Essays in Behavioural Economics

by

Sergio Alessandro Castagnetti

Thesis

Submitted to the University of Warwick

for the Degree of

Doctor of Philosophy in Economics

Department of Economics

December 2021

Contents

List of Figures

List of Tables

Acknowledgements

This thesis could not have been written without the invaluable support of my supervisors. I thank Robert Akerlof for his constant support and insightful feedback on my research. I also thank Eugenio Proto who taught me how to do research on behavioral and experimental economics. Finally, I thank Sebastiano Massaro who made me understand the importance and relevance of doing interdisciplinary research. I also want to thank my coauthors from whom I learned different skills on how to do better research: Alessandro Bucciol, Giovanni Burro, James Fenske, Renke Schmacker, Karmini Sharma, and Jan Potters. Likewise I thank the faculty at the University of Warwick. Special thanks go to my PhD colleagues, and in particular, Roberto Asmat and Lajos Kossuth. Finally, I want to thank my family and friends for their support throughout this journey.

I am grateful to the Department of Economics of the University of Warwick which allowed me to pursue my research interests. I also thank the Leverhulme Trust for the financial support, through a Bridges Doctoral Scholarship. I am thankful to Andy Schotter for hosting me at the Center for Experimental Social Science of the New York University.

Declaration

This thesis is submitted to the University of Warwick in accordance with the requirements of the degree of Doctor of Philosophy in Economics. I declare that it has not been submitted for a degree at another university. I am the sole author of Chapter 1. Chapter 2 is co-authored with Karmini Sharma (PhD candidate at the University of Warwick) and James Fenske (Professor at the University of Warwick). Chapter 3 is co-authored with Eugenio Proto (Alec Cairncross Professor at the University of Glasgow). In both cases, all authors contributed equally to the design of the experiment, the analysis of the data, as well as the writing of the paper.

December 2021

Abstract

This thesis consists of three essays on Behavioural Economics. The first two chapters study how individuals form beliefs and whether they systematically end up with biased (posterior) beliefs due to psychologically motivated biases (e.g. the desire to hold positive views about oneself) and cognitive biases (e.g. cognitive failures) in information processing. The third chapter studies how (negative) emotions drive behavior. In particular, it investigates the implications of anger and sadness on strategic reasoning and performance.

In "the Ego is No Fool: Absence of Motivated Belief Formation in Strategic Interactions" (Chapter 1), I use an online experiment to investigate whether individuals are more easily fooled by others when they enhance their personal characteristics and abilities. Literature in economics and psychology suggests that individuals may want to believe good news about themselves, even if it comes from people who will gain economically from inducing such beliefs. I use an experiment in which participants complete an IQ test and then play a sender-receiver game. I find that receivers are not more likely to believe senders when they provide news that carries positive information about their IQ, compared to the cases in which the news carries no ego-relevant information or negative information about themselves. These results show that the desire to form favorable beliefs about oneself does not make individuals blind to the motives of the person who sends the information.

In "Attribution Bias by Gender: Evidence from a Laboratory Experiment" (Chapter 2), I conduct a laboratory experiment to study whether principals are prone to attribution bias by gender (i.e., if they reward male agents for good luck, while punishing female agents for bad luck). In the experiment, agents perform tasks for the principals and the realized outcomes depend on both the agents' performance and luck. Principals then assess agents' performances and decide what to pay the agents and are asked their beliefs. Our experimental results do not show evidence consistent with attribution bias by gender. While principals' payments and beliefs about agent performance are heavily influenced by realized outcomes, they do not depend on the gender of the agent.

In “Anger Impairs Strategic Behavior: A Beauty-Contest Based Analysis” (Chapter 3), I look at whether anger is a credible commitment device because it limits the capacity for strategic reasoning. In the lab experiment, I externally induce anger in a subgroup of subjects following a standard procedure (treatment) and no emotion in the other subgroup (control). Results show that angry subjects choose numbers further away from the best response level and earn significantly lower profits in a beauty contest game, compared to subjects in the control. This suggests that anger does indeed impair the individual's capacity to think strategically. Moreover, in a second experiment, I find that this effect is not common to all negative emotions: sad subjects do not play significantly further away from the best response level than the control group.

Chapter 1

The Ego is No Fool: Absence of Motivated Belief Formation in Strategic Interactions

I investigate whether individuals are more easily fooled by others when they enhance their personal characteristics and abilities. Individuals may want to believe good news about themselves, even if it comes from people who will gain economically from inducing such beliefs. In fact, economics and psychology literature suggests that motivated belief formation may shape economic interactions. I use an experiment in which participants complete an IQ test and then play a sender-receiver game. The experiment has a 2x2 factorial design. First, I determine the state either by the receiver's relative performance or by 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) rank, while the receiver benefits from selecting the action that matches his true rank. I find that receivers are not more likely to believe senders when they provide news that carries positive information about themselves, compared to the cases in which the news carries no ego-relevant information or negative information about their IQ. These results reveal important boundary conditions to motivated belief formation.

1.1 Introduction

There is mounting evidence that belief formation is subject to motivated biases: individuals process information in ways that serve their ego. What is less known is whether the desire to hold positive views about oneself shapes interactions with economic and financial consequences. In particular, individuals might be more easily fooled if others flatter their personal characteristics to get an economic advantage. For instance, an entrepreneur might be contemplating taking out a bank loan. A bank intermediary might be willing to convince him to request such loan by making him believe that he has the skills and abilities to succeed in the business. The intermediary's reward is a year-end bonus based on the number of new loans issued. If the entrepreneur does not take into account the economic incentives of the bank intermediary, he might be too credulous about the praise for his abilities. He might take out the loan because of his inaccurately instilled belief that he has the skills and abilities to start the business.

Understanding whether motivated beliefs shape social and economically relevant interactions is crucial. Indeed, while previous literature has shown that individuals process self-serving ego-relevant information that comes from objective and impartial sources, little is known about how individuals process such information when it comes from people they are transacting with. Specifically, we do not know how they process information, when their counterpart profits if the former engages in motivated belief formation. If the individual fails to account for potential bias in the information received, because he wants to believe well of himself, there could be many economically relevant implications. For instance, motivated biases may lead to suboptimal market equilibria in which individuals are constantly fooled about their personal qualities and skills. This may also help explain why individuals are generally overconfident about their personal characteristics (Moore and Healy, 2008). If, on the other hand, individuals acknowledge that information may be biased because of others' incentives and are skeptical about, we might be able to conclude that motivated information processing has bounds that limit its impact on economic interactions.

In this paper, I provide first evidence that motivated belief formation does not affect economically relevant interactions. I specifically investigate whether individuals are more easily fooled when others positively and strategically praise the personal characteristics that they care about in an exaggerated way. While this dynamic may be prominent in many economic interactions, it is difficult to identify it cleanly. An experiment in a controlled environment is useful for assessing motivated belief formation and its economic implications.

I conduct a simple experiment with an economic interaction in which: 1) there is social transmission of ego-relevant information; 2) incentives are misaligned between who sends and who receives the information; and, 3) the outcome variables make it possible to study the influence of the information transmitted on actions and beliefs. In the experiment subjects play a sender-receiver game. The sender sends a message about the state of the world to the receiver, who then takes an action. While the sender profits most when the receiver takes a specific action, the receiver's optimal action is to match the state. The experiment has a 2x2 factorial design. I vary whether the state of the world is determined by the receivers' performance in an IQ test (ego-relevant condition), or by a random number (non-ego-relevant condition). I then vary whether the sender's incentive is for the receiver to take an action that corresponds to him being of a high rank (positive condition) or of a low rank (negative condition). Before and after the game I elicit beliefs about the receiver's relative ranking.

The experimental results show that receivers are not more likely to follow and believe the messages from senders when they carry good news (or bad news) about their relative ability. In fact, while (as expected) news has a strong impact on actions played and on receivers' posterior beliefs about their rank, it does not have any differential impact by the ego-relevance of the news. Additional analyses further confirm that there are no systematic differences across experimental conditions. However, it is important to point out that the experimental results cannot exclude that there might be relatively small effect sizes of ego-relevance in strategic settings. Nevertheless, the experimental results suggest that motivated belief formation is bounded, and while individuals desire to hold positive views about themselves, it does not make them fully oblivious to the economic context of the information exchange. In fact, individuals in the experiment account for the strategic incentives of who is sending the information irrespective of whether the information is ego-relevant or not.

This paper contributes to: the literature on motivated cognition in psychology, specifically, the experimental evidence on how information is processed in light of self-enhancement motives; the literature in economics on motivated beliefs about ego-relevant personal characteristics, and the literature on communication games in economics.

An extensive literature in psychology shows that individuals have a basic desire to believe good things about themselves (self-enhancement motive), while protecting themselves against having negative self-views (self-protection motive).[1] There are many possible mechanisms through which individuals can engage in motivated reasoning, one of which is information processing. There is experimental evidence that information that is consistent with a preferred conclusion is examined less critically than information that is not (Ditto and Lopez, 1992; Kunda, 1990; Pyszczynski

[1] For a comprehensive review and a summary of its emergence in the psychological literature, see Alicke and Sedikides (2009).

and Greenberg, 1987). Similarly, experimental evidence show that individuals tend to accept positive statements about the self without giving much thought to the motives of the person making such statement (Vonk, 2002). These studies strongly suggest that individuals may be misled by others about their personal characteristics in economic interactions. I contribute to this literature by analysing the effects of motivated reasoning where there is flattery and ingratiation in a game that captures important features of many economically relevant interactions.

A relatively recent literature in economics has drawn interest and inspiration from the psychological evidence on motivated cognition and has studied closely how individuals process ego-relevant information. Theoretical work has emphasized how ego motives may affect the way people process information (see, e.g., Bénabou and Tirole (2016)). Moreover, experimentally many of these advanced mechanisms have been shown to fuel overconfidence about individuals' personal characteristics. These include: selective recall, motivated errors, and asymmetric updating. Usually these papers look at how individuals process ego-relevant information that is provided by an objective and precise mechanism. The main finding is that they process ego-relevant information self-servingly. That is, subjects are more likely to remember positive than negative performance feedback, they are more likely to commit (motivated) mistakes to reach more flattering beliefs about themselves, and they tend to update more strongly to positive than negative signals about their ability.[2] This paper is closest in spirit to the asymmetric updating literature since I study individuals' reaction to ego-relevant news. To this literature I crucially add the social exchange of ego-relevant information in a strategic environment.[3,4]

I also contribute to the experimental literature on communication experiments. In particular, 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 I implement

[2] For evidence on selective recall see Chew, Huang and Zhao (2018) and Zimmermann (2020), while for evidence on motivated errors see Exley and Kessler (2018). The evidence on asymmetric updating is less compelling. While initially Eil and Rao (2011) and Möbius et al. (2014) found evidence of asymmetric updating, other papers have failed to find it (Buser, Gerhards and van der Weele, 2018; Coutts, 2019; Ertac, 2011; Schwardmann and Van der Weele, 2019). However, more recent papers have identified conditions under which asymmetric updating is more likely to arise (Castagnetti and Schmacker, 2020; Coutts, Gerhards and Murad, 2019).

[3] In not strategic settings, Oprea and Yuksel (2020) studies how individuals jointly update beliefs about their IQ performance, while Gneezy et al. (2017) examine the conditions under which people provide accurate feedback to others about their physical appearance. In this setting, the paper finds suggestive evidence that individuals update their beliefs in a self-serving fashion.

[4] Schwardmann and Van der Weele (2019) and Solda et al. (2019) show experimentally that individuals' level of (over)confidence is shaped by whether it helps in social interactions. Using a different approach, I look at whether one is deceived by others because of the ego-relevance of the messages.

here is based on a cheap talk game that borrows features from Cai and Wang (2006) and Wang, Spezio and Camerer (2010).[5] In brief, the experimental findings are that senders reveal more information and receivers react more to the messages than predicted by the theory (Crawford and Sobel, 1982). Relatedly, in similar experimental settings (e.g., Sánchez-Pagés and Vorsatz (2009); Serra-Garcia, Van Damme and Potters (2011)) and in settings in which lying is not permitted but subjects can withhold or make the information transmitted more complex to interpret (e.g., Jin, Luca and Martin (2015, 2018)), there is ample evidence that receivers make inferential mistakes when assessing the senders' messages. That is, they are insufficiently skeptical to false, empty, vague, or complex messages. I add to this literature the study of receivers' behavior in communication games where the state is ego-relevant, and study how it affects behavior compared to the standard case where the state is not ego-relevant.[6]

The closest paper to mine is Ho and Yeung (2014). Their experiment 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. Ho and Yeung (2014) find that agents inflate the feedback and clients report higher levels of happiness. While there are many common features with my experiment, there are also significant differences, of which two stand out. First, in my experiment the incentives are deliberately misaligned across the two roles. This allows me to study whether motivated reasoning can cause one party to be insufficiently skeptical in settings where subjects have conflicting interests. Second, I look at the effect of messages holding prior beliefs constant. In this way, I am able to learn whether individuals are more easily fooled upon receiving good news about their personal characteristics and that cannot, therefore, be accounted by their initial (over)confidence levels.[7]

Overall, I contribute a single, important finding to the literature. Despite extensive literature in both economics and psychology that shows that individuals engage in motivated reasoning, my results show that this effect is constrained by the envi-

[5] For a comprehensive literature on cheap talk games, see Blume, Lai and Lim (2020).

[6] The study of senders' behavior has started a literature on deception and lying aversion (Gneezy (2005); Erat and Gneezy (2012)). In this paper, I focus on receivers' actions and, therefore, I do not analyze senders' behavior. However in a companion paper (Castagnetti and Burro, 2021) I look at deception rates by ego-relevance of the state. In short, I do not find any evidence that lying depends on the ego-relevance of the state.

[7] Examples of other important differences are the following. I look at relative performance (i.e., receivers' IQ performance relative to the performance of other individuals) instead of absolute performance (i.e., the number of questions solved correctly by the receiver). Arguably, the former is more ego-relevant than the latter. I also study the effect of messages on beliefs. In particular, I am able to do this since I elicit prior and posterior beliefs.

ronment. There are limits to motivated reasoning that prevent people from forming biased beliefs about themselves. In particular, in the case of this experiment, receivers react to their environment by realizing others' stakes in the game and not internalising the positive messages senders send. These boundary conditions of motivated reasoning are explained by Bénabou and Tirole's (2016) framework and what they refer to as the constraints of reality. By this they mean that while individuals are willing to engage (consciously or not) in motivated reasoning (the demand of motivated beliefs), their ability to do so is not infinite. In fact, they are constrained by the environmental cues (e.g. the supply side) limiting their ability to form high beliefs of themselves.

The remaining of the paper is organized as follows. In Section 1.2, I provide a detailed description of the experiment. In Section 1.3, I present the experimental hypotheses, while in Section 1.4 I show the results. Finally, in Section 1.5 I discuss the results and conclude.

1.2 Experimental Design

To causally investigate whether subjects are more easily deceived when others inflate their personal characteristics, an experiment with the following features is required. First, a game with at least two players and the ability to transmit messages. Second, an action that measures ones' propensity to follow the messages and 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). This environment can be created in a laboratory. There are two parts to the experiment. Subjects are asked to complete an IQ test and then play a sender-receiver game. Senders are informed about the state of the world and send a message about the state to receivers, who then take an action. 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 2x2 between-subject design. First, I 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). Second, I vary the payoffs in the game. In the positive (negative) condition, senders profit the most when the receivers take a high (low) ranked action in the game, corresponding to them being in the top (bottom) of the rank.[8] Before and after the

[8] I will call a "high ranked action" (or "high ranked message") an action (message) that cor-

game, I elicit receivers' beliefs about their ranking, and senders' beliefs about the receivers' beliefs. Figure 1.1 shows the timeline of the experiment.

Figure 1.1: Timeline of the Experiment

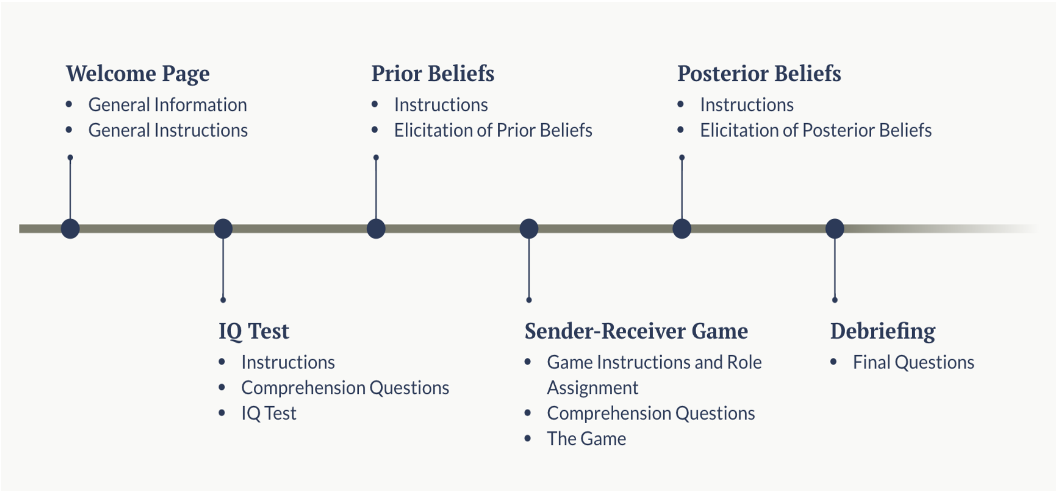

responds to rankings closer to the top. Conversely, for "low ranked action" (or "low ranked message").

1.2.1 The IQ Test

The experiment started with an IQ test, the Raven Advanced Progressive Matrices (APM) test. I administered 20 matrices from Set II of the APM. This set is appropriate for adolescents and adults of average intelligence, because it differentiates across the entire range of adult ability.[9] In each question, subjects were shown a 3x3 matrix of pictures with the one in the bottom right corner missing and asked to find the image (out of 8 possible choices presented below the matrix) that completes the pattern. Figure 1.2 shows one example.

Figure 1.2: Raven Matrix Example

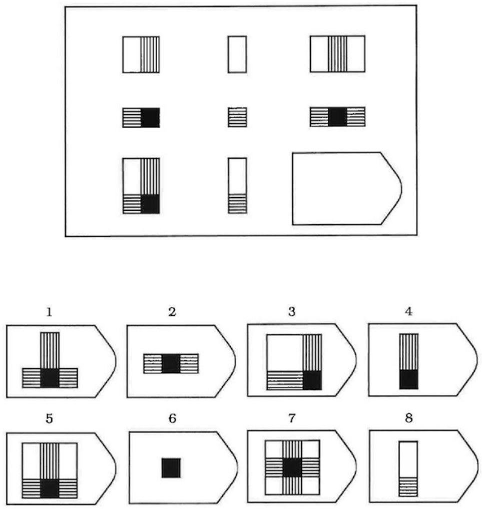

Notes: The figure displays matrix 11 from Set II of the APM test. Image number 5 completes the pattern.

Subjects were given detailed instructions. They had 10 minutes to answer the 20 questions. They could be answered in any order and answers could be changed, within the time limit. Financial rewards are not usually given with this test, but I decided to pay subjects S/. 5.00 per correct answer out of three randomly chosen questions.[10] I did this to increase subjects' motivation to perform well. This meant that poor performance in the test could not be ex-post rationalized by lack of attention, effort, or willingness to perform well.

I shared true information that this test is often used to measure fluid intelligence (i.e., reasoning ability) and general intelligence, and that high scores in this test correlate highly with economic variables (e.g., income and occupation) and health variables (e.g., health quality and longevity) (see Sternberg, Grigorenko and Bundy

[9]This is particularly suitable for university students, who are on average of high IQ ability.
[10]At the time of the experiment (June, 2020), the exchange rate was: $ 1.00 = S/. 3.54).

(2001)). This was accomplished to increase the ego-relevance of the task.

1.2.2 The Sender-Receiver Game

After the test, subjects were randomly and evenly sorted into: senders and receivers. The software randomly created sender-receiver pairs that played the sender-receiver game.[11] In the game, the sender is informed about the state of the world (the "realized state"), which can take any value in the state space, $S_s = \{1, 2, 3, ..., 10\}$, 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: $M_s = S_s = \{1, 2, 3, ..., 10\}$. The receiver chooses an action $a_r \in A_c = \{1, 2, 3, ..., 10\}$, after receiving the message. Payoffs are determined by both the receiver's action in the game and the state of the world. Figure 1.3 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.

