism, pessimism, mood swings and dishonest behavior. *Journal of Economic Psychology*, *72*, 54–63.

Stanley, T. D. (2005). Beyond publication bias. *Journal of Economic Surveys*, *19*(3), 309–345.

Sternberg, R. J., Grigorenko, E. L., & Bundy, D. A. (2001). The predictive value of IQ. *Merrill-Palmer Quarterly*, *47*(1), 1–41.

Wang, J. T.-y., Spezio, M., & Camerer, C. F. (2010). Pinocchio's pupil: Using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American Economic Review*, *100*(3), 984–1007.