Figure 1.3: The Sender-Receiver Game

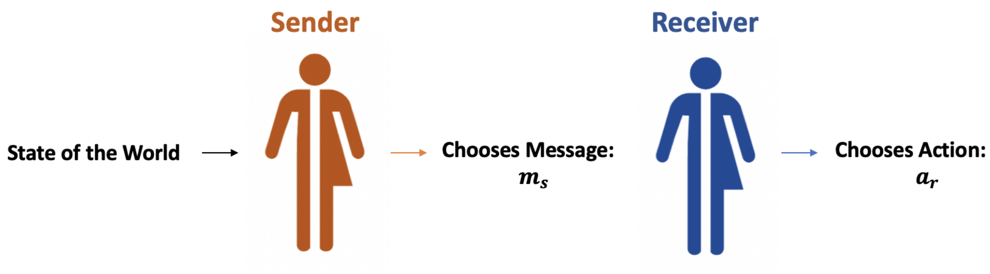

Treatment Variations

To study causally whether individuals are more likely to be fooled when the news they hear positively enhances their ego, the experiment features a 2x2 factorial design. The factors correspond to variations in the ego-relevance of the state and how payoffs are determined in the game. I 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 9 other

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

subjects who took part in a pilot session. They were informed of this ranking procedure and that the scores elicited a strict ordering.[12] In the non-ego-relevant condition, the ranking was determined by the random draw. 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; and, 3) a negatively skewed distribution where lower rankings were drawn with higher probability. I varied the distributions to have exogenous variation in prior beliefs in the non-ego-relevant condition. In Appendix A.1, I provide a detailed description of the distributions.

With this experimental variation I could study whether receivers are more likely to follow high ranked messages (and "good" news) when the state is about their relative performance in the IQ test, compared to the case in which the state has been randomly determined.

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 higher ranked actions in the game. In the negative condition, payoffs were reversed: the sender's payoff monotonically increased as the receiver played lower ranked actions. In both conditions, therefore, 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 both 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. Figure 1.4 shows the payoff structure for both players and by condition.[13] Both players knew the payoff structure in the game and they were explicitly made aware of the misalignment of interests in game incentives across roles.

With this variation I can investigate the link between the ego-relevance of the state and being fooled. In particular, I can study whether there are asymmetric responses to negative news by ego-relevance of the state. These analyses are crucial as they will allow me to exclude other confounding effects that may be affecting differences in actions by the ego-relevance of the state. For instance, receivers might believe that senders are more trustworthy when the messages they send are about their personal characteristics. If, instead, receivers' actions are driven by ego-

[12] If two or more subjects had the same score, then it was randomly determined whose rank was higher.

[13] The payoffs for both senders and receivers are similar to the ones in Jin, Luca and Martin (2015) and Wang, Spezio and Camerer (2010).

Figure 1.4: Payoff Table by Condition

(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

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.

relevant motives, then we should expect the opposite predictions in the negative conditions: receivers in the ego-relevant treatment will be less likely to follow "bad" news, relative to the non-ego-relevant condition. Figure 1.5 provides a summary of the resulting experimental treatments.

Figure 1.5: Summary of Experimental Conditions and Corresponding Treatments

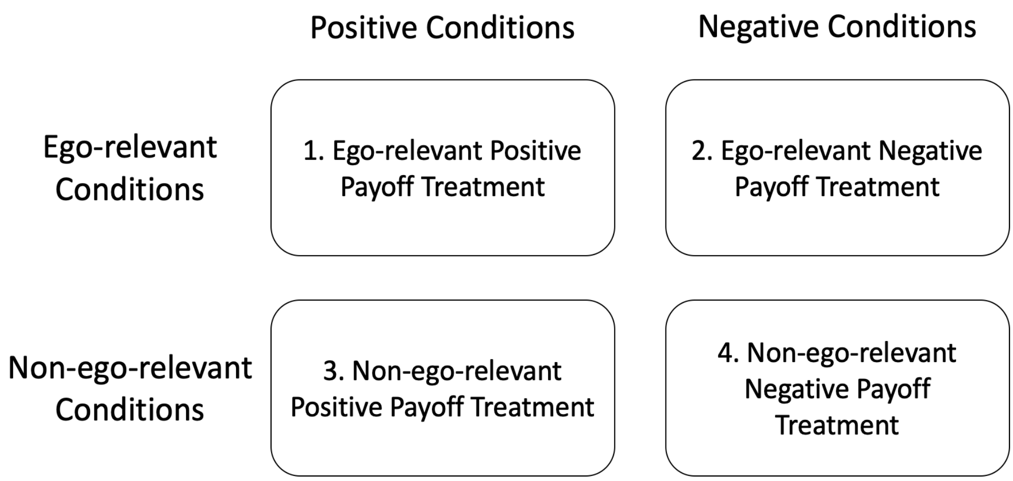

1.2.3 Prior and Posterior Beliefs

Before and after the sender-receiver game, I asked participants about the following set of beliefs.

Receivers' beliefs I asked receivers their prior (posterior) beliefs about their relative ranking before (after) the sender-receiver game.[14] In the ego-relevant condition this corresponded to their relative ranking in the IQ test. In the non-ego-relevant condition, it was determined by a random draw. I elicited the full distribution of these prior beliefs. That is, receivers had to write down their estimated probability of being in each of the 10 ranks.

Senders' beliefs I asked senders to report their beliefs about what their matched receivers thought their mean rank was before and after the game. Again, I elicited the entire distribution of the prior (posterior) beliefs.[15] Importantly, before the belief elicitation stage, senders were informed about their matched receiver's rank, how the rank was determined, and that the receivers did not know their true rank.

[14] At the time I asked participants their prior beliefs, they did not know what they would be doing in the next step of the experiment.

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

The elicitation of prior and posterior beliefs is crucial to the experiment. First, when analyzing game play, prior beliefs make it possible to define messages as carrying "good" or "bad" news. It will be also crucial to study receivers' actions in the game controlling for prior beliefs. Conversely, one could acknowledge differences in game play that ultimately are not driven by ego motives, but, instead, by differences in prior beliefs. Second, posterior beliefs allow me to analyze whether messages in the game influence not only actions but also beliefs about the receivers' rankings.

I used a financial incentive for the elicitation procedure. It consisted in the Binarized Scoring Rule proposed by Hossain and Okui (2013), and a fixed price of S/ 20.00. Under this method, truthful reporting is orthogonal to subjects' risk preferences and it does not rely on expected utility theory.[16] I explicitly and truthfully told participants that the elicitation mechanism guaranteed that it was in their best interest to report their true beliefs. I did not explain to subjects how the procedure worked, as withholding the description of the mechanism increases truthful reporting (see Danz, Wilson and Vesterlund (2020)).[17]

1.2.4 Debriefing

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. I 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.

1.3 Research Hypotheses

The experiment is designed to test whether individuals are more likely to follow "good" news that is ego-relevant, compared to the case in which the same news is not. In other words, the experimental conjecture is that individuals will be more easily fooled when they hear positive news about their personal characteristics, relative to the case in which the messages do not carry any ego-relevant content. Similarly, if ego-relevance drives behavior in the positive conditions, an opposite effect should emerge in the negative conditions. That is, "bad" news will not be followed as much

[16] For a detailed explanation of this elicitation procedure see also Schotter and Trevino (2014).
[17] The interested participants, however, could click on a button to read a detailed description of the elicitation method.

in the ego-relevant treatment compared to the non-ego-relevant negative treatment. The first hypothesis is therefore:

Hypothesis 1. *Individuals (receivers) in the positive ego-relevant treatment will be more likely to follow messages that carry "good" news, relative to those in the positive non-ego-relevant one. Conversely, individuals (receivers) in the negative ego-relevant treatment will be less likely to follow messages that carry "negative" news, relative to those in the negative non-ego-relevant one.*

Moreover, if individuals desire to interpret information in a self-serving way, the effects of messages might not be circumscribed to actions in the game, but they may affect (posterior) beliefs as well. In particular, individuals will interpret information in the ego-relevant treatments self-servingly (i.e., they will be more likely to update their beliefs following "good" news, while they will stick to their priors following "bad" news), while this effect will not hold in the non-ego-relevant conditions. Thus, the second experimental hypothesis is:

Hypothesis 2. *Individuals (receivers) in the positive ego-relevant treatment will be more likely to update their beliefs downwards (of being in a higher rank) following "good" news, relative to those in the positive non-ego-relevant one. Conversely, individuals (receivers) in the negative ego-relevant treatment will be less likely to update their beliefs upwards (of being in a lower rank) following "bad" news, relative to those in the negative non-ego-relevant one.*

1.4 Results

1.4.1 Implementation

The experiment took place in June 2020. I recruited subjects through the Orsee recruitment system and I used the pool of participants registered at the economics laboratory of Universidad Catolica del Perú in Lima, Perú. I recruited 514 participants, but after removing those participants who did not complete the experiment (mainly due to internet connection issues) and their matched partners, I have 504 participants.[18] I conducted eleven sessions (N=162) for the ego-relevant and positive treatment, six (N=104) for the non-ego-relevant and positive treatment, nine (N=134) for the ego-relevant and negative treatment, and six (N=104) for

[18] The experiment comprised another condition, not part of this paper, for which I recruited 122 additional subjects. This condition is part of a different paper and studies whether overconfidence in the IQ test increases by telling subjects that the test indeed measures IQ ability. The inclusion of this condition and its results does not change significantly any result in this paper.

the non-ego-relevant and negative treatment. On average, sessions lasted 45 minutes. Participants earned an average of S/. 12.00, including the show-up fee of S/. 5.00. I programmed and conducted the experiment in oTree (Chen, Schonger and Wickens, 2016).[19] Descriptive statistics of the sample of receivers are provided in Appendix A.2.

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

1.4.2 Prior Beliefs

I begin the analysis by describing prior beliefs.[20] Table 1.1 shows two moments of receivers' distribution of prior beliefs by treatment: the mean rank belief and the standard deviation. Overall, I find that subjects' mean rank belief is to be about rank five across treatments (mean rank prior belief = 5.19). I do not find significant differences in mean prior beliefs across them.[21] The mean standard deviation across treatments is 2.05; in the ego-relevant treatments they are lower than those in the non-ego-relevant ones. The differences are highly significant (p-values$<$0.01).[22] This finding is not surprising. In fact, receivers are likely to have more information about their own relative ranking in an IQ test compared to the realization of a random process. This translates to lower standard deviations in the distributions of IQ prior beliefs.

1.4.3 Sender-Receiver Game

I now analyze receivers' actions in the game. In Table 1.2, I start by providing summary statistics of the two main variables of interest by treatment: messages received and actions taken in the game. As expected in the positive payoff treatments, messages and actions are lower compared to messages and actions in the negative payoff treatments. There does not seem to be a difference in messages and actions

[19] Appendix A.4 shows the experimental instructions (translated from Spanish).

[20] From now on I will study receivers' behavior only. These analyses come from the 252 subjects who played in the role of receivers.

[21] Statistical significance is assessed by running a regression of the mean prior belief on the treatment variable with robust standard errors. The reference category is the positive ego-relevant treatment.

[22] Statistical significance is assessed by running a regression of the standard deviation on the treatment variable with robust standard errors. The reference category is the positive ego-relevant treatment.

Table 1.1: Receivers' Prior Beliefs by Treatment

	Mean prior belief	Std. Dev. prior distribution
Positive ego-relevant treatment	5.31 (0.17)	1.82 (0.06)
Negative ego-relevant treatment	5.15 (0.21)	1.85 (0.07)
Positive not-ego-relevant treatment	5.21 (0.14)	2.40 (0.06)
Negative not-ego-relevant treatment	5.02 (0.14)	2.31 (0.06)

Notes: Standard errors are shown in parentheses.

that are driven by the ego-relevance of the state. But to confirm this I carry out two analyses. First, I study whether the news in the messages influences actions differently by ego-relevance of the state in each payoff condition. Second, I look at how these relationships differ across payoff conditions in a unique estimation. In other words, I conduct a difference in differences analysis.

Table 1.2: Messages and Actions in the Game by Treatment

	Mean action	Mean message
Positive ego-relevant treatment	4.73 (0.23)	4.05 (0.24)
Negative ego-relevant treatment	5.52 (0.27)	6.90 (0.29)
Positive non-ego-relevant treatment	4.25 (0.27)	4.14 (0.28)
Negative non-ego-relevant treatment	5.69 (0.29)	6.67 (0.31)

Notes: Standard errors are shown in parentheses.

I therefore start by studying how subjects respond to messages that carry "positive" ("negative") news in the positive (negative) conditions. That is, I study whether responders in the positive (negative) payoff treatment are more (less) likely to follow "positive" ("negative") ego-relevant news compared to the same news when they are not ego-relevant. To perform these analyses, separately by positive and negative conditions, I run the following econometric specification:

$$\begin{aligned} Action_i = \beta_0 + \beta_1 news_i + \beta_2 treatment_i + \beta_3 news_i \times treatment_i \\ + \beta_4 rank\, prior_i + \beta_5 std.\, prior_i + x_i' \beta_6 + \epsilon_i \end{aligned} \tag{1.1}$$

Here, i is the receiver. $Action_i$ is the dependent variable and corresponds to receiver's i action in the game. $news_i$ is a dummy variable. It is equal to 1 if the receiver received "good" news in the positive payoff treatment or "bad" news in the negative payoff treatment. "Good" ("bad") news is defined as receiving a message about the state of the world that is strictly lower (higher) than the mean prior belief of the receiver i.[23] The $treatment_i$ variable is the treatment to which the receiver was randomly allocated. Then I use the following variables as controls. $rank\, prior_i$ is the receiver's mean rank prior belief and $std.\, prior_i$ is the standard deviation of the prior belief distribution. x_i is a vector of receiver's demographic variables (age, gender, risk preferences, and student status). I report robust standard errors in all specifications.

β_1 captures the effect of the news on the receiver's action. This corresponds to the effect of "expected" news on actions, given that receivers are aware of senders' strategic incentives. β_2 captures average differences in actions across the treatments. β_3 captures the interaction effect of (expected) news with being in the ego-relevant treatment. Then, β_4 and β_5 capture the effects of the mean rank belief and the standard deviation of the prior belief distribution on actions. β_6, is a vector of coefficients that captures the association between demographic variables and actions.

My main coefficient of interest is thus β_3. Following the conjectures of the previous section, I expect β_3 to be negative in both payoff conditions. In other words, I expect that subjects in the positive payoff conditions react more strongly to positive news when it is ego-relevant compared to when it is not, by playing lower actions in the game. Similarly, I expect subjects to react less strongly to negative news when it is ego-relevant in the negative payoff conditions.

The results of the estimation of Equation (1.1) are shown in Table 1.3. In columns (1)-(3), I report the results for the positive treatments, whereas in columns (4)-(6) I report those for the negative treatments. The different regressions, that differ in the number of control variables, show similar results. In the following, I will refer to the estimates of the most comprehensive models that are shown in columns (3) and (6). As expected, news has a strong effect on actions. Receiving good news decreases the action in the game by 2.212 ranks (p-value=0.001) in the positive conditions and increases it by 1.881 (p-value<0.001) in the negative conditions. The treatment variable increases actions in the positive conditions, while it decreases them in the negative conditions. However, the effects are not statistically significant.[24]

As per the main coefficients of interest, I find that they are small in magnitude and not distinguishable from zero. This is particularly true in the positive condition, where the estimated coefficient is equal to 0.043 (p-value=0.952), whereas in the negative condition it is larger in magnitude, -0.319, but still not significant (p-value=0.591). These initial results suggest that subjects are not more (less) likely to follow news that carries positive (negative) information about themselves.

[23]The mean of the *news* variable is 0.730, and its standard deviation is 0.445. It does not significantly vary by treatment.

[24]p-values are equal to 0.780 in the positive and to 0.936 in the negative conditions.

Table 1.3: Regression Results for Actions by Ego-relevance of the State

	Positive Condition			Negative Condition		
	(1)	(2)	(3)	(4)	(5)	(6)
	Action	Action	Action	Action	Action	Action
news	-1.417**	-2.043***	-2.212***	1.442***	1.827***	1.881***
	(0.557)	(0.567)	(0.625)	(0.544)	(0.501)	(0.512)
treatment	0.354	0.281	0.169	0.534	-0.196	-0.039
	(0.627)	(0.600)	(0.601)	(0.610)	(0.478)	(0.491)
news $\times$ treatment	0.099	-0.080	0.043	-0.899	-0.188	-0.319
	(0.738)	(0.697)	(0.712)	(0.772)	(0.601)	(0.593)
rank prior		0.681***	0.706***		0.919***	0.986***
		(0.128)	(0.132)		(0.112)	(0.118)
std. prior		-0.175	-0.174		-0.179	-0.300
		(0.319)	(0.319)		(0.292)	(0.308)
constant	5.286***	2.619**	2.944	4.583***	0.087	-0.378
	(0.475)	(1.011)	(1.923)	(0.423)	(0.971)	(1.413)
Demographics			✓			✓
R2	0.110	0.293	0.308	0.044	0.376	0.408
N	133	133	133	119	119	119

Notes: The table shows regression results of Equation (1.1). The regressions are estimated separately by positive and negative conditions. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The prior belief has a significant effect on actions. A one-unit increase in the rank prior belief is associated with a 0.706 (0.986) increase in actions in the positive (negative) condition. The standard deviation of the prior belief distribution, instead, is not significantly associated with actions in the game. I also do not find any significant association between the demographic variables and actions.

I now take advantage of the 2x2 factorial design and conduct a second analysis by running the following equation:

$$
\begin{aligned}
react_i = \beta_0 + \beta_1 news_i + \beta_2 treatment_i + \beta_3 payoff_i^{+} + \beta_4 treatment_i \times news_i \\
+\beta_5 payoff_i^{+} \times news_i + \beta_6 payoff_i^{+} \times treatment_i \\
+\beta_7 treatment_i \times payoff_i^{+} \times news_i \\
+\beta_8 rank\, prior_i + \beta_9 std.\, prior_i + x_i' \beta_{10} + \epsilon_i
\end{aligned}
\tag{1.2}
$$

There are two changes in this version of the model. First, now the dependent variable, $react_i$, is the difference between the mean prior rank belief and the action played in the game in the positive treatments. It is the opposite of the same difference in the negative payoff treatments. It thus captures how distant actions are from beliefs, or, in other words, how strongly subjects react to news. Second, this model adds the following variables. The $payoff_i^{+}$ variable, which is a dummy equal to 1 if the subjects belongs to the positive payoff condition, and 0 otherwise. The $payoff_i^{+} \times news_i$ which is the interaction between the news received and being in the positive payoff condition. The interaction between $payoff_i^{+}$ and the $treatment_i$ variables. Finally, the $treatment \times payoff_i^{+} \times news_i$ variable, which interacts the treatment and payoff variables with the news received in the game.

β_1 captures the effect of the news variable on the dependent variable. β_2 and β_3 capture the effect of both the ego-relevant and payoff treatments on $react_i$, while β_4 and β_5 capture whether there is a differential effect of news on the dependent variable in either the ego-relevant or payoff conditions. β_6 captures whether there is an interaction effect of being in the positive ego-relevant treatment. β_7 is the main coefficient of interest. It compares differences in reaction to positive news in the positive conditions (ego-relevant and not) with the difference in reaction to negative news in the negative conditions. Realize that this coefficient will be positive if subjects react: a) more strongly to positive ego-relevant news compared to positive non-ego-relevant news; b) less strongly to negative ego-relevant news compared to negative non-ego-relevant news; or c) a combination of these two. Thus, if this coefficient is not positive, this would strongly suggest that whether people follow the news received is not dependent on the ego relevance of the state.

Table 1.4 shows the results. From the first row of the table, it is clear that (expected) news has a strong and positive effect on the dependent variable. That is, in the positive payoff treatments, positive news lowers actions in the game from the mean prior belief, whereas in the negative payoff treatments, negative news increases actions in the game from the mean prior belief. The effect is of almost 2 ranks (p-value<0.000). However, if we look at the β_7 coefficient, we again find no evidence of asymmetric responses to ego-relevant news. In fact, the coefficient is small in magnitude, 0.175, and it is not statistically different from zero (p-value=0.852). These results further confirm that subjects are not more (less) likely to follow news when they carry positive (negative) news about their ability, compared to the same news when it is about some random number.

In Appendix A.3, I conduct further analyses to show the robustness of the results. In particular, I conduct alternative econometric specifications to assess the impact of ego-relevant messages on actions. I find similar results of no significant effects of ego-relevant messages on actions.

1.4.4 Posterior Beliefs

I now study whether the news received in the sender-receiver game have an impact on posterior beliefs. In particular, I study whether subjects interpret ego-relevant information self-servingly. In fact, even if subjects are not more (less) likely to follow positive (negative) news in the ego-relevant treatments, it still could be the case that their beliefs are shaped by ego-relevant information in a self-serving fashion. To study this, I now conduct the same analysis as in Equation (1.1) but by using the posterior beliefs as dependent variable. That is:

$$\begin{aligned} Rank\,post_i = \beta_0 + \beta_1 news_i + \beta_2 treatment_i + \beta_3 news_i \times treatment_i \\ + \beta_4 rank\,prior_i + \beta_5 std.\,prior_i + x_i'\beta_6 + \epsilon_i \end{aligned} \tag{1.3}$$

The coefficient of interest is again β_3 and if news influences beliefs differentially by ego relevance of the state, we should find the coefficient to be negative across all models. In fact, this would imply that news is processed in a self-serving way: subjects update their beliefs to a greater (lesser) extent when they hear positive (negative) ego-relevant news.

In Table 1.5, I report the results. As before, I refer to the estimates of the most comprehensive models that are shown in columns (3) and (6) and that include

Table 1.4: Regression Results for the React Variable by Ego-relevance and Payoff Treatments

	(1) React	(2) React	(3) React
news	1.869***	1.912***	1.943***
	(0.504)	(0.505)	(0.514)
treatment	-0.162	-0.284	-0.158
	(0.435)	(0.461)	(0.494)
payoff^{+}	-0.024	0.062	0.055
	(0.597)	(0.599)	(0.632)
news $\times$ treatment	-0.135	-0.047	-0.129
	(0.602)	(0.611)	(0.613)
news $\times$ payoff^{+}	0.518	0.386	0.400
	(0.750)	(0.760)	(0.801)
treatment $\times$ payoff^{+}	-0.186	-0.100	-0.210
	(0.735)	(0.737)	(0.777)
news $\times$ treatment $\times$ payoff^{+}	0.230	0.095	0.175
	(0.920)	(0.921)	(0.938)
rank prior		0.110	0.112
		(0.089)	(0.091)
std. prior		-0.092	-0.110
		(0.214)	(0.220)
constant	-0.763**	-1.136	-1.523
	(0.357)	(0.763)	(1.205)
Demographics			✓
R2	0.250	0.255	0.262
N	252	252	252

Notes: The table shows regression results of Equation (1.2). Robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.5: Regression Results for Posterior Beliefs by Ego-relevance of the State

	Positive Condition			Negative Condition		
	(1)	(2)	(3)	(4)	(5)	(6)
	Posterior	Posterior	Posterior	Posterior	Posterior	Posterior
news	0.026	-0.796***	-0.782***	-0.045	0.338	0.353
	(0.345)	(0.281)	(0.284)	(0.301)	(0.230)	(0.231)
treatment	0.025	0.060	0.075	0.476	-0.030	-0.001
	(0.442)	(0.333)	(0.321)	(0.456)	(0.204)	(0.206)
news × treatment	-0.139	-0.271	-0.240	-0.250	0.409	0.367
	(0.531)	(0.398)	(0.392)	(0.565)	(0.281)	(0.271)
rank prior		0.827***	0.829***		0.876***	0.900***
		(0.072)	(0.074)		(0.054)	(0.063)
std. prior		0.103	0.129		0.209	0.164
		(0.191)	(0.187)		(0.150)	(0.149)
constant	5.074	1.122**	1.705*	5.146***	-0.039	-0.537
	(0.277)	(0.544)	(0.992)	(0.219)	(0.381)	(0.593)
Demographics			✓			✓
R2	0.001	0.518	0.535	0.013	0.663	0.673
N	133	133	133	119	119	119

Notes: The table shows regression results of Equation (1.3). The regressions are estimated separately by positive and negative conditions. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

all variables as described in Equation (1.3). The table shows that in the positive condition β_3 is -0.240, while in the negative condition it is 0.367. Thus, I find only limited support of self-serving belief updating in the positive condition. However, both estimated coefficients are not statistically distinguishable from zero (p-values equal to 0.542 and 0.179, respectively). In sum, there is no (strong) evidence that ego-relevant news in the game affects beliefs differentially by the ego-relevance of the state.

Not surprisingly, the results also show that the mean prior belief is a strong and significant predictor of the posterior belief. A one-point increase in the mean rank prior is associated with a 0.9 increase in the posterior belief. Similarly, I find that the news variable shapes posterior beliefs and, therefore, the effect of the news is not limited to the game. In the positive condition, hearing good news is associated with a decrease in the action by about 0.8 units (p-value=0.007), while hearing bad news is associated with an increase in the action by about 0.353, although this latter effect is not statistically significant (p-value=0.129).

To conclude, I conduct a belief-updating analysis that is based on Equation (1.2). That is:

$$\begin{aligned} updating_i^r = \beta_0 + \beta_1 news_i + \beta_2 treatment_i + \beta_3 payoff_i^+ \\ +\beta_4 treatment_i \times news_i + \beta_5 payoff_i^+ \times news_i + \beta_6 payoff_i^+ \times treatment_i \\ +\beta_7 treatment_i \times payoff_i^+ \times news_i \\ +\beta_8 rank\, prior_i + \beta_9 std.\, prior_i + x_i'\beta_{10} + \epsilon_i \end{aligned} \tag{1.4}$$

The only difference with Equation (1.2) is the dependent variable, $updating_i$. It captures beliefs updating and is defined as the difference between the mean prior rank belief and the post rank belief in the positive payoff conditions, while it is the opposite of the same difference in the negative payoff conditions. The coefficient of interest is β_7. It captures how subjects react to positive news differentially by ego-relevance of the state in the positive payoff conditions, compared to how they react to negative news differentially by ego-relevance in the negative ones. A positive coefficient thus would suggest that individuals interpret information self-servingly to reach a belief that they have a higher IQ.

The results are shown in Table 1.6. Interestingly, I find that news shapes posterior beliefs to a greater extent when it carries ego-relevant information ($\beta_4 = 0.565$, p-value=0.045), however this does not depend on whether it carries positive or negative information. In fact, the results here confirm that there is no evidence in support of asymmetric updating. Indeed, the estimated coefficient of interest is even negative (β_7=-0.231), although not significantly different from zero (p-value=0.633).

In Appendix A.3, I conduct further analyses. These findings are consistent with the results shown in this section: the ego-relevance of the state does not influence the way subjects form their posterior beliefs.

1.5 Discussion and Concluding Remarks

1.5.1 Discussion

I now discuss potential threats to the interpretation of the results. In doing so, I also provide evidence of why they cannot convincingly account for the null effect of ego-relevance of the state on actions and (posterior) beliefs.

In the main specifications of the previous section, I have looked at how receivers

Table 1.6: Regression Results for the Updating by Ego-relevance and Payoff Treatments

	(1) Updating	(2) Updating	(3) Updating
news	0.382*	0.396*	0.381*
	(0.221)	(0.229)	(0.230)
treatment	-0.220	-0.241	-0.312
	(0.171)	(0.186)	(0.199)
payoff$^+$	-0.375	-0.347	-0.346
	(0.237)	(0.240)	(0.239)
news $\times$ treatment	0.514*	0.539*	0.565**
	(0.283)	(0.283)	(0.280)
news $\times$ payoff$^+$	0.562*	0.515	0.508
	(0.336)	(0.359)	(0.373)
treatment $\times$ payoff$^+$	0.202	0.224	0.282
	(0.375)	(0.374)	(0.382)
news $\times$ treatment $\times$ payoff$^+$	-0.182	-0.212	-0.231
	(0.484)	(0.481)	(0.483)
rank prior		0.033	0.042
		(0.046)	(0.050)
std. prior		0.005	-0.002
		(0.120)	(0.120)
constant	-0.201	-0.386	-0.921
	(0.126)	(0.343)	(0.615)
Demographics			✓
R2	0.173	0.175	0.183
N	252	252	252

Notes: The table shows regression results of Equation (1.4) in two subsamples of interest. The regressions are estimated separately by positive and negative conditions. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

react to news depending on the ego-relevance of the state. It could be the case that subjects react differently to *messages* in the game and not to *news* itself. However, as shown in Appendix A.3, econometric analyses of the effect of messages on actions and beliefs find similar results. In sum, the main results are not driven by the way in which I analyze the information that is being socially transmitted.

Importantly, receivers' beliefs about the truthfulness of the messages received might vary with the ego-relevance of the state. For instance, it could be the case that receivers think senders are more honest about the information they send when the state is ego-relevant.[25] This could explain why in the negative payoff conditions, I do not find that subjects respond less to negative news that is ego-relevant. However, the 2x2 factorial design remedies this concern. In fact, if this effect is present, we should also expect that even in the absence of motivated reasoning, receivers are more likely to follow positive news that is ego-relevant in the positive payoff conditions. This is not what results show, however. In sum, the null results cannot be explained by receivers' differential expectations about the truthfulness of news by ego-relevance of the message.

It is often problematic to disentangle cognitive biases (i.e., due to cognitive constraints and limitations) from motivated biases. In particular, it is difficult to unravel the effects of confirmation bias (i.e., the tendency to put more weight on information that confirms one's prior beliefs relative to information that contradicts them) from those of motivated belief formation.[26] In this paper, confirmation bias cannot be confounded with motivated reasoning. Indeed, the research questions (and experimental analyses) are related to whether subjects are more (less) likely to believe "good" ("bad") news, which is defined as positive (negative) deviations from prior beliefs. Thus, in this environment there is little room for confirmation bias.

Finally, the ego-relevant treatments are about IQ ability. One may wonder whether the results in the ego-relevant treatments are dependent on IQ ability. In particular, it is possible that individuals of higher IQ ability are more likely to receive "good" news, or vice versa. Whereas in the non-ego-relevant treatments this is likely not to be the case. This would result in spurious treatment comparisons. Two pieces of evidence show that this is not the case. First, I do not find strong evidence that the type of news received is significantly shaped by receivers' IQ scores.[27]

[25]There is some evidence of this effect in Table 1.4. In fact, subjects update their beliefs to a greater extent in the ego-relevant treatments.

[26]This is particularly true in the asymmetric updating literature where the experimental signals are either positive or negative. The coarse structure of the signal structure makes it more difficult to disentangle confirmation bias from motivated beliefs.

[27]In fact, when I run separate regressions for each condition of news received on receivers IQ scores, the coefficients are small in magnitude and not statistically significant throughout.

Second, if I repeat all the analyses controlling for IQ scores, the main results and conclusions hold.

One important concern that limits the interpretation of the results is however statistical power. In particular, at a significance level of 5% and with 80% power, the minimum detectable effect size is relatively large. However, it is anyway lower than the effect of news on actions and beliefs.[28] Thus, although not precise, the experimental results are informative as they allow to put bounds on the effect of ego-relevant news on strategic interactions.

1.5.2 Concluding Remarks

Theoretical work in economics has put forward reasons why subjects may end up with overconfident beliefs about their personal characteristics. These include self-esteem concerns (Bénabou and Tirole, 2002; Köszegi, 2006) and anticipatory utility (Brunnermeier and Parker, 2005). Recent experimental literature has provided evidence of different mechanisms that allow individuals to interpret ego-relevant information self-servingly. Yet, in most of this work, information processing takes place in abstract settings in which there is no social exchange of information. Thus, there is much scope for work that studies the implications of ego-relevant concerns in economically relevant interactions.

In this paper, I conduct an experiment that fills this gap and that presents key features of many real-world economic relationships. In particular, here I study whether ego concerns leave individuals more easily fooled by others in a setting of strategic information transmission. To study this research question, the experimental design varies two relevant dimensions (ego-relevance and payoff incentives), resulting in a 2x2 factorial design. The findings of the experiment clearly show that the recipients of information are not more likely to be fooled when the messages they hear provide good news about their relative ability. In fact, subjects' actions in the game and their belief formation do not diverge depending on the ego-relevant content of the news, nor on whether the news carries positive or negative content. One limitation of the study is however power, which precludes the possibility of ruling out small effect sizes of ego-relevance in strategic settings.

On the whole, this research brings a new perspective to the literature on motivated belief formation. In particular, it shows that the desire to form favourable beliefs about oneself does not make individuals blind to the motives of the person

[28] The minimum detectable effect is calculated as follows: $(t_{1-\kappa}+t_{\alpha})\times SE_{\hat{\beta}}$. For example, for the estimation of equation (1.1) regarding the positive condition (column 3), it is $(0.84+1.96)\times 0.712 \approx 2$.

who sends the information. This is important. In fact, it shows limits to motivated belief formation and, specifically, that this previously documented force does not readily translate to strategic economic interactions. The findings presented here, therefore, are consistent with Bénabou and Tirole (2016)'s framework in which the supply side of motivated beliefs creates reality constraints that prevent individuals from freely forming their desired beliefs. Future research should study the implications of motivated belief formation in other economic interactions with social exchange of ego-relevant information, to better understand how the demand and supply sides of motivated beliefs interact.

Chapter 2

Attribution Bias by Gender: Evidence from a Laboratory Experiment

In many settings, economic outcomes depend on the competence and effort of the agents involved, and also on luck. When principals assess agents' performance they can suffer from attribution bias by gender: male agents may be assessed more favorably than female agents because males will be rewarded for good luck, while women are punished for bad luck. We conduct a laboratory experiment to test whether principals judge agents' outcomes differently by gender. Agents perform tasks for the principals and the realized outcomes depend on both the agents' performance and luck. Principals then assess agents' performance and decide what to pay the agents. Our experimental results do not show evidence consistent with attribution bias by gender. While principals' payments and beliefs about agent performance are heavily influenced by realized outcomes, they do not depend on the gender of the agent. Our evidence suggests that the interaction between the gender of the principal and the agent plays a role. In particular, principals are more generous to agents of the opposite gender.

2.1 Introduction

In many economic settings, outcomes depend on dispositional factors such as effort and ability, as well as on situational factors, such as luck. This creates room for attribution bias. Attribution bias is the tendency for people to under-emphasize situational explanations for outcomes while over-emphasizing dispositional explanations (Ross, 1977). Attribution bias by gender is understood as the tendency of observers to attribute successes to ability for males and to luck for females, and to attribute failure to luck for males and to ability for females (Deaux and Emswiller, 1974). Two strands of literature in social psychology have investigated attribution bias by gender. One has focused on how men and women differ in accounting for their *own* successes or failures, and has found that men are more likely to attribute their own successes to ability while women are more likely to attribute their failures to ability (McMahan, 1982; Stipek and Gralinski, 1991). The other strand studies attribution of success and failure to *others*, and has found mixed evidence on whether observers are more likely to attribute men's successes in some tasks to ability or more likely to attribute their failures to luck, compared to women (Hill and Augoustinos, 1997; Räty et al., 2002). This literature has not focused on cases where outcomes realized by the individual being evaluated affect the payoff of the individual making the evaluation, but recent empirical evidence in economics suggests that there may be attribution bias by gender in such contexts. These include referrals to surgeons after the death of a patient (Sarsons, 2019), executive pay in the finance sector (Selody, 2010), firing of corporate executives (Landsman, 2019), and punishment for misconduct (Egan, Matvos and Seru, 2017). However, there may be other variables in these real-world environments that cannot be controlled for, including agents' real contributions to outcomes, prior experience, and unobserved characteristics. Factors other than attribution bias might drive these differences in outcomes by gender.

We present evidence from a laboratory experiment that tests the presence of attribution bias by gender in a controlled environment. We use a controlled setting so that other factors are less likely to influence participants' behavior. We use a principal-agent setup. Participants are first randomly divided into two roles: principals and agents. In each round (out of 20), they are randomly matched into pairs. Agents perform one of four tasks for their principals in each round: a maths task, a Raven task, an effort task, or a memory task. There is a random component in how agents' performance produces the outcome. Principals are rewarded based on this outcome, and agents are paid by their principals after the outcome is revealed. In each interaction, principals are shown information that allows them to identify the agent's gender. This information is conveyed through agents' (nick-)names, and presented along with other demographic information to minimize demand effects. After each interaction, we elicit agents' and principals' beliefs about the agents' performance.

Our main tests follow from the concept of attribution bias by gender. Following a high outcome, we test whether principals are more likely to attribute it to the agent's ability if male and to luck if female. This would result, in turn, in greater payments being made to males relative to females conditional on a high outcome. Similarly, we test whether principals attribute a low outcome to the agent's luck if male and to the agent's ability if female. Thus, again we test whether female agents receive lower payments as compared to male agents conditional on a low outcome.

Our experimental results do not show evidence of attribution bias by gender. While principals' payments are heavily influenced by the realized outcomes, they do not differ by the gender of the agent. Similarly, principals' beliefs about agents' performance do not differ by gender, although they are heavily influenced by the realized outcomes. Our results, therefore, suggest that gender is not a driving force when principals assess the agents' performance, at least in a laboratory environment. We do, however, find evidence that the interaction between principal and agent gender affects payment decisions. In particular, principals pay higher wages to agents of the opposite gender.

We show that our results are robust to including session fixed effects and round fixed effects, to restricting the sample to the first ten rounds of the experiment, to discarding the first five rounds from the sample, and to alternative definitions of the dependent variable. We provide evidence that principals did treat payments as relevant, that they were aware of the gender of the agent, that principals' prior beliefs did not differ by agent gender, that our results are unlikely to be due to sample selection, and that payments were not driven by agents' ages.

2.1.1 Contribution

How individuals attribute causes of behavior and outcomes to both dispositional and situational factors has received considerable attention in social psychology. In particular, the fundamental attribution error, the tendency of observers to assign too much weight to dispositional factors (e.g., preferences and ability) and too little weight to situational factors (e.g., constraints and luck) when interpreting others' behaviour and performance, has been the focus of several studies (e.g., Jones and Harris (1967), Moore et al. (2010), Ross (1977)). Agents take the perceived role of luck into account when rewarding performance (e.g. Erkal, Gangadharan and Nikiforakis (2011), Rey-Biel, Sheremeta and Uler (2018), Rubin and Sheremeta (2016)). Here, however, we are interested in a specific manifestation of this bias: attribution bias by gender. Evidence in social psychology, for example, shows that observers are more likely to attribute good performance of males to skill and females to luck in certain tasks. Parents and teachers have been shown to suffer from attribution bias too (Deaux and Emswiller, 1974; Dweck et al., 1978; Espinoza, da Luz Fontes and Arms-Chavez, 2014; Fennema et al., 1990; Yee and Eccles, 1988). We contribute to this literature by testing whether attribution bias by gender exists in a general framework. First, our design features variation in tasks so we can study whether gender-biased attributions are task-dependent. Second, our principal-agent setting allows us to mimic a variety of real-world environments such as workplaces and educational institutions. Our experiment explicitly controls for the output-generating process to isolate the dispositional and situational factors affecting the outcome.

Experimental and applied work within economics has emphasized gender differences in preferences (risk and ambiguity, competition, social preferences, negotiation, among others) as possible explanations for differences in economic outcomes such as income, education, and types of occupation (Adams and Funk, 2012; Flory, Leibbrandt and List, 2014; Saccardo, Pietrasz and Gneezy, 2018).[1] However, a vast literature also shows that discrimination contributes to differences in labor market outcomes by gender at several stages, including screening, hiring, and promotion.[2] Taste-based discrimination (Becker, 2010) and statistical discrimination (Phelps, 1972) have been widely studied (Bertrand and Duflo, 2017). Individuals who anticipate discrimination may change their own behavior, intensifying group differences along dimensions such as productivity (Glover, Pallais and Pariente, 2017), self-beliefs (Beyer, 1998; Keller, 2001) and perceived performance (BenYishay et al., 2020; Leibbrandt, Wang and Foo, 2018). Another mechanism that has been put forward to explain differences in economic outcomes is that of attribution bias. Sarsons (2017) shows, for example, that women are given less credit for group work than men. Other examples, cited above, come from the markets for surgeons, executives, and financial advisors (Sarsons, 2019; Selody, 2010; Landsman, 2019; Egan, Matvos and Seru, 2017). Complementary to these studies, we develop an experiment where there is uncertainty about an individual's contribution to output and explicitly model the uncertainty, which is known to principals in our setup. We further elicit principals' beliefs. Thus, we are able to perform a clean test for attribution bias by gender.

A large literature emphasizes the importance of stereotypes and their influence on judgements about performance or ability (Alan, Ertac and Mumcu, 2018; Bohnet, Van Geen and Bazerman, 2016; Bordalo et al., 2019; Carlana, 2019; Coffman, 2014; Coffman, Collis and Kulkarni, 2019; Milkman, Akinola and Chugh, 2013). In particular, this literature finds that stereotypes about tasks lead to biased judgments

[1]See Bertrand (2011) and Croson and Gneezy (2009) for reviews.
[2]See Azmat and Petrongolo (2014) and Bertrand and Duflo (2017) for reviews.

about others' and own ability to perform gender-incongruent tasks. Women performing male-typed tasks and men performing female-typed tasks are expected to perform worse than the opposite gender. To understand whether stereotypes also drive attribution bias by gender, we introduce variation in tasks performed by the agents.

2.2 Experimental Design

An experiment that studies attribution bias by gender requires several ingredients. First, it requires two roles: an agent whose performance is to be evaluated, and a principal who evaluates the agent's performance. Second, the outcome of the agent's performance needs to be a function of both dispositional and situational factors. Third, the principal must be aware of the gender of the agent.

Our experimental sample consists of 84 students from a university in India. First, we asked participants to fill out a demographic questionnaire. Second, we randomized participants into two roles: principals and agents and matched them for each round into pairs using the stranger-matching protocol. That is, at every round principals and agents were randomly rematched. We opted for the stranger-matching protocol since it avoids reputation building and related strategic concerns. It also prevents the experiment from becoming a Bayesian game, in which principals update their priors about agents after repeated interaction. While the experiment was conducted in person, participants were not given the actual identities of their matched partners. The agent performed a task for the principal. The agent's performance influenced the output, but not deterministically. This output determined the principal's earnings in that round. The principal then paid the agent for his performance. In each of the 20 rounds, we elicited agents' and principals' beliefs about the agent's contribution to the realized outcome. Finally, subjects were asked questions about the experimental task.

2.2.1 The Experiment

At the outset of the experiment, participants completed a demographic questionnaire that included information about their gender, field of study, level of study, country and state of origin, age, caste, and religion. Participants were then randomly assigned to be principals or agents. They were told that these roles were fixed for the session, that the experiment consisted of two tasks, and that the tasks would be played one after the other and for ten rounds each.[3] We then explained to our participants the structure of a round. While the general features of each round were read aloud by the experimenter, we asked participants to read the specific details on their computer screens. To make sure participants understood the experiment, we encouraged them to ask questions if anything was unclear and we asked them comprehension questions. Participants could not continue with the experiment until they had answered all questions correctly.

Description of a Round

After being matched, participants were informed that although they could earn money in each of the 20 rounds, at the end of the experiment only one would be randomly selected to count for the final payments.

Agents' performance and output produced At the beginning of each round, the agent performs a task which was to answer a fixed number of questions within 45 seconds. The agent's performance determines the lottery that is assigned to the principal. Each lottery has only two possible outcomes: High and Low output. The agent's performance (i.e., the number of correctly solved questions) affects the lottery assigned to the principal by increasing the probability that the high output is realized. However, even in the case an agent had solved all questions correctly, there is a positive probability that the resulting output is low.

Principals' payments to agents After 45 seconds the principal is shown her payoff for that round (i.e., the output produced by the agent). Importantly, the principal is not informed about the number of questions solved correctly by the agent. However, the principal knows the mapping between the number of correctly answered questions and the probability of high output. The principal then chooses a reward for her agent. In particular, the principal is given access to a pot of ₹350 and she is free to choose how to divide this amount between her agent, a random agent in the session, and the experimenter.[4] This pot is independent of the realized outcome in that round. Importantly, the agent does not see the payment he receives until the end of the session. In this way, his performance is not dependent on the

[3] In two sessions, participants played 9 rounds per task, rather than 10, due to time constraints.

[4] At the time of the experiment, this amount corresponded to £3.92 (exchange rate as of July 2018: £1.00 = ₹89.21). We followed the same considerations as in Gurdal, Miller and Rustichini (2013). Two features are worth noting. First, not allowing the principal to keep any unassigned money for herself shuts down any (financial) incentive for the principal to keep all the money. Second, having the option to also pay a random agent allows to eliminate any efficiency motives (in terms of subjects versus experimenter considerations) that the principal might have.

history of payments he has received, and the principal's payments will not be driven by an underlying motive of incentivizing the agent to perform well.

Principal's beliefs After the payment decision, we elicited the principal's beliefs about the performance of the agent. In particular, we asked the principal how many questions she thought the agent had answered correctly. We incentivized this question by paying ₹50 if the answer was correct. In some sessions we asked this question while the agent was performing the task and so before the outcome of the lottery was known. That is, we asked the principal to indicate her prior belief about the number of questions that the agent would answer correctly. Finally, we also asked the principal two unincentivized questions. We asked the principal to guess how many questions she thought that the agent attempted, and whether she would like to be paired for another round with the same agent.

Agent's beliefs We asked the agent three unincentivized beliefs' questions. First, we asked him to guess the number of questions that he solved correctly. Second, we asked him to guess whether the principal earned the high or low output. Finally, we asked him what percentage of the ₹350 he expected to receive from the principal.

Debriefing

When the tasks were completed, we asked participants to answer two sets of questions. First, we asked participants to guess our research questions. Second, we asked participants questions about the previous tasks. For instance, we asked them which task was more difficult and whether the agents were anxious or stressed while performing the task; and, similarly, whether the principals were anxious or stressed while the agents were performing the task.

2.2.2 Gender Information

In each round, while the agent was performing the task the principal was shown some demographic information about the agent. In particular, the principal was given information about whether the agent was a university student, the agent's age, and the agent's gender. We disclosed gender information of the agent via nicknames. By using software to assign a gender-congruent nickname to the agent. We did the experiment in India and randomly selected from a list of popular Indian names. Since we used only first names, they did not signal caste. We selected the most popular Hindu names. Female names included "Akansha", "Neha" and

"Priya" and male names included "Amit", "Ashish" and "Nitin".[5]

We used nicknames instead of a direct statement of the agent's gender to mask the fact that our research question was gender-related and prevent potential distortions due to demand effects and social desirability concerns. We opted for nicknames rather than real names because we wanted to preserve anonymity and control more carefully for the type of information disclosed via names. For instance, we wanted to make sure that names did reveal the gender of the agent, and that they did not prime religion or caste-related information. Principals were instructed that these were pseudonyms, not real names.

The Tasks and Output

We had four tasks that we varied across sessions: a math task (39% of sessions), a Raven task (31% of sessions), an effort task (19% of sessions), and a memory task (11% of sessions).

The math task We implemented a variation of the Niederle and Vesterlund (2007) math task. In each round of this task, agents were asked to perform 7 additions. Each addition consisted of three two-digit numbers.

The Raven task In each round of this task, agents were asked to solve three Raven Matrices. For our experiment, we used the matrices from the Raven Advanced Progressive Matrices (APM). This test is commonly used to measure fluid intelligence (Carpenter, Just and Shell, 1990).

The effort task We used a variation of the Abeler et al. (2011) effort task. In this task, agents were shown ten 5×5 matrices that were randomly filled with zeros and ones. Agents were asked to solve as many grids as possible by counting the number of ones in each matrix.

The memory task This is a working memory exercise. Agents were shown 16 common English words (e.g., cat, umbrella, house) for 25 seconds. After that, the words disappeared from the screen and they had to write down as many words as they could remember.

[5] Showing other religious groups or full names would have primed religion and/or caste.

Lotteries and Output

Each correct answer increased the probability of the high output being realized. In each task and for each round we had variation in two dimensions: the mapping of correct answers into the probability of the high output (i.e., the set of lotteries) and the level of the high output.[6] These were randomly assigned and orthogonal to each other.

The lotteries Given that the number of questions asked by task differed, the precise mapping of correct questions into the probability of the high output occurring changed by task. However, the overarching feature across tasks was that the probability of the high output was always increasing in the number of correct questions solved by the agent. Moreover, for each task, we had two different mappings: The high and low calibrations. In the former, the probability of the high output started at 50% had the agent solved one question correctly and, as the agent solved more questions correctly, it could exceed 90%, though it could never reach 100%. In the latter, the probability of the high outcome started at 5% and could at most reach 60% had the agent solved all questions correctly. We varied the mapping in order to understand whether this feature affects payments, beliefs, and gender-biased attributions.

Output level The high output could take three different levels: ₹400, ₹550, or ₹700. We varied the level of the high output to see whether principals' payments and beliefs were affected by the potential value of the high output. Importantly, both the agents and the principals had access to this information in each round. Agents were shown the mapping and the output level before they performed the task, while principals were shown this information at the time the agents were performing the task. Both agents and principals were given unlimited time to read and process the information, which was provided in table form for intuitive exposition and ease of understanding.

2.2.3 Attribution Bias by Gender

Our experiment is designed to determine whether principals make biased attributions regarding the performance of the agents. To capture attribution bias, we designed an environment in which outputs represent noisy signals of the agent's performance. The output in each round is a function of the number of questions

[6] The low output was always set equal to ₹0.

answered correctly by the agent and luck. The principal therefore has to base her payment on the basis of outcome of the lottery.

In this environment, we test whether the gender of the agent plays a crucial role in the principal's payment and in shaping her beliefs about how much the agent's competence contributed to the output. In particular, our empirical tests follow directly from the concept of attribution bias by gender. A principal exhibiting attribution bias by gender will attribute a high output to the agent's performance if the agent is male and to luck if the agent is female. Similarly, following a low output, the principal will attribute it to misfortune if the agent is male and to performance if the agent is a female. This difference in the principal's beliefs by gender would affect payments principals make to agents. We implement the following tests:

1. We test whether the principal's belief about the number of correctly solved questions is higher for male than for female agents, following both high and low outputs.

2. We test whether the principal's payments are higher for male agents than for female agents, following both high and low outputs.

3. We test whether the sensitivity of the principal's beliefs about the number of correctly solved questions and the sensitivity of the principal's payments to the realization of output differ for male agents and female agents.

2.3 Experimental Results

2.3.1 Implementation

The experiment was conducted in July 2018 in the computer lab at the Delhi School of Economics. Invited participants belonged to the departments of Commerce, Economics, Geography, and Sociology. We recruited 84 subjects and conducted 5 sessions of about 75 minutes each. The participants earned on average ₹510, which includes the show-up fee of ₹250. We programmed the experiment with oTree (Chen, Schonger and Wickens, 2016).

We begin by examining agents' performance and their beliefs in Section 2.3.2. In particular, we look at agents' performance across tasks and by gender. In this section, we also analyze agents' beliefs about their own performance and their beliefs about their principals' payment decisions. We then investigate principals' payment decisions and their beliefs about their agents' performances in Section 2.3.3. This tells us whether principals make biased attributions and payments depending on the gender of their matched agents. We present our econometric specifications in Section 2.3.4, and the results of these estimations in Section 2.3.5. In Section 2.3.6 we discuss alternative factors that might be driving our results: the salience of gender information, principals' prior beliefs about the agents' performances, selection of our sample, and whether principals' payments are driven by other demographic information about the agents such as age. Summary statistics are reported in Tables B.1, B.2, and B.3 in Appendix B.1.

2.3.2 Agents

Agents' Performances and their Beliefs

Performance The mean proportion of correctly answered questions across tasks was 39% (s.d. 0.23). Performance, defined as the proportion of questions solved correctly, varies by task: it is highest in the math and effort tasks with over 50% of questions solved correctly, while it is lowest for the Raven task with 22% of correct answers. Looking at performance broken down by gender in Table 2.1, we find no difference in performance by gender across tasks.[7]

Table 2.1: Mean performance by task and gender

	All Agents	Female Agents	Male Agents	P-value Difference
All Tasks	0.39	0.39	0.40	0.85
Math Task	0.52	0.51	0.53	0.73
Raven Task	0.22	0.21	0.23	0.62
Memory Task	0.33	0.33	0.33	0.85
Effort Task	0.53	0.54	0.52	0.81

Statistical significance is assessed by running regressions of performance (proportion of questions solved correctly) on the gender of the agent. Standard errors are clustered at the individual level.

Beliefs about performance As can be seen in Table 2.2, agents' beliefs about their own performances differ by task, but there are no differences in beliefs by gender. If we compare performance and beliefs, we can see that agents are overconfident in the math and the Raven tasks. Indeed, they overestimated their performance by about 15%.

[7]There are also no significant differences in performances' distributions nor in the variance of the number of correct questions by gender and across tasks.

Table 2.2: Mean beliefs about performance by task and gender

	All Agents	Female Agents	Male Agents	P-value Difference
All Tasks	0.51	0.49	0.52	0.41
Math Task	0.60	0.58	0.62	0.34
Raven Task	0.47	0.46	0.48	0.78
Memory Task	0.36	0.35	0.36	0.71
Effort Task	0.55	0.54	0.57	0.45

Statistical significance is assessed by running regressions of performance beliefs on the gender of the agent. Standard errors are clustered at the individual level.

Agents' Beliefs about Outcomes and Expected Payments

Beliefs about realized outcomes and expected principals' payments At the top of Table 2.3, we look at agents' beliefs about realized outcomes (i.e. the perceived probability that the high outcome was realized) by gender and task. We find essentially the same patterns as with beliefs about performance: agents are overconfident, but their beliefs do not differ by gender. On the other hand, at the bottom of Table 2.3, we can see that female agents believe that, on average, principals will allocate 59% of the ₹350 to them, while male agents believe they will receive roughly 51%. However, the difference is not statistically significant.

Table 2.3: Mean beliefs about realized outcomes and expected principals' payments

	All Agents	Female Agents	Male Agents	P-value
	Beliefs about Realized Outcomes			
All Tasks	0.74	0.73	0.75	0.78
Math Task	0.77	0.76	0.77	0.87
Raven Task	0.65	0.62	0.68	0.47
Memory Task	0.75	0.78	0.72	0.71
Effort Task	0.87	0.85	0.88	0.72
	Beliefs about Principals' Payments			
All Tasks	0.55	0.59	0.51	0.15
Math Task	0.59	0.60	0.58	0.74
Raven Task	0.46	0.48	0.45	0.66
Memory Task	0.61	0.68	0.49	0.03
Effort Task	0.55	0.65	0.47	0.23

Statistical significance is assessed by running regressions of either beliefs about realized outcomes or beliefs about principals' payments on the gender of the agent. Standard errors are clustered at the individual level.

In sum, we find that, while performance differs by task, it does not significantly differ by gender. Similarly, beliefs about own performance and principals' actions do not differ by the gender of the agent. Male and female agents' beliefs about principals' actions differ but the difference is not statistically significant.

2.3.3 Principals

We now turn to our main outcome variables: the principals' payment decisions and beliefs. We start by considering principals' beliefs and choices depending on the outcome produced by their agents. We then analyze these variables depending on the agent's gender. Importantly, in the following analysis we looked at pooled results that take into account all tasks and calibrations. The outcome produced in each round is coded as 0 when the realized outcome of the lottery was low (₹0) and 1 if it was high (₹400, ₹550, or ₹700).

Wages and Principals' Beliefs

Wages In Figure 2.1, we show the distribution of principals' payments depending on the realized outcome. From the figure it is clear that payments depended heavily on the realized outcome: higher payments were made following a high outcome and lower payments were made following low outcomes. A Mann-Whitney test confirms that the distribution of principals' payments differs significantly by the realized outcome (p-value$<$ 0.00).

Figure 2.1: Principals' payment decisions by realized outcome

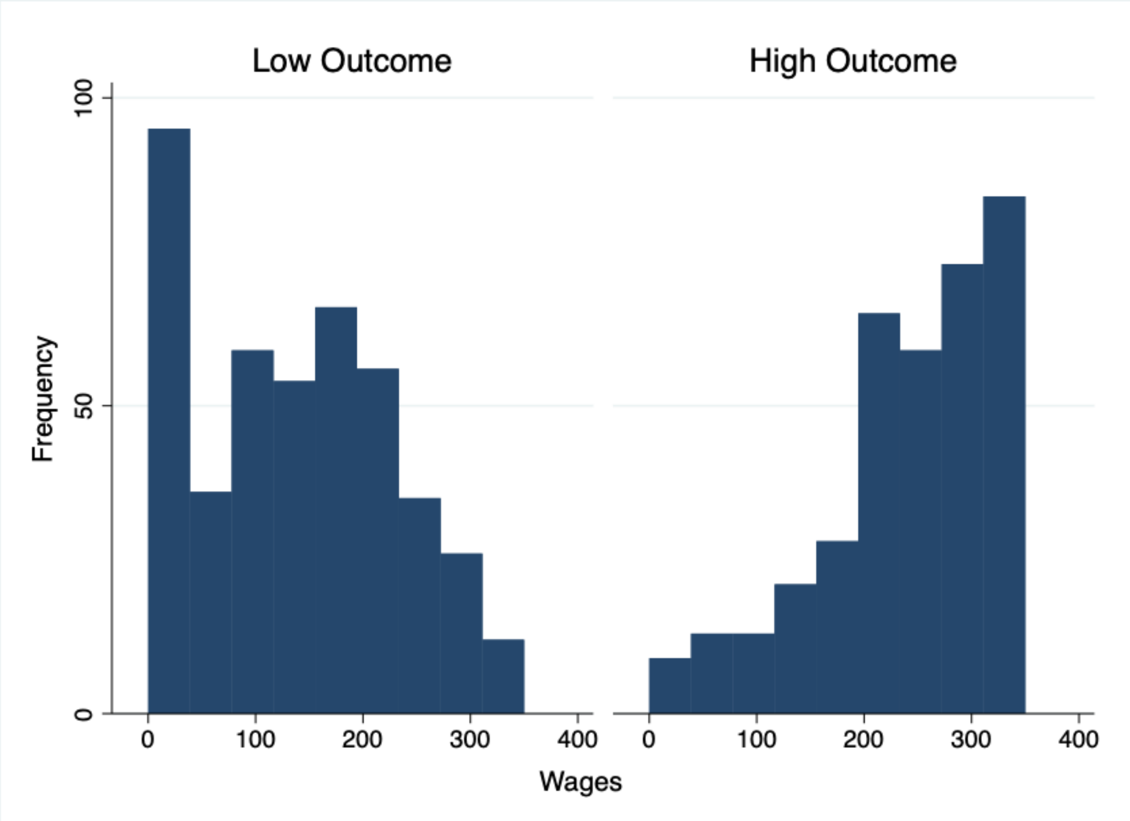

Notes: the histograms show the distribution of principals' payment decisions by realized outcome.

Beliefs Figure 2.2 shows that principals' beliefs about their agents' performances follow a similar pattern. The principals' beliefs here correspond to the proportion of questions that they think the agents have solved correctly in a given round. Principals' beliefs are higher when the output is high compared to when it is low. A Mann-Whitney test shows the the distribution of principals' beliefs differs significantly according to the realized outcome (p-value$<$ 0.00).

Figure 2.2: Principals' beliefs by realized outcome

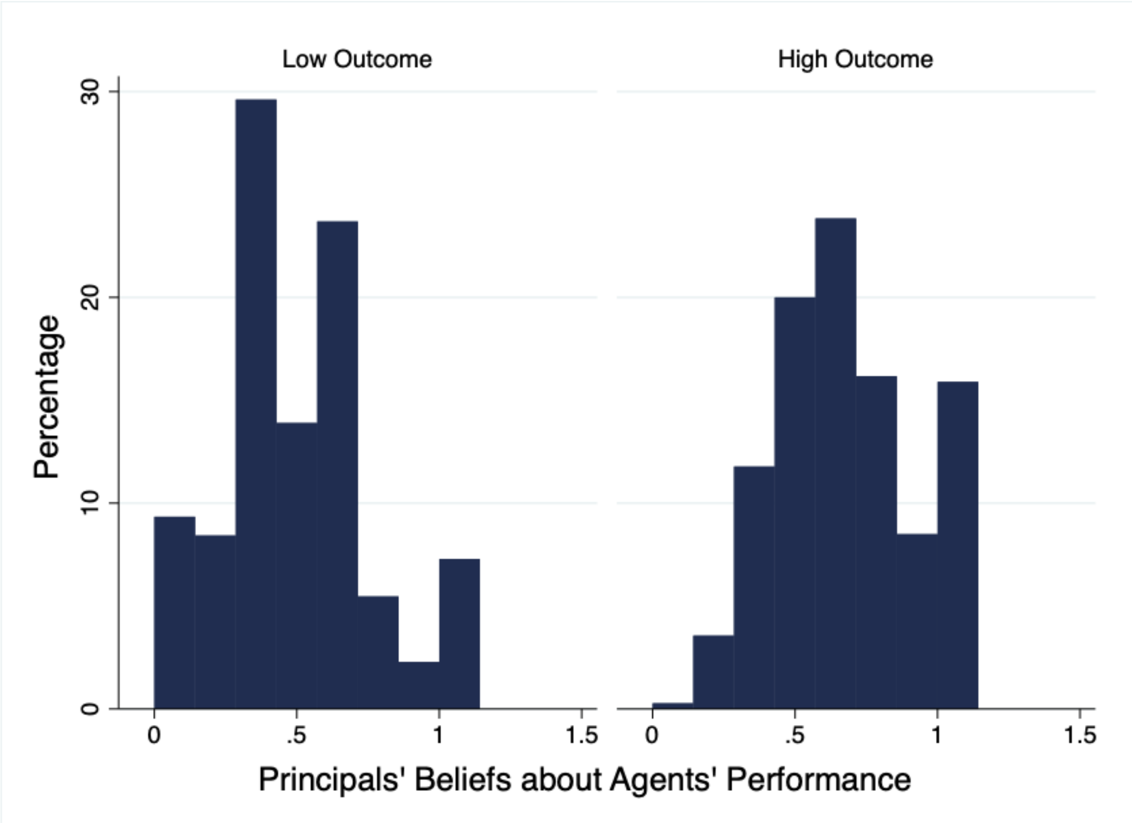

Notes: the histograms show the distributions of principals' beliefs about the matched agents' performance by realized outcome.

Wages and Principals' Beliefs by the Gender of their Matched Agents

Wages by gender of the agent We now analyze whether there are differences in principals' wages depending on the gender of their matched agents. Figure 2.3 shows that, while payments respond to the outcome of the lottery, they do not differentially respond by the agents' gender. Indeed, a Mann-Whitney test fails to reject the null hypotheses of equality in distributions following either a low (p-value= 0.611) or a high outcome (p-value= 0.883).[8]

[8] In appendix B.2, we also show mean payments (from the separate pot) made to the other randomly matched agent and to the experimenter by realized outcome and the gender of the agent.

Figure 2.3: Principals' wages by realized outcome and agents' gender

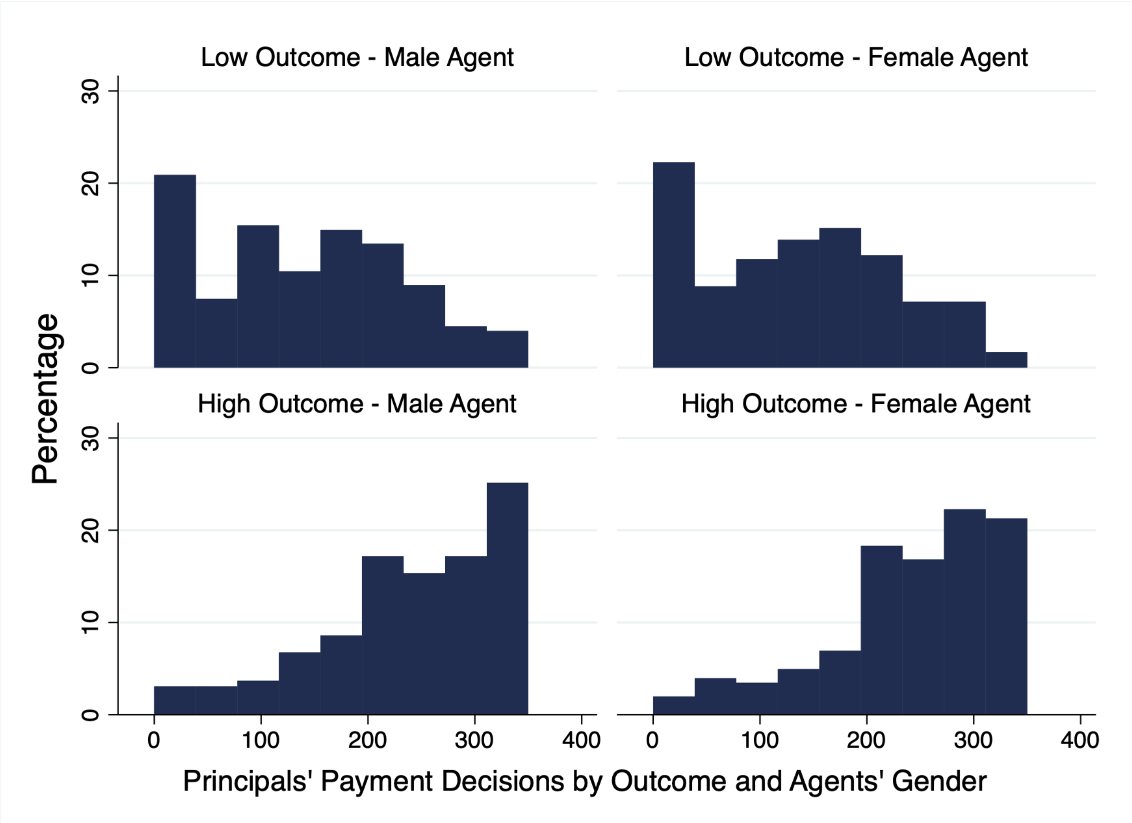

Notes: the histograms show the distribution of principals' payment decisions by realized outcome and the gender of the matched agents.

Beliefs by gender of the agent The results for beliefs match those for wages. Figure 2.4 shows that, while beliefs about the number of questions solved correctly are heavily influenced by the realized outcome, they do not shift according to the agent's gender. Results of a Mann-Whitney test show no significant difference in distributions irrespective of whether the outcome is low or high (p-value=0.514 and p-value=0.884).

Figure 2.4: Principals' beliefs by realized outcome and agents' gender

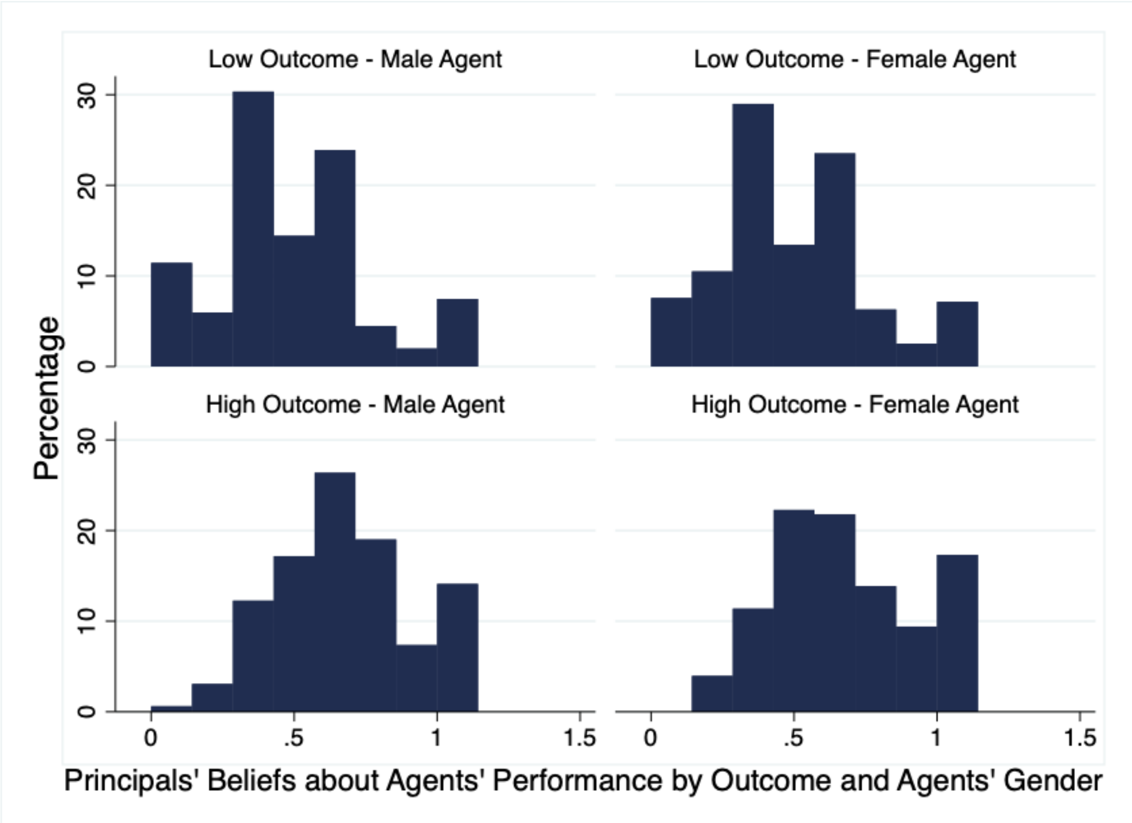

Notes: the histograms show the distributions of principals' beliefs about the matched agents' performance by realized outcome and gender of the matched agents.

In sum, our experimental results do not provide evidence that principals' payment decisions and beliefs are influenced by their agent's gender.

2.3.4 Econometric Specifications

We next conduct parametric analyses to further analyze the variables affecting the principals' payment decisions. In particular, we estimate the following regression:

$$Y_{ij} = \beta_0 + \beta_1 Z_{ijr} + \beta_2 Female_i + \beta_3 Female_i \times Z_{ijr} + x'_{ijr}\beta_4 + \epsilon_{ij} \tag{2.1}$$

Here, i is the agent, j is the principal and r is the round, Y_{ij} is the dependent variable. This is either the principal's payment to the agent or her belief about the agent's performance. Z_{ijr} is a dummy for a high outcome in the lottery produced by agent i matched with principal j in round r, $Female_i$ is a dummy equal to 1 if the agent is a female. x_{ijr} is a vector of controls that includes principals' demographic variables (age, caste, religion, field of study, education level, and state of birth) and task characteristics (task, calibration of the lottery, and level of the high outcome). We report standard errors clustered at the principal level in all specifications. β_2 captures whether there are any average differences in payments made to female versus male agents when the outcome produced is low, while β_3 captures if there is any difference in the increase in the payment made to female agents in response to a high outcome, compared to the increase for male agents.

We then also control for the principal's gender to check whether this variable affects the payments made to the agent. We use the following econometric specification:

$$\begin{aligned} Y_{ij} = \beta_0 + \beta_1 Z_{ijr} + \beta_2 Female_i + \beta_3 Female_i \times Z_{ijr} + \\ \beta_4 Female_j + \beta_5 Female_j \times Z_{ijr} + x'_{ijr}\beta_6 + \epsilon_{ij} \end{aligned} \tag{2.2}$$

$Female_j$ is a dummy equal to 1 if principal j was a female. β_4 captures if there are any average differences in payments made by male versus female principals in the case of a low outcome (holding everything else constant), while β_5 captures whether there is any difference in the increase in payments made by female principals in response to the high outcome, relative to the increase made by male principals. Our random matching design also allows us to test for an interaction between agent's gender and the principal's gender. Hence we also report estimates from the following specification:

$$\begin{aligned} Y_{ij} = \beta_0 + \beta_1 Z_{ijr} + \beta_2 Female_i + \beta_3 Female_i \times Z_{ijr} + \beta_4 Female_j \\ + \beta_5 Female_j \times Z_{ijr} + \beta_6 SameGender_{ij} + x'_{ijr}\beta_7 + \epsilon_{ij} \end{aligned} \tag{2.3}$$

β_6 here captures whether being matched to an agent of the same gender leads to any differential effect on payments made by principals. In Section 2.3.6, we will show the robustness of our results to alternative specifications, including fixed effects for principals, sessions, and rounds.

2.3.5 Econometric Results

The results for principals' payment decisions (mean ₹184 and s.d. 107.15) are shown in Table 2.4. As is apparent from columns 1 to 4, the outcome of the lottery is

Table 2.4: Regression results for principal's payments

	(1)	(2)	(3)	(4)
High Outcome	101.57***	101.59***	112.79***	110.01***
	(14.38)	(17.59)	(33.53)	(33.47)
Female Agent		1.92	1.31	10.01
		(9.25)	(9.06)	(9.54)
Female Agent × High Outcome		-0.04	0.96	-0.27
		(13.20)	(12.41)	(12.54)
Female Principal			-34.15	-33.59
			(40.89)	(39.92)
Female Principal × High Outcome			-15.49	-12.46
			(35.42)	(35.31)
Same Gender				-17.22**
				(6.86)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.33	0.33	0.34	0.35
N	804	804	804	804

Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Task controls are dummy variables for each task. Standard errors are clustered at the principal level.

important in determining the payment made to the agent.[9] Going from a low to a high outcome increases the principal's payment to the agent by around ₹100 in all specifications. On the other hand, the agent's gender does not play a role. The coefficient on the female dummy and the interaction with the outcome variable are small and not significant. The coefficient on the female agent dummy varies between ₹1 and ₹10 across specifications, compared to the average of the dependent variable, which is ₹184. Thus, principals did not make payments differently to women as compared to men conditional on the same outcome. We do not, then, find evidence of attribution bias by gender. An F-test for whether the total effect of a high outcome on payments made to female agents is similar to that made to men cannot be rejected (p value=0.736). At a significance level of 5% and with 80% power, our minimum detectable effect size is $(t_{1-\kappa} + t_\alpha) \times SE_{\hat{\beta}}$, or approximately $(0.84 + 1.96) \times 9.25 \approx 25.9$, which is less than one quarter the estimated coefficient on a high outcome and less than 15% of the mean of the outcome variable.

If we perform the same regressions for beliefs (mean 0.54 and s.d. 0.26), we find the same patterns. Table 2.5 shows that principals' beliefs are significantly shaped

[9] The number of observations is 804 because in two sessions we had 9 rounds per task and hence 18 rounds instead of 20. This gives us 180, 160, 140, 162 and 162 observations for each session.

by outcomes while they are not affected by agent gender. Female principals are more likely to believe that agents solved a smaller proportion of questions correctly in case of a low outcome while they inflate their beliefs significantly more than the male principals in response to a high outcome. Their payments react less than those of male principals, as seen in Table 2.4, though this difference is not significant. Here, our minimum detectable effect size for the coefficient on female is $(0.84 + 1.96) \times 0.02 \approx 0.06$, which is slightly larger than one quarter of the estimated coefficient on a high outcome and less than 12% of the mean of the outcome variable.

Table 2.5: Regression results for principal's beliefs

	(1)	(2)	(3)	(4)
High Outcome	0.21***	0.21***	0.16***	0.16***
	(0.03)	(0.03)	(0.03)	(0.03)
Female Agent		0.02	0.03	0.04*
		(0.02)	(0.02)	(0.02)
Female Agent × Outcome		-0.02	-0.02	-0.02
		(0.03)	(0.03)	(0.03)
Female Principal			-0.13**	-0.13**
			(0.06)	(0.06)
Female Principal × Outcome			0.09*	0.09**
			(0.05)	(0.05)
Same Gender				-0.02
				(0.02)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.30	0.30	0.31	0.31
N	804	804	804	804

Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Standard errors are clustered at the principal level.

Thus, we find that a high outcome leads to a smaller increase increase in payments for female agents, a difference equivalent to roughly 2% of the outcome mean. In comparison to our results, Sarsons (2019) finds that the increase in referrals was 17% less for women after a positive outcome, while Selody (2010) shows that the growth rate in female executive pay is 25% lower than that of male executives after their company faces an unexpected good outcome. Similarly following a low outcome, we find that there is a larger payment made to female agents equivalent to 1.3% of the mean, but that is not significant. Egan, Matvos and Seru (2017), by contrast, find a 20% higher likelihood of punishment for women following a misconduct incident, while Selody (2010) estimates a decrease in top executive pay of about 68% for women relative to men after a negative change in the firm's market value. These differences in results could be due to differences in outcome variables. However we also measure beliefs and we find that, while there is a smaller increase in beliefs about ability for women than men after a high outcome, the difference is not significant and is roughly 9% of the belief mean. Thus, we find no evidence of biased beliefs by gender.

While the gender of the agent alone does not influence payments, its *interaction* with the principal's gender does and significantly so as shown in column 4 of Table 2.4. In particular, payments are significantly higher to agents of the opposite gender, irrespective of the realized outcome. In other words, principals pay around ₹17 less to their matched agents if they belong to the same gender. In Table 2.5, with beliefs as the outcome variable, we find that the coefficient on same gender is negative, in line with the evidence for payments, however, this is not significant.

Taken together, our results show that principals' beliefs and payment decisions are heavily influenced by realized outcomes. We do not, however, find evidence of gender-biased attributions.

2.3.6 Robustness Checks

Principal, Session, and Round Fixed Effects

In the appendix, we show that our results are robust to the inclusion of several sets of fixed effects. In Table B.4, we show that results for beliefs and payments are similar when fixed effects are included for principals. Results for principals' payments are robust to including session fixed effects (see Table B.5) and round fixed effects (see Table B.6). Running regressions without controls shows that the main result for principals' payments is robust to not including controls (see Table B.7). The results regarding the same gender of the principal and agent also continue to hold. Similarly, the main results for beliefs are also robust to including session fixed effects (see Table B.12) and round fixed effects (see Table B.13). Results without controls are similar in magnitude and significance, as shown in Table B.14.

Restricting the Sample to the First Ten Rounds

If principals' beliefs in early rounds are more biased than in later rounds, for example if they have not yet been influenced by observing the realized outcomes, then it is possible that attribution bias by gender was only present in the initial rounds of the experiment. To test for this, in Tables B.8 and B.15 we report similar regressions restricting the sample to the first ten rounds. The results for attribution bias by gender hold. The coefficient on the female agent dummy stays small relative to the mean payment of ₹184. The coefficient on the interaction of the female dummy with the high outcome dummy is negative but still small and not statistically significant. For the case of payments, the coefficient on same gender becomes larger by ₹7 than the estimate obtained using the whole sample. In Table B.15, high outcomes no longer make female principals update their beliefs significantly more than males. The results for attribution bias by gender, however, are the same and there is no significant effect of being matched with a female agent on the beliefs of the principal after a high or a low outcome. In summary, we do not find evidence that attribution bias by gender arises even in the initial rounds of the experiment.

Removing the First Five Rounds from the Sample

We further look at the results after removing the first five rounds of each session to account for the possibility that participants may not have fully understood the experiment in the first few rounds. In Tables B.9 and Table B.16, we can see that there is no evidence of attribution bias by gender. The coefficients on the female agent dummy become larger than in the main results, but are still not statistically significant and are small in comparison to the mean of the outcome variable. The coefficient for same gender remains large and significant. Thus our results do not appear to be driven by principals not understanding the experiment in the initial five rounds.

Alternative Dependent Variables

To evaluate whether attribution bias manifests in changes in the shape of the payment distribution, we have estimated Equation (2.1) using alternative definitions of the dependent variable. We thus estimate the following econometric specification:

$$Y_{ijx} = \beta_0 + \beta_1 Z_{ijr} + \beta_2 Female_i + \beta_3 Female_j + x'_{ijr}\beta_3 + \epsilon_{ij} \tag{2.4}$$

This equation is the same as the one in Equation (2.1) except that the dependent variable Y_{ijx} is a dummy variable which is equal to one if the payment to the agent takes a value greater than or equal to x. We vary x from ₹50 to ₹300 in increments of ₹50. The estimates are shown in Table B.10. For all these different cutoffs, we find similar results in that the gender of the agent does not play any role in driving payments.

Results for Different Tasks

We show in Table B.11 that our results are not driven by any one task. In particular, if tasks that are known to be male stereotypical could lead to attribution bias, we would find attribution bias for these tasks. Although our tasks are not as gender stereotypical as those used in papers focused directly on stereotypes (e.g. Coffman (2014); Coffman, Collis and Kulkarni (2019)), it is possible that our participants perceived some tasks as gender stereotypical. However, the coefficients on female agent and its interaction with high outcome remain small in comparison to the coefficient on high outcome across all tasks. The coefficient on female agent interacted with high outcome becomes significant at the 10% level for the memory task, though the coefficient is now positive. Similarly, in Table B.17 we find that there is no significant effect of the female dummy on the beliefs of principals and that coefficients remain small relative to effect of a high outcome on beliefs across tasks.

2.4 Possible Threats

2.4.1 Irrelevance of Payments

Given that the principals' payments were made from a separate pot of money, and thus they were payoff-irrelevant for the principals, one possible explanation for the results could be that these payments did not vary or were chosen randomly. However, the results of section 2.3.3 show that this was not the case: principals understood that their choices had economic implications for their matched agents and were responsive to the realized value of output.

2.4.2 Gender Information

Our lack of experimental evidence for attribution bias by gender could be explained by the way in which we disclosed gender information about the agents. A failure to find attribution bias by gender could be driven by the possibility that principals understood that our research question was about gender discrimination and, therefore, they were particularly cautious in preventing such bias from arising during the experiment (e.g., due to a social desirability bias). However, when subjects in the role of principals were asked to guess our research questions at the end of the experiment,[10] none guessed it was about gender. The most common guesses included

[10] In particular, we asked the following open-text question: "Please, guess what our research questions are."

answers such as: "the sharing tendency of people", "a study on how individual decision making is affected when their possible returns are contingent on the actions of another person", and "assessing contracts".[11]

Alternatively, one might worry that displaying gender information using nicknames is not salient enough to induce gender discrimination. While this is a possible interpretation of our results, if it were indeed the case that principals did not pay attention to the gender of the agent, that would itself be a finding. This would imply that principals did not judge that this piece of information was important in making their payments and attributions. Further, the results regarding the principals' gender in *interaction* with the gender of the agents shows that principals *did* pay attention to the agent's gender. In other words, this result provides evidence that information about the agent's gender was salient and principals did take it into account, although not in a manner consistent with attribution bias by gender.

2.4.3 Prior Beliefs by Gender of the Agent

In two sessions, we also elicited principals' beliefs about the agents' performance prior to any knowledge regarding the realized outcome. While prior beliefs that principals have are slightly higher for male agents than for female agents (70% of questions solved correctly vs. 67%), the difference is not statistically significant (p-value$= 0.29$). We therefore do not believe that our results are driven by differences in prior beliefs.

2.4.4 Selection of our Sample

Since we conducted experimental sessions at the Delhi School of Economics, one may wonder whether our "null" results might be driven by sample selection: women at this university may be positively selected relative to the population, which would affect how principals behave. Two considerations are worth emphasizing. First, it is not the case that they did better on the tasks than male participants. Second, our sample of positively selected females resembles the same samples (e.g. highly educated female physicians, CEOs) in which observational studies have found patterns consistent with attribution bias by gender.

[11] Importantly, our subject pool was new to experiments. Therefore, they were not aware that in standard experiments subjects' personal characteristics (such as nicknames and age) are not usually disclosed.

2.4.5 Agents' Ages

When we showed gender information about the agent, we also showed the principal the age of the agent as a way to mask our research question. We chose age in particular given the relatively small variation in age among university students. We can check therefore whether payments and beliefs are affected by this piece of information. That is, we can test whether the principal's payments and beliefs are driven by the agent's age. When we run the same regressions as before, replacing gender with age in the set of control variables, we find that neither the agent's age nor that of the principal affect payments or beliefs (see Tables B.18 and B.19).

2.5 Conclusion

Recent literature has suggested that a particular form of discrimination – attribution bias by gender – might affect assessments of actors' outcomes in economic environments. We conduct a laboratory experiment to test for this effect. Our results do not show evidence consistent with attribution bias by gender. While in our experiment principals' beliefs and payments are influenced by realized outcomes, we find no evidence that they differ by the agent's gender. With the caveat that we have a relatively small sample size, our findings suggest that attribution bias by gender does not arise in a controlled environment. However our findings need not imply that attribution bias by gender does not play a role in real-world settings. In other environments, where gender may be more salient, this bias may emerge naturally.

Chapter 3

Anger Impairs Strategic Behavior: A Beauty-Contest Based Analysis

The frustration-aggression hypothesis posits that anger affects economic behaviour essentially by changing temporally the individual preferences. Here, we test a different channel in an experiment where we externally induce anger to a subgroup of subjects (following a standard procedure that we verify by using a novel method of textual analysis). We show that anger can impair the capacity to think strategically in a beauty contest game, in a pre-registered experiment. Angry subjects choose numbers further away from the best response level and earn significantly lower profits. Using a finite mixture model, we show that anger increases more than 30% the number of level-zero players. Furthermore, with a second pre-registered experiment, we show that this effect is not common to all negative emotions. Sad subjects do not play significantly further away from the best response level than the control group and sadness does not lead to more level-zero play. We analyze the implications of this finding on a bargaining game.

3.1 Introduction

Anger is an important emotion that pervasively affects many basic interactions among people in daily life. A growing economic literature has examined the effects of anger on economic decisions. The repercussions of anger on economic behavior that have emerged in the literature are generally based on the frustration-aggression hypothesis (Dollard et al., 1939; Selten, 1978), which is based on the idea that anger can lead a person to behave in hostile ways to someone else – regardless of whether the person targeted is the source of the anger.[1] In general, this literature hinges on the hypothesis that anger and frustration generate preferences for punishment.

This paper examines the effects of anger from another angle: by looking at how such emotion affects the strategic ability. We see our approach as complementary rather than a substitute of the former.

Our results show that anger negatively affects the capacity of thinking strategically. This finding is puzzling because anger is pervasive in human behavior. As the literature in psychology argues, a plausible explanation relies on the notion that anger, like other emotions, can serve as a credible commitment device in situations of conflict (e.g. Elster, 1998; Frank, 1987, 1988; Hirshleifer, 1987), and thus can lead to greater evolutionary success in strategic interactions. If angry individuals do not think carefully about the consequences of their actions – as our results suggest – then others have reason to be wary of the angry person, to avoid conflict. We will use the standard ultimatum game to illustrate this argument (see Section 3.5).

Our work is in line with the literature on cognitive psychology positing that anger may be linked to the impairment of cognitive processes. For example, it has been experimentally shown that anger promotes heuristic processing of information at the expense of more systematic processing.[2] For instance, Tiedens and Linton (2001) find that being angry leads to lower information processing by making individuals rely more on persuasive messages and stereotypes, rather than on the strength of the arguments.[3] Moreover, Gneezy and Imas (2014) show how individuals may sometimes strategically exploit this cognitive impairment of the opponent.

To test our hypothesis that anger decreases the capacity to reason strategically, we conduct two experiments involving a beauty contest game. Prior to the start of the game, we use written exercises to exogenously induce anger in one treatment group, and we compare the play of these subjects with those of a control group, among whom a placebo exercise is conducted (Experiment 1). To evaluate whether any effects on strategic reasoning stem specifically from anger, rather than from negative emotions, we conduct a second, identical experiment in which we instead

[1] For research examining another player in a game as a source of the anger, see, e.g., Xiao and Houser (2005); Rotemberg (2005); Anderson and Simester (2010); Carpenter and Matthews (2012); Winter (2014); Winter, Méndez-Naya and García-Jurado (2016); Akerlof (2016); Van Leeuwen et al. (2017); Passarelli and Tabellini (2017). For research examining the situation in which the other player is not the source of anger, see, e.g., Card and Dahl (2011); Munyo and Rossi (2013). Gurdal, Miller and Rustichini (2013) conduct research examining a situation in which the other player is probably the source of anger. Battigalli, Dufwenberg and Smith (2019) develop a general framework to analyse the frustration-aggression hypothesis in the above-mentioned situations.

[2] For a review, see Litvak et al. (2010).

[3] More generally, research in cognitive and affective sciences has emphasized strong interactions between emotions and cognitive processes (see Engelmann et al., 2018, for a review).

induce sadness among participants in the treated group (Experiment 2).[4]

Our emotion-induction procedures involve asking participants to recall and write about previous experiences that led them to feel angry or sad. These procedures rely on methods and techniques commonly used in and previously validated by the social psychology literature. Furthermore, to analyse the emotional content of subjects' answers to the induction questions, we use a method of textual analysis (to the best of our knowledge we are the first to do that). We underscore that we induce incidental anger rather than provoking a conflict between players; this is designed to achieve our aim of distinguishing our mechanism of interest rather than anger that hinges on social preferences for punishment.[5]

We choose the beauty contest game to test the ability in strategic reasoning because social preferences are unlikely to affect behavior in this game (Eyster, 2019). In fact, as Carpenter, Graham and Wolf (2013) and Gill and Prowse (2016) show, the capacity to play in this game depends on cognitive skills. The beauty contest allows to obtain a rather precise characterization of the level a player's strategic ability, which can be assumed to depend on the ability to form higher-order beliefs (the so-called level-k thinking) (see Nagel, 1995; Duffy and Nagel, 1997; Stahl, 1996) and on another cognitive capacity that characterizes strategic ability – the *theory of mind*. This can be defined as the ability to think about others' thoughts and mental states to predict their intentions and actions (Coricelli and Nagel, 2009).

Our experimental findings from Experiment 1 show a strong negative effect of anger on the strategic ability in the beauty contest game. Subjects who participate in the anger-inducing exercises play the game further away from the best response, given other players' choices – and significantly so compared to those participants among whom no emotion is induced. Furthermore, our structural analysis estimates that there is an increase in level-zero players in the anger treatment group relative to the share of these players in the control group.

Results from Experiment 2 show that sadness has no significant impact on guesses in the game. In the structural analysis we find that, if anything, sadness decreases the number of level-zero players. We also observe that sadness has only a weakly significant negative effect on profits. We thus conclude that anger – rather than negative emotions in general – lowers subjects' capacity of strategic thinking.

Establishing a clear link between anger and capacity of strategic reasoning is im-

[4] According to a common characterization of emotions (Ekman, 1999), the basic negative emotions besides anger are: disgust, fear, and sadness. Among these, we chose sadness because of its closeness with anger; this allows us to maintain the same induction procedure in both experiments.

[5] Incidental emotions are externally generated and unrelated to the process under consideration. See Loewenstein and Lerner (2003) for a discussion of the distinction between incidental and anticipatory emotions.

portant at least for three reasons. First, from a theoretical perspective, it provides some insights on how to include anger in economic models of behavior. Second, as we will see in the conclusion, it has implications for behavioral policies at the individual level. Third, because the effects of incidental emotions are pervasive in many economically relevant decisions (Lerner, Small and Loewenstein, 2004; Tice, Bratslavsky and Baumeister, 2001), and because such emotions may have an enduring and unconscious impact (Vohs, Baumeister and Loewenstein, 2007; Andrade and Ariely, 2009), anger may represent a relevant negative externality for social interactions. This, in turn, represents a potentially important negative externality for a poor economy, who may be more likely to experience negative shocks (see e.g. Koren and Tenreyro, 2007), or among poorest socioeconomic classes, as they likely have little capacity to insure themselves from negative shocks that can generate potentially vicious, widespread cycles of anger and frustration.

The remainder of the paper is organized as follows. Section 3.2 describes our experimental design. Section 3.3 presents our experimental results. Section 3.4 structurally estimates the proportion of level-k thinking by condition and experiment. Section 3.5 discusses the implications of our results on bargaining power in an ultimatum game. Section 3.6 discusses the possible underlying mechanisms and potential confounding factors that may explain the results. Section 3.7 offers final remarks and conclusions. The appendix provides supplementary analysis and tables.

3.2 Experimental Design

To study whether and how anger affects strategic reasoning, and to disentangle the effect of anger from the potential effects of negative emotions more in general, we devise an experiment with the following three features: the ability to exogenously manipulate emotions of participants; to assess the effect of the emotional induction on strategic thinking; and, to compare the effects of the inducement of two different, negative emotions, or the lack of emotional inducement.

Our experimental design is as follows: We set up two experiments with treatment and control groups. Those in the treatment group participated in emotional induction tasks (detailed in Section 3.2.1), which consisted of writing exercises designed to elicit anger (Experiment 1) or sadness (Experiment 2). Those in the control group were asked to complete similar exercises that did not elicit any emotion. Participants representing both treatment and control groups were then matched in groups of three players. They played together the p-beauty contest game for ten rounds. This allowed us to cleanly assess the effects of anger (Experiment 1) and sadness

(Experiment 2) on strategic reasoning. In this way, we tested whether any effect of anger on cognitive reasoning is unique to the emotion of anger itself, or whether it is a more general effect of negative emotions.

Both experiments have been pre-registered. Registration numbers in the AEA Registry: AEARCTR-000426 for the first (anger) experiment and AEARCTR-0004729 for second (sadness) experiment. The University of Warwick Economics Department IRB approval was obtained on March 2019.

3.2.1 The Experiment

At the outset of the experiment, participants were randomly allocated to a computer terminal. The participants, university students, were asked to complete a demographic questionnaire that included information about their age, country of birth, department, year of study, gender, and high school marks. We then asked them to complete the Positive and Negative Affect Schedule (PANAS) questionnaire, (Schamborg, Tully and Browne, 2016; Watson, Clark and Tellegen, 1988). This questionnaire consists of two ten-item scales that measure positive and negative affect. Each item is rated on a five-point Likert scale. We added three further questions from the PANAS–X questionnaire (Watson and Clark, 1999). Two of these additional questions assess the emotions of interest: anger and sadness.[6] This questionnaire provides us with a measure of their overall emotional state before the induction took place.[7]

Emotional Induction

Participants were randomly assigned to be part of the treatment or control group in both experiments. Those in the treatment group then participated in an induction task designed to elicit the emotion of interest: anger (Experiment 1) or sadness (Experiment 2). Participants in the treatment group were asked to answer two questions about past life experiences. They had 10 minutes to answer the questions, and they could not proceed to the next part of the experiment until this time expired. Before participants read the emotional induction questions, we provided them with two pieces of information. First, we informed them that the exact questions they would be answering were randomized and, therefore, that the questions they would read might be different from those that another participant would. Second, we told

[6] We also added a question about happiness as a further control check.

[7] From now on, for the sake of simplicity, we refer to this set of questions as the PANAS questionnaire.

them that they could answer these questions in their native language (the screenshot where we provided this information is in section C.2.1 of the appendix).[8]

In the treatment of Experiment 1, the first question asked participants to recall and list up to five events in which they had experienced angry feelings. The second question asked them to carefully describe one of these events. Similarly, in the treatment of Experiment 2, we asked participants to recall up to five events that made them feel sad, and to write in detail about one of these events.[9] By contrast, participants in the control groups of both experiments were asked to recall up to five things they did earlier in the day, and to describe in detail how they typically spend their evenings.

This method to induce emotions is known as "autobiographical recall" in the psychological literature. It is based on the idea that recalling an event that caused an individual to feel a specific emotion will make that individual recreate and relive that emotion.[10] We chose this method for several reasons. First, ample evidence in the psychology literature shows that autobiographical recall effectively induces anger and sadness.[11] Second, as compared to other methods, autobiographical recall more specifically induces the emotion of interest with limited arousal of related but different emotions (Lerner et al., 2003; Strack, Schwarz and Gschneidinger, 1985; Tiedens and Linton, 2001). Finally, this method allows us to ex-post perform text analyses to further assess the effectiveness of the induction.

Matching Protocol

Following the induction procedure, participants were randomly sorted in groups of three. Each group consisted of three subjects that received the two different inductions used in the experiment; that is, each threesome consisted of either one treatment recipient and two members of the control group, or two treatment recipients and one member of the control group. Importantly, participants were told only that they were matched with two other players. As mentioned above, participants did not know the exact questions we asked other participants in the session, and

[8] Given the relatively large number of subjects who are not native English speakers, we offered this option to encourage writing. In this way, for instance, we could prevent people from not answering because they were not comfortable with writing in English. Only one subject chose to write in a language other than English.

[9] Appendix C.2.1 provides the screenshots of the induction task for each experimental condition. The induction questions are based on Small and Lerner (2008).

[10] Other methods, generally referred to as "mood induction procedures," include asking individuals to look at images, watch video excerpts, listen to music, and imagine certain scenarios. Other methods rely on situational procedures, too (e.g., consumption of bitter drink to induce disgust). For a review see Lench, Flores and Bench (2011) and Westermann et al. (1996).

[11] For a recent review see Siedlecka and Denson (2019).

they were not informed about the matching protocol.

The Beauty Contest Game

Subjects played the p-beauty contest (Nagel, 1995). We follow closely the design in Gill and Prowse (2016). Participants in groups of three played the p-beauty contest for ten rounds with fixed-group matching. In each round, participants had to choose an integer between zero and 100. The participants whose chosen number was closer to 70% of the mean of the three numbers earned £10.00, whereas all others earned nothing. If there was more than one winning number, then the winners equally split the £10.00, while the loser earned nothing. In each round of this game, the unique Nash equilibrium is to choose zero.[12] To avoid wealth effects, we told participants that at the end of the session that only one round would be randomly drawn to count for payments.

In each round, participants had to type the number in a given box; they did not face any time constraint. After all participants in a group had made their choices, each participant was shown the following information about the game in that round: 1) the three chosen numbers in the group; 2) the 70% of the mean of the chosen numbers; 3) the winning number(s); and 4) one's own earnings in that round.[13]

Other Tasks

Upon completion of the game, we asked participants to fill in a set of questionnaires. We asked them to complete the PANAS questionnaire again. We could thus assess the induction procedure as the difference in participants' responses to this questionnaire before and after the induction took place. We then asked them to self-report, on a nine-point scale, the degree to which they experienced different discrete emotions while they were writing about their personal past-life experiences (Rottenberg et al., 2007). Third, in Experiment 1, we asked participants to complete the State-Trait Anger Expression Inventory 2 (STAXI-2) questionnaire (Schamborg, Tully and Browne, 2016), which assesses one's disposition to anger (i.e., anger as a trait). We then asked two final questions. The first asked whether they had previously played the p-beauty contest game (prior to participation in our experiment). The second was a non-incentivized general willingness to take risks question (Dohmen et al., 2011).

[12] See e.g. Gill and Prowse (2016) and López (2001) for the formal proof that takes into account that numbers in the game can only be discrete.

[13] In Appendix C.2.2 we provide the screenshots of the instructions for the p-beauty contest game.

3.2.2 Experimental Procedure

The experimental sessions were conducted from May to October 2019 at the economics laboratory of the University of Warwick. Overall, we recruited 351 subjects through the university's SONA System for recruitment of participants in experiments. We conducted 11 sessions with 171 subjects in Experiment 1, and 12 sessions with 180 subjects in Experiment 2. Sessions lasted roughly 35 minutes. Participants earned an average payment of £8.33 including the show-up fee of £5.00. We coded and conducted the experiment using the oTree software (Chen, Schonger and Wickens, 2016). Tables C.1 and C.2 provide descriptive statistics of the sample. Tables C.3 and C.4 show that there are no significant differences across conditions in observable characteristics in the experiment.

At the onset of each session, subjects were randomly allocated to a computer terminal, and the experimenter read aloud general instructions about the session. After that, detailed instructions about the experimental tasks were shown on the computer screens. A reminder of the instructions for each part of the experiment was shown at the bottom of each page. Participants were encouraged to ask questions to the experimenter at any point.

3.3 Results

In this section we show that the anger treatment induced subjects to play farer from the best response, given the choices of the other two players, and, consistently, led them to earn lower profits, compared to subjects in the control group. The effects of the sadness treatment on guesses and profits are insignificant. Before presenting these results, we analyze whether the induction procedure itself was successful in inducing the emotions of interest.

3.3.1 Emotion Induction

To analyse the emotional content of subjects' answers to the induction questions, we use the Linguistic Inquiry and Word Count (LIWC) 2015 software. The LIWC software reads a specific text, and counts each time a word in the text corresponds to a word present in the built-in dictionary.[14] The dictionary matches each word with psychologically relevant categories (e.g., affect word, social word, etc.). For instance, the word "cried" matches the following dictionary categories: Sadness, Negative Emotion, Overall Affect, Verb, and Past Focus. The software then computes the

[14] The dictionary recognizes about 6,400 English words.

percentage of total words that match one of these categories. Following our example, if the world "cried" were found in the text, the scores for these five categories would increase.

For the purposes of the induction assessment, we look at two categories: Anger and Sadness. Table 3.1 shows some examples of words in these categories, and the number of words included in each of these.

Table 3.1: Anger and Sadness categories in LIWC

	Example	Words in Category
Anger	Hate, Kill, Annoyed	230
Sadness	Crying, Grief, Sad	136

Figure 3.1 shows the results of the text analysis. Subjects, on average, followed the instructions of the induction. The difference in the percentage of angry words in the texts for the treated subjects compared to those in the control group is 2.077, and it is statistically different from zero (p-value<0.001) (top-left panel).[15] Similarly, the subjects treated with the sadness induction used more words associated with sadness than subjects in the control group ($\Delta = 2.628$, p-value<0.001). Importantly, in Experiment 1, the difference in the anger content of the writings between participants in the treated and control groups is significantly larger than the difference in the sadness content between the treated and control group members ($\Delta = 1.633$, p-value<0.001) (top-left panel compared with bottom-left). Similarly, in Experiment 2, the difference in the use of sad words between treated and control group subjects in Experiment 2 is significantly larger than the difference in the angry words between them ($\Delta = 2.094$, p-value<0.001) (top-right panel compared with bottom-right).

These results highlight that, at a minimum, participants wrote about the episodes that they had been asked to address in the exercises.[16] This, in turn, should have caused participants to relive the recalled event and experience once again the emotions related to it (see Section 3.2.1). Our data allows us to look further into this link. We thus analyze self-reported measures of the levels of an array of discrete emotions felt in the experiment.

[15] In this subsection, we compute Mann-Whitney tests to assess the differences in the emotions across conditions and experiments.

[16] In Appendix C.3 we perform further text analyses by measuring total and negative affect in the texts by condition and experiment. Results are similar. Interestingly, if we look at the presence of words indicative of anxiety in the texts, levels are significantly lower compared to those for anger in the treated group of Experiment 1, and for sadness in the treated group of Experiment 2 (see Figure C.3.3).

Figure 3.1: Anger and Sadness measured using text analyses of the inductions

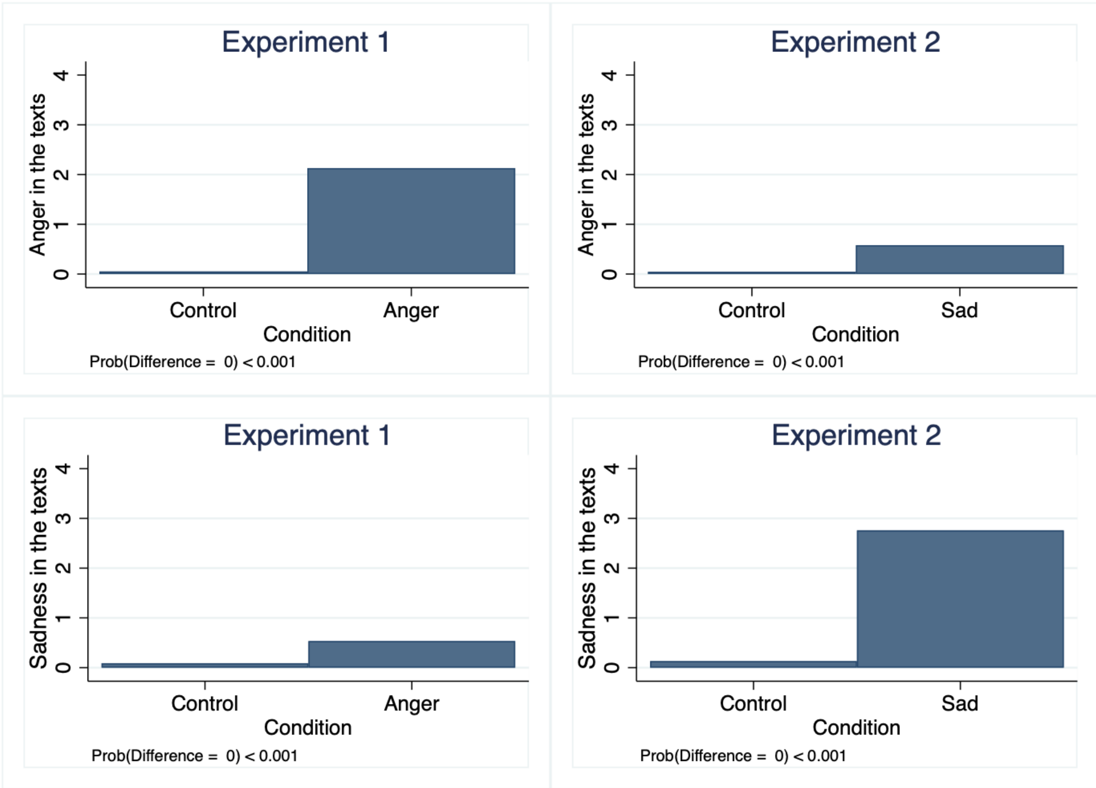

Notes: 1) The bars report the averages of anger and sadness from the subjects' written words for the different inductions. 2) Analysis based on the LIWC2015 (Linguistic Inquiry and Word Count) dictionary (Pennebaker, 2015). 3) The notes report the results of the corresponding Mann-Whitney test.

In particular, we now study before and after responses to the PANAS questionnaire. Figure 3.2 shows the results. There is a significant difference in reported anger before and after the anger treatment, compared to the control ($\Delta = 0.280$, p-value=0.004) (top-left panel), while the difference about sadness is not statistically significant in the same groups ($\Delta = -0.071$, p-value=0.478) (bottom-left panel). The opposite is true in Experiment 2. Here the difference in subjects' reported levels of sadness, before and after the induction, is significantly greater in the treatment compared to the control ($\Delta = 0.263$, p-value=0.014) (bottom-right panel). We also find a positive marginal effect in the difference in reported anger across the two groups ($\Delta = 0.156$, p-value=0.078) (top-right panel).

Figure 3.2: Anger and Sadness felt before and after the induction

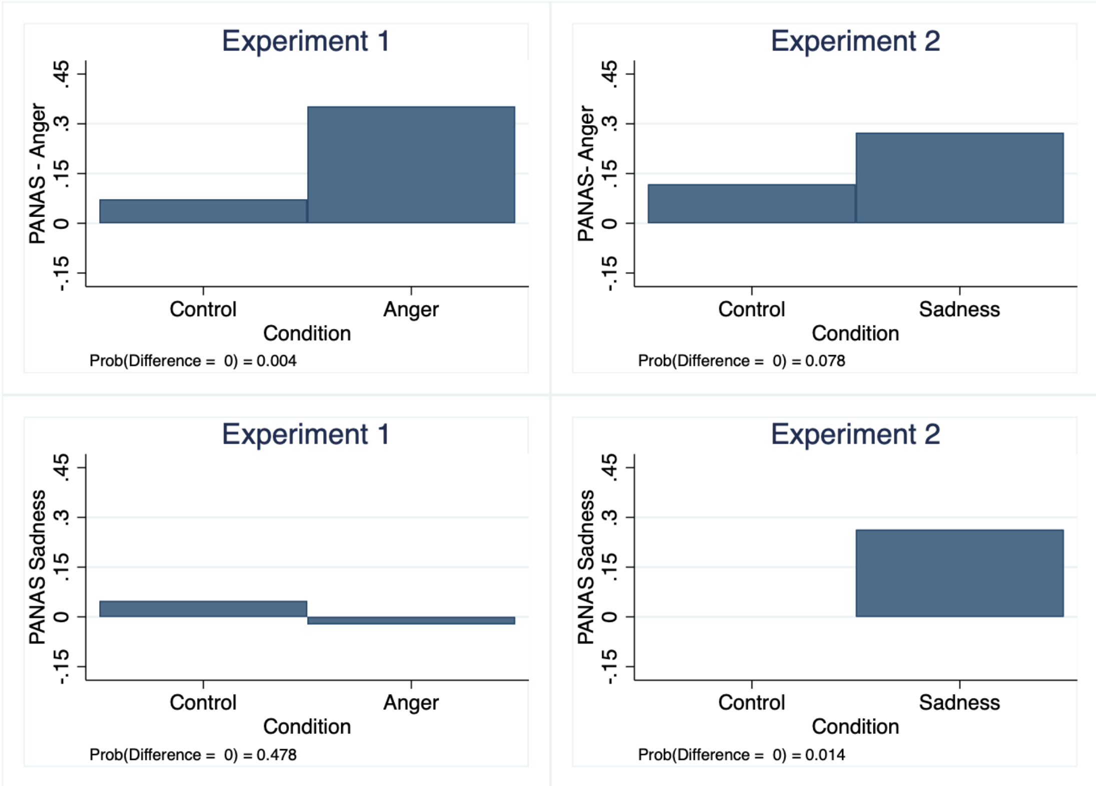

Notes: 1) The bars report the average difference between the levels of anger or sadness participants said they felt before the induction and at the end of the sessions. 2) The question posed was "How much anger (sadness) you are feeling now?". Responses were coded from zero (low) to five (high). 3) The notes report the results of the corresponding Mann-Whitney test.

In Appendix C.4.1 we show that our conclusions of the effectiveness of the induction procedure holds if we look at the answers to the Rottenberg et al. (2007)'s questionnaire. Additionally, through the responses to the first PANAS questionnaire, we can check whether participants' general affect was similar at the baseline by condition and experiment. In Appendix C.4.2 we present the results. We find no differences in either positive or negative affect at the baseline across conditions in the two experiments. This is a further check that randomization into the treatment conditions produced balanced groups. And, more importantly, it shows that our experimental results cannot be driven by differences in affect at the baseline.

3.3.2 p-beauty Contest Game

Having shown that the induction procedure was successful in inducing the emotions of interest, we now analyze the effect of anger and sadness in the p-beauty contest game. This game allows us to assess how the treatments affect strategic skills. These skills are characterized by the capacity of optimally reacting to other individuals' actions and in predicting those actions.[17]

Accordingly, in what follows we assess strategic ability by measuring the capacity to best respond to other people's guesses and the payoffs in the game. They are both measures of subjects' ability to choose their guesses optimally and predicting others' behavior in the game.

Best Responses and Payoffs

We analyze the effects of the treatments on best responses and payoffs in the two experiments. As argued above, these variable proxies for an individual's capacity to understand other players' behaviors, and to respond optimally by following these beliefs.

Given the rules of the game, the best response of subject $i \in 1,2,3$, of group j in round $t \in 1,2...10$ is $BR_{i,j,t} = \frac{0.7}{3}(BR_{i,j,t} + Guess_{j,t,1} + Guess_{j,t,2})$. Where $Guess_{j,t,1}$ and $Guess_{j,t,2}$ are are the guesses made by the other two subjects in the same group play in the current round t. We are interested in the capacity of guessing a number as close as possible to the best response, hence we will consider the difference in absolute value, say $\Delta BR_{i,j,t}$, between the guess of player i,j at time t, $Guess_{i,j,t}$, and BRi,j,t:

$$\Delta BR_{i,j,t} = |Guess_{i,j,t} - Bri,j,t|. \tag{3.1}$$

We show the average of this difference over each round and treatment in Figure 3.3. The top-left panel of Figure 3.3 shows that the anger treatment results in a higher absolute distance from the best response in every round. On the other hand, the sadness treatment does not have any clear effect as we can observe from the top-right panel. As shown by the bottom panels of the figure, only the anger treatment significantly increases this distance when we pool all data together.

[17] This last capacity is often referred as theory of mind, the process of mental modeling about others' beliefs and actions (Coricelli and Nagel, 2009). That is, when thinking about others' characteristics and beliefs, people build mental models that change and develop through continuous interactions, and these can be used to anticipate others' behavior.

Figure 3.3: The effect of anger and sadness on the average distance in absolute value from the best response

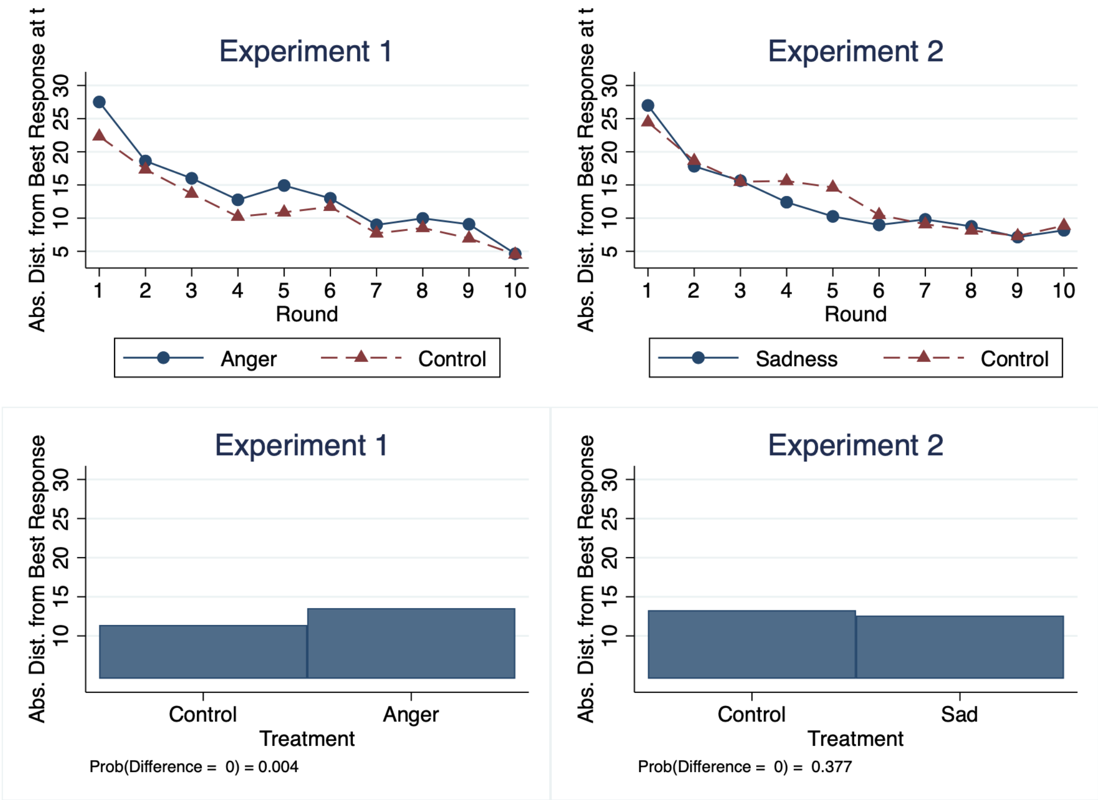

Notes: 1) The lines in the top panels report the average distance in absolute value from the best response for each round of play by condition and experiment. 2) The bottom panels report the average distance in absolute value from the best response across all rounds by condition and experiment. 3) The notes report the results of the corresponding t-test.

In the first round, the distance from the best response for those who experienced the anger treatment is 27.51, significantly larger than the 22.32 found in the control group (with p-value =0.0204). We also find lower average distance for those who experienced the sadness treatment compared to those in the control group, but the effect is substantially smaller (24.46 vs. 26.99) and insignificant (p-value= 0.1741).

The previous analysis on average distance from the best response across all rounds does not consider that guesses are influenced by past behavior. Therefore, in order to take into account previous game play and group fixed effects we estimate the following model:

$$\Delta BR_{i,j,t} = \beta_0 + \beta_1 Treatment_i + \beta_2 AverageGuess_{j,t-1} + \beta_3 t + \gamma_j + \epsilon_{i,j,t}; \qquad (3.2)$$

where, as usual, i indicates the subject in group j, while t is the round of play. Our independent variable of interest is $Treatment_i$, which is a dummy variable indicating the emotion treatment individual i received in one of the two experiments. Control variables include: $AverageGuess_{j,t-1}$ that is the average guess in the previous round, t is the round of play, γ_j is the group-level effect, while $\epsilon_{i,j,t}$ is the error term.

We estimate Equation (3.2) by using an OLS model with group fixed effects. We cluster standard errors at the group level. The results are reported in Table 3.2. Column (1) reports the results for Experiment 1; Column (2) reports those for Experiment 2; Column (3) reports the combination of the two. The anger treatment has a positive and significant effect on distance from the best response, thus a negative effect on strategic ability. The deviation from the best response of subjects who experienced the anger treatment is about 1.9 units higher on average, compared to subjects' guesses in the control. The sadness treatment has an insignificant effect of the opposite sign.

Table 3.2: The effect of the treatment on the distance from the best response (in absolute value) in both experiments

	Experiment 1 Distance BR Abs. Val. b/se	Experiment 2 Distance BR Abs. Val. b/se	Experiments 1 & 2 Distance BR Abs. Val. b/se
Anger Treatment	1.867**		1.867**
	(0.795)		(0.791)
Sadness Treatment		-0.552	–0.552
		(0.980)	(0.976)
Average Guess at $t-1$	-0.129	-0.006	-0.010
	(0.016)	(0.015)	(0.011)
Round	-1.557***	–1.378***	–1.469***
	(0.308)	(0.257)	(0.201)
Group Exp FE	Yes	Yes	Yes
N	1539	1620	3159
Individuals	171	180	351
R2	0.164	0.150	0.156

Notes: 1) OLS estimator; 2) Standard errors (shown in parentheses) are clustered at the group level. 3) * p-value<0.1, ** p-value<0.05, *** p-value<0.01.

Consistent with the previous results, the top panels of Figure 3.4 show that the anger treatment results in lower levels of payoffs in almost every round. The sadness treatment seems to have some detrimental impact on payoffs, though the results are not as clear cut as those that stem from the anger treatment. As shown by the bottom panels of the figure, both treatments decrease payoffs. The effect of the anger treatment, however, is larger in magnitude.

Figure 3.4: The effect of anger and sadness on payoffs in both experiments

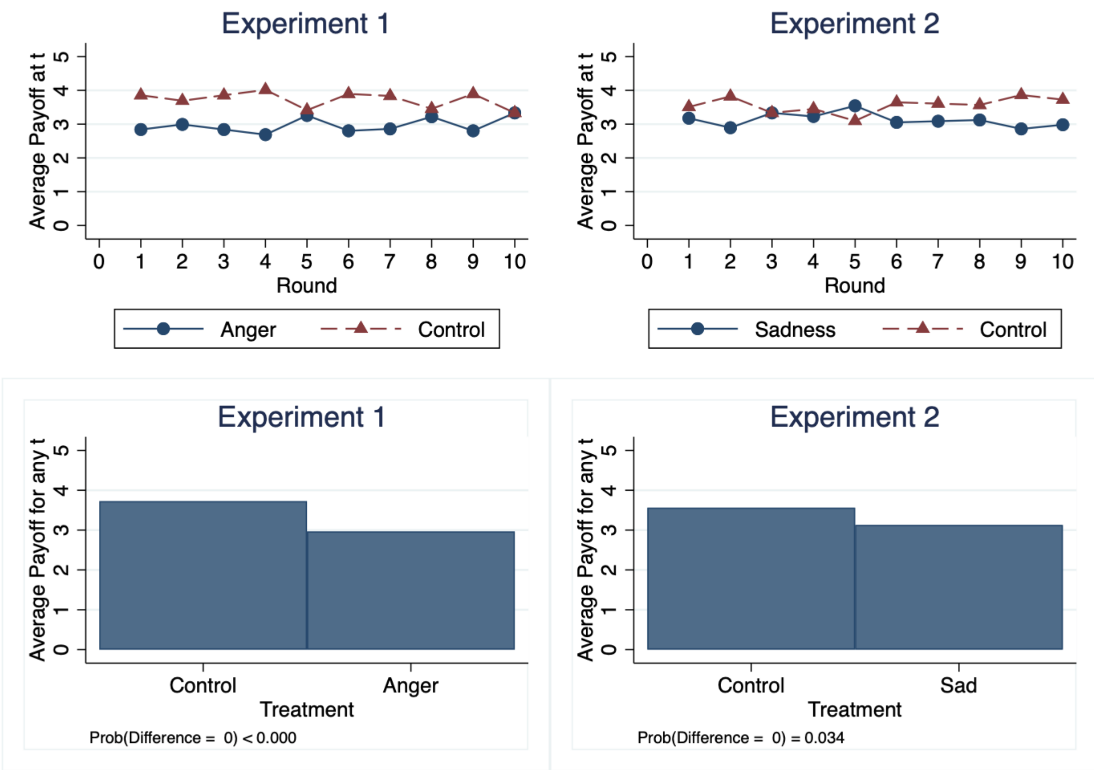

Notes: 1) The lines in the top panels report average payoffs for each round of play by condition and experiment. 2) The bottom panels report average payoffs across all rounds by condition and experiment. 3) The notes report the results of the corresponding t-test.

In the first round, the average payoff for those who experienced the anger treatment is £2.84, smaller than £3.86 in the control group, but the difference is insignificant (p-value=0.139). We also find lower average payoffs for those who experienced the sadness treatment compared to those in the control group, but the effect is substantially smaller (£3.18 vs. £3.51, p-value=0.628).

As before, we perform OLS regressions to analyze the impact of the treatments on payoffs, while also taking into account previous game play. We estimate the following model:

$$\begin{aligned} Payoff_{i,j,t} = \beta_0 + \beta_1 Treatment_i + \beta_2 AverageGuess_{j,t-1} \\ + \beta_3 Guess_{j,t,1} + \beta_4 Guess_{j,t,2} + \gamma_j + \epsilon_{i,j,t}; \end{aligned} \tag{3.3}$$

This model is similar as the one in Equation (3.2) with two further controls: the guesses made by the other two subjects in the same group play in the current round t, $Guess_{j,t,1}$ and $Guess_{j,t,2}$. We need these variables because the payoffs are influenced by the guesses of the two other players. We cluster standard errors at the group level.

We present the estimation in Table 3.3: Column (1) reports the results for Experiment 1, Column (2) reports those for Experiment 2. Column (3) combines the two. Column (1) shows that the anger treatment significantly reduces payoffs per round by about £0.66 (p-value=0.047). The coefficient of the sadness treatment is smaller (£0.53) and non-significant in Experiment 2 (p-value=0.147). We find the same results when the two experiments are pooled together in a unique regression (Column (3)).

Table 3.3: The effect of the treatment on payoffs in both experiments

	Experiment 1 Payoff b/se	Experiment 2 Payoff b/se	Experiments 1 & 2 Payoff b/se
Anger Treatment	–0.659**		–0.700**
	(0.324)		(0.323)
Sadness Treatment		–0.527	–0.530
		(0.359)	(0.358)
Guess other player (1) at t	0.060***	0.042***	0.049***
	(0.009)	(0.007)	(0.006)
Guess other player (2) at t	0.034***	0.051***	0.042***
	(0.007)	(0.006)	(0.005)
Average Guess at t-1	–0.018***	–0.019***	–0.018***
	(0.003)	(0.002)	(0.002)
Group Exp FE	Yes	Yes	Yes
N	1539	1620	3159
Individuals	171	180	351
R2	0.056	0.055	0.054

Notes: 1) OLS estimator; 2) Standard errors (shown in parentheses) are clustered at the group level. 3) * p-value<0.1, ** p-value<0.05, *** p-value<0.01.

The results in this section show that anger impairs the capacity of thinking strategically, while sadness does not seem to have any clear effect. We used the distance from the best response level and the payoffs to measure the strategic ability.

For completeness, in section C.4.3 of the appendix we also show that angry subjects play further away from 0, the unique Nash equilibrium level.[18]

[18] This is a measure often used in the literature to assess the level of strategic sophistication of the subjects in a Beauty Contest game, (e.g. Eyster, 2019, for a survey). However, this measure does not take into account the capacity of the subjects to predict other subjects' guesses (what we called theory of mind). For this reason we preferred the absolute distance from the best response as an index of strategic ability in the current work.

3.4 Structural Level-k Model

In the beauty contest game, players have incentive to win the game and not to play the Nash equilibrium number, zero. As it is well known, this discrepancy is essentially due to the fact that players do not typically expect that other players are fully rational (i.e. the common knowledge or rationality condition is violated). Therefore the ability to win is determined by the capacity of forming higher order beliefs *plus* the capacity of correctly predicting other people's behaviour. In order to understand the effect of anger on the capacity of forming higher order beliefs only, we present and estimate a finite mixture model, in which individuals are grouped according to the different latent k-rule chosen. Using this model, we estimate the impact of anger on players' level-k rule. We first describe the model. We then present the estimation strategy and the results.

3.4.1 The Model

We assume that choices $x_{i,g,t}$, in which i is the subject in group g and round t, are independent draws over rounds and subjects. Let $k_{i,g,t} \in [0, 1, ..\overline{k}]$ be the rule followed by subject i in group g and round t.

When the choice rule is $k = 0$, we assume that individuals choose randomly. Thus, the choice per individual and round is uniformly distributed, with probability:

$$Pr(x \mid k_{i,g,t} = 0) = 1/101. \tag{3.4}$$

When, on the other hand, the choice rule is $k > 0$, we assume that the choice of subject i in group g and round t follows a normal distribution $g(x \mid \mu_{k,t,g}, \sigma)$, characterized by the mean $\mu_{k,t,g}$ and the variance σ.

- For any round $t > 1$, let $\overline{x}_{g,t-1}$ be the average choice in group g for round $t-1$. We then assume that individuals at round t start their iteration using the mean guess in their group, g, in round $t-1$ (that they directly observe at the end of round $t-1$ of the experiment).[19] Accordingly, because subjects choosing a strategic rule k best respond to the ones choosing a rule $k-1$, we assume that the mean of the distribution of their guesses is $\mu_{k,t,g} = (\frac{7}{10})^k \cdot \overline{x}_{g,t-1}$.
- In round $t = 1$, subjects choosing $k = 1$ best respond to the one with $k = 0$ by calculating the mean of the uniform distribution of $k = 0$'s choices; hence

[19] Gill and Prowse (2016) find that assuming that subjects do not take into account the effect of their own guesses on the mean leads to a better fit of the estimation. We then follow them in this assumption.

$\overline{x}_{g,0} = 50$ and $\mu_{1,1,g} = (\frac{7}{10}) \cdot 50$. Subjects choosing a rule $k = 2$ best respond to the ones choosing $k = 1$, with $\mu_{2,1,g} = (\frac{7}{10})^2 \cdot 50$, and so on.

As a result, the probability of choosing any x for an individual i in g at round t and a rule $k_{i,g,t} > 0$ is:

$$Pr(x \mid k_{i,g,t} > 0) = g(x \mid \mu_{k,t,g}, \sigma) \tag{3.5}$$

Let $z(k)$ be the distribution of the choice rule among the different subjects in the different rounds. The unconditional probability of any choice x of any i in group g at round t is then:

$$Pr(x) = \sum_{k>0}^{\overline{k}} z(k)g(x \mid \mu_{k,t,g}, \sigma) + z_{k=0}(k)(1/101) \tag{3.6}$$

3.4.2 Estimation Strategy

Using the experimental data that includes 351 subjects, who make a total of 3,510 choices across the 10 rounds of play, we estimate the parameter vector $\theta = [\sigma, z(k)]$ for each treatment and experiment.

Here, we assume that there are up to level-4 subjects (hence $\overline{k} = 4$) and that the distributions $g(x \mid \mu_{k,t,g}, \sigma)$ are normally distributed with mean, $\mu_{k,t,g}$, and variance, σ. Thus, θ consists of five parameters (remember that for $k = 0$ the distribution of x is uniform).

Given our assumptions, the probability of all observed choices for any individual and round, $\mathbf{x}$, is $Pr(\mathbf{x}) = \prod_{t=1}^{10} \prod_{g=1}^{N} \prod_{i=1}^{3} Pr(x_{i,g,t})$. The likelihood we maximize is then: $L(\theta, \mathbf{x}) = Pr(\mathbf{x})$; where $\theta = [\sigma, z(k)]$.

We maximize the sample log likelihood function using a standard Matlab routine. Table 3.4 shows the estimated level-k types by condition in Experiment 1 and Experiment 2.[20] We note that the anger treatment increases the share of level-0 play of about 35 percent (from 0.253 to 0.343) against decreases of about 4 percent (from 0.532 to 0.509) and 50 percent (from 0.177 to 0.089) in the shares of level-1 and level-2 play. Overall, the anger treatment leads to a decrease in the average level-k from 1.001 to 0.864. We do not observe the same patterns in Experiment 2. In fact, the sadness treatment, if anything, leads to a lower share of level-0 choices.

In Table C.7 of the appendix, we present the results using an alternative hypothesis on the distributional form of choices. We assume that choices follow a Poisson distribution. We qualitatively observe similar results.

[20]We omit to report the estimated variances for expositional simplicity. However, they are available upon request.

Table 3.4: Estimated Level-k types by treatment and experiment

	Experiment 1		Experiment 2	
	Anger	Control	Sadness	Control
Level 0	0.343	0.253	0.218	0.302
Level 1	0.509	0.532	0.691	0.538
Level 2	0.089	0.177	0.048	0.118
Level 3	0.059	0.038	0.043	0.039
Level 4	0.000	0.000	0.000	0.002
σ	3.501	3.967	4.806	3.410
Log likelihood	-3,408	-3,114	-3,547	-3,216
Average Level-k	0.864	1.001	0.915	0.901

3.5 Anger and Bargaining Power

The above results can be puzzling. If anger negatively affects the level of strategic sophistication, why is it so pervasive in human behaviour? In what follows, we will argue that anger represents a powerful commitment device. We consider a simple sequential game, the ultimatum game, to illustrate this argument. In the ultimatum game the proposer offers a share of $x \in (0, 100)$ to a responder, who decides whether to accept or reject the offer. If the responder accepts, she earns x and the proposer earns $100 - x$. Otherwise, they respectively receive the outside payoffs of V_R and V_P.

We proceed by following Ho and Su (2013), who analyze a model in which the level-k logic is applied to sequential games and to the ultimatum game in particular.[21] They argue that their model implies that:

- Level-0 players are assumed to choose a number between zero and 100 randomly. For the proposer, this randomly chosen number is the initial demand, while for the responder this number is the acceptance threshold (i.e., only offers that are above this threshold are accepted).

- Level-1 players best respond to level-0 players' randomization strategies. The responder accepts any offer greater than V_R. The proposer chooses x so to maximize the expected payoff:

$$\frac{x}{100}(100 - x) + \frac{100 - x}{100} V_P. \tag{3.7}$$

- Level k$\geq$ 2 players best respond to level k-1 players. The responder accepts any offer greater than V_R. Hence the proposer offers the minimum acceptable

[21] See pages 456-467.

amount to a level k-1 responder, which is V_R.

The implications of our results in this environment can be appreciated as follows. Assume that there is a threshold $\overline{x} > 0$, such that if any offer $x < \overline{x}$, the responder grows angry, and thus becomes level-0 player.[22] We also assume that the proposer at the beginning of the process chooses a level k > 1. Finally, for expositional simplicity let $V_R = V_P = 0$.

Therefore, a level $k > 1$ proposer knows that if he offers $x \geq \overline{x}$, this will be accepted for sure by any level k-1 responder, obtaining a payoff $100 - x$. If she offers $x < \overline{x}$, then the responder will use the level-0 rule of decision and will accept the offer with probability $x/100$. Given that $100 - x \geq \frac{x}{100}(100 - x)$ for all $x \in (0, 100)$. The proposer's optimal choice is $x^* = \overline{x}$, hence the responder payoff is $\overline{x}$.

As $\overline{x}$ can be thought of as a measure of propensity to anger that can also be affected by external factors, the higher $\overline{x}$ is, the higher the payoff for the responder. This can explain why anger is so pervasive in human behavior. In fact, a plausible explanation relies on the notion that anger, like other emotions, can serve as a credible commitment device in situations of conflict (e.g. Elster, 1998; Frank, 1987, 1988; Hirshleifer, 1987), and thus can lead to greater evolutionary success in strategic interactions.

Furthermore, it is interesting to note that in this model, a possible strategy for the responder would be to try to convince the proposer than he is angry (whether or not this is true), so that $\overline{x}$ is high and he is easy to upset. This strategy explains the to so-called *madman theory*, where signaling anger is a way to signal irrationality.[23]

The predictions of this model are generally consistent with findings in laboratory experiments, in which negative mood induction increases the probability of rejection of unfair offers (see e.g. Forgas and Tan, 2013). It is important to note than in any sequential game, such as the ultimatum game, it is virtually impossible to distinguish the anger effect of the desire to punish the opponent (i.e., due to negative social preferences) from the effect on strategic sophistication tested in this paper with the beauty contest game. Thus, this is a key factor motivating our experimental strategy.

[22]The existence of such a threshold is consistent with the framework of psychological game theory (see e.g. Geanakoplos, Pearce and Stacchetti, 1989; Battigalli and Dufwenberg, 2009), in which individuals can become angry when their beliefs about others' choices or beliefs are not fulfilled. Therefore, we can interpret $\overline{x}$ as the minimal offer the responder expects to receive.

[23]Niccolò Machiavelli argued that sometimes it is "a very wise thing to simulate madness" (Machiavelli, 2009, book 3, chapter 2).

3.6 Mechanism and Potential Confounding Factors

Our results show that anger negatively affects strategic reasoning and theory of mind. Although, in this paper we are agnostic about the exact cognitive mechanism, it is instructive to briefly discuss it.

One possibility is that anger leads the player to use System I thinking, which is faster and more instinctive, hence less elaborated. Another possibility is disengagement, angry players (but not sad ones) may feel disinterested on the ongoing process they are involved and take a less elaborated decision. A third possibility is that anger can more generally impairs cognitive processes. Obviously, these mechanisms are not mutually exclusive and can reinforce each other; the bottom line is that they all point toward less elaboration in the decision making process.

A potential concern of the experiment is that the induction changed participants' beliefs about the other participants' play in the game. For instance, participants could have anticipated the effects of the induction and adapted their guesses in the game accordingly. We find this prospect unlikely however. This is because in the instructions of the induction we did not inform the participants about the nature of other subjects' induction; furthermore we explicitly told the participants that the questions they would receive in the induction were not the same across participants. Moreover, these second-order beliefs should only matter for first-round guesses, while they should be much less relevant once subjects have played the first round and have seen others' guesses.

3.7 Conclusions

Our results provide strong evidence that anger impairs the capacity to think strategically. Our findings show that angry participants make significantly worse choices in a p-beauty contest game. Angry players earned lower profits than players who did not participate in the exercise to elicit anger. These players also use level-0 thinking more often. Our follow-up experiment, which exposed a group of players to exercises to induce sadness, does not produce the same effects.

The fact that anger is so pervasive in human relationships and has a negative effect on the capacity of strategic thinking and on cooperation between individuals is puzzling. However, the literature (e.g. Elster, 1998; Frank, 1987, 1988; Hirshleifer, 1987) has emphasized that anger can serve as an efficient commitment device in strategic interactions. We discussed the implications of the link of this together with our results in a sequential game. Here, the propensity to become angry and the inability to think strategically (or be seen as someone who does not think strategically) can represent an effective commitment device able to increase individuals' bargaining power.

Our results have implications for behavioral policies at the individual level. Ex-post anger is detrimental, so it is optimal to control it. On the other side, ex-ante anger (i.e., before the event/action takes place), and showing high anger propensity – the so-called madman theory – may represent a bargaining advantage.

As we mentioned in the introduction, the incidental effect of anger on cognitive abilities may represent a negative externality for those economies in which subjects are more exposed to shocks leading to anger. Proto, Rustichini and Sofianos (2019) show that cooperation rates on a non zero-sum complex game, such as the repeated prisoner's dilemma, positively depend on the cognitive abilities of the players. Following this finding and the results obtained in the current article, it is natural to hypothesize that anger has a detrimental effect on cooperation. This hypothesis also finds empirical confirmation in a repeated prisoner's dilemma game (Castagnetti, Massaro and Proto, 2018), in which participants in the treatment were induced to feel anger through the use of a standard video induction procedure.

We can then argue that the negative effect of anger on strategic reasoning can represent a negative externality for an economy and a society in aggregate because it can potentially reduce cooperation in situations in which cooperation is likely beneficial. Therefore, anger can generate self-sustaining, vicious cycles, particularly in environments in which anger-producing events are more frequent – such as in poorer countries, during times of negative economic shocks, and among poorer, disadvantaged socioeconomic classes.

Bibliography

Abeler, Johannes, Armin Falk, Lorenz Goette and David Huffman. 2011. "Reference points and effort provision." *American Economic Review* 101(2):470–92.

Adams, Renée B and Patricia Funk. 2012. "Beyond the glass ceiling: Does gender matter?" *Management science* 58(2):219–235.

Akerlof, Robert. 2016. "Anger and enforcement." *Journal of Economic Behavior & Organization* 126:110–124.

Alan, Sule, Seda Ertac and Ipek Mumcu. 2018. "Gender stereotypes in the classroom and effects on achievement." *Review of Economics and Statistics* 100(5):876–890.

Alicke, Mark D and Constantine Sedikides. 2009. "Self-enhancement and self-protection: What they are and what they do." *European Review of Social Psychology* 20(1):1–48.

Anderson, Eric T and Duncan I Simester. 2010. "Price stickiness and customer antagonism." *The Quarterly Journal of Economics* 125(2):729–765.

Andrade, Eduardo B and Dan Ariely. 2009. "The enduring impact of transient emotions on decision making." *Organizational Behavior and Human Decision Processes* 109(1):1–8.

Azmat, Ghazala and Barbara Petrongolo. 2014. "Gender and the labor market: What have we learned from field and lab experiments?" *Labour Economics* 30:32–40.

Battigalli, Pierpaolo and Martin Dufwenberg. 2009. "Dynamic psychological games." *Journal of Economic Theory* 144(1):1–35.

Battigalli, Pierpaolo, Martin Dufwenberg and Alec Smith. 2019. "Frustration, aggression, and anger in leader-follower games." *Games and Economic Behavior* 117:15–39.

Becker, Gary S. 2010. *The Economics of Discrimination.* University of Chicago Press.

Bénabou, Roland and Jean Tirole. 2002. "Self-confidence and personal motivation." *The Quarterly Journal of Economics* 117(3):871–915.

Bénabou, Roland and Jean Tirole. 2016. "Mindful economics: The production, consumption, and value of beliefs." *Journal of Economic Perspectives* 30(3):141–64.

Bénabou, Roland and Jean Tirole's. 2016. "Mindful economics: The production, consumption, and value of beliefs." *Journal of Economic Perspectives* 30(3):141–64.

BenYishay, Ariel, Maria Jones, Florence Kondylis and Ahmed Mushfiq Mobarak. 2020. "Gender gaps in technology diffusion." *Journal of Development Economics* 143:102380.

Bertrand, Marianne. 2011. New perspectives on gender. In *Handbook of Labor Economics.* Vol. 4 Elsevier pp. 1543–1590.

Bertrand, Marianne and Esther Duflo. 2017. Field experiments on discrimination. In *Handbook of Economic Field Experiments.* Vol. 1 Elsevier pp. 309–393.

Beyer, Sylvia. 1998. "Gender differences in causal attributions by college students of performance on course examinations." *Current Psychology* 17(4):346–358.

Blume, Andreas, Ernest K Lai and Wooyoung Lim. 2020. "Strategic information transmission: A survey of experiments and theoretical foundations." *Report.[1457]* .

Bohnet, Iris, Alexandra Van Geen and Max Bazerman. 2016. "When performance trumps gender bias: Joint vs. separate evaluation." *Management Science* 62(5):1225–1234.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli and Andrei Shleifer. 2019. "Beliefs about gender." *American Economic Review* 109(3):739–73.

Brunnermeier, Markus K and Jonathan A Parker. 2005. "Optimal expectations." *American Economic Review* 95(4):1092–1118.

Buser, Thomas, Leonie Gerhards and Joël van der Weele. 2018. "Responsiveness to feedback as a personal trait." *Journal of Risk and Uncertainty* 56(2):165–192.

Cai, Hongbin and Joseph Tao-Yi Wang. 2006. "Overcommunication in strategic information transmission games." *Games and Economic Behavior* 56(1):7–36.

Card, David and Gordon B Dahl. 2011. "Family violence and football: The effect of unexpected emotional cues on violent behavior." *The Quarterly Journal of Economics* 126(1):103–143.

Carlana, Michela. 2019. "Implicit stereotypes: Evidence from teachers' gender bias." *The Quarterly Journal of Economics* 134(3):1163–1224.

Carpenter, Jeffrey, Michael Graham and Jesse Wolf. 2013. "Cognitive ability and strategic sophistication." *Games and Economic Behavior* 80:115–130.

Carpenter, Jeffrey P and Peter Hans Matthews. 2012. "Norm enforcement: anger, indignation, or reciprocity?" *Journal of the European Economic Association* 10(3):555–572.

Carpenter, Patricia A., Marcel A. Just and Peter Shell. 1990. "What one intelligence test measures: A theoretical account of the processing in the Raven Progressive Matrices Test." *Psychological review* 97(3):404.

Castagnetti, Alessandro and Giovanni Burro. 2021. "Will I tell you that you are smart (dumb)? Deceiving Others about their IQ or about a Random Draw." Working Paper.

Castagnetti, Alessandro and Renke Schmacker. 2020. "Protecting the Ego: Motivated Information Selection." Working Paper.

Castagnetti, Sergio Alessandro, Sebastiano Massaro and Eugenio Proto. 2018. The Influence of Anger on Strategic Cooperative Interactions. In *Academy of Management Proceedings*. Vol. 2018 Academy of Management Briarcliff Manor, NY 10510 p. 14162.

Chen, Daniel L, Martin Schonger and Chris Wickens. 2016. "oTree—An open-source platform for laboratory, online, and field experiments." *Journal of Behavioral and Experimental Finance* 9:88–97.

Chew, Soo Hong, Wei Huang and Xiaojian Zhao. 2018. "Motivated false memory." *Working Paper* .

Coffman, Katherine B, Manuela Collis and Leena Kulkarni. 2019. "Stereotypes and Belief Updating." *Working paper* .

Coffman, Katherine Baldiga. 2014. "Evidence on self-stereotyping and the contribution of ideas." *The Quarterly Journal of Economics* 129(4):1625–1660.

Coricelli, Giorgio and Rosemarie Nagel. 2009. "Neural correlates of depth of strategic reasoning in medial prefrontal cortex." *Proceedings of the National Academy of Sciences* 106(23):9163–9168.

Coutts, Alexander. 2019. "Good news and bad news are still news: Experimental evidence on belief updating." *Experimental Economics* 22(2):369–395.

Coutts, Alexander, L Gerhards and Z Murad. 2019. No one to blame: Self-attribution bias in updating with two-dimensional uncertainty. Technical report Working Paper.

Crawford, Vincent P and Joel Sobel. 1982. "Strategic information transmission." *Econometrica* 50(6):1431–1451.

Croson, Rachel and Uri Gneezy. 2009. "Gender differences in preferences." *Journal of Economic Literature* 47(2):448–74.

Danz, David, Alistair J Wilson and Lise Vesterlund. 2020. "Belief Elicitation: Limiting Truth Telling with Information or Incentives." *CESifo Working Paper* .

Deaux, Kay and Tim Emswiller. 1974. "Explanations of successful performance on sex-linked tasks: What is skill for the male is luck for the female." *Journal of Personality and Social Psychology* 29(1):80–85.

Ditto, Peter H and David F Lopez. 1992. "Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions." *Journal of Personality and Social Psychology* 63(4):568.

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp and Gert G Wagner. 2011. "Individual risk attitudes: Measurement, determinants, and behavioral consequences." *Journal of the European Economic Association* 9(3):522–550.

Dollard, John, Neal E Miller, Leonard W Doob, Orval Hobart Mowrer and Robert R Sears. 1939. "Frustration and Aggression.".

Duffy, John and Rosemarie Nagel. 1997. "On the robustness of behaviour in experimental "beauty contest" games." *The Economic Journal* 107(445):1684–1700.

Dweck, Carol S, William Davidson, Sharon Nelson and Bradley Enna. 1978. "Sex differences in learned helplessness: II. The contingencies of evaluative feedback in the classroom and III. An experimental analysis." *Developmental Psychology* 14(3):268–276.

Egan, Mark L, Gregor Matvos and Amit Seru. 2017. "When Harry fired Sally: The double standard in punishing misconduct." *National Bureau of Economic Research Working Paper No. 23242* .

Eil, David and Justin M Rao. 2011. "The good news-bad news effect: Asymmetric processing of objective information about yourself." *American Economic Journal: Microeconomics* 3(2):114–38.

Ekman, Paul. 1999. "Basic emotions." *Handbook of cognition and emotion* 98(45-60):16.

Elster, Jon. 1998. "Emotions and Economic Theory." *Journal of Economic Literature* 36(1):47–74.

Engelmann, Jan B, Todd A Hare, Andrew S Fox, Regina C Lapate, Alexander J Shackman and Richard J Davidson. 2018. "Emotions can bias decision-making processes by promoting specific behavioral tendencies.".

Erat, Sanjiv and Uri Gneezy. 2012. "White lies." *Management Science* 58(4):723–733.

Erkal, Nisvan, Lata Gangadharan and Nikos Nikiforakis. 2011. "Relative earnings and giving in a real-effort experiment." *American Economic Review* 101(7):3330–48.

Ertac, Seda. 2011. "Does self-relevance affect information processing? Experimental evidence on the response to performance and non-performance feedback." *Journal of Economic Behavior & Organization* 80(3):532–545.

Espinoza, Penelope, Ana B Arêas da Luz Fontes and Clarissa J Arms-Chavez. 2014. "Attributional gender bias: Teachers' ability and effort explanations for students' math performance." *Social Psychology of Education* 17(1):105–126.

Exley, Christine L and Judd B Kessler. 2018. Motivated Errors. Technical report Working Paper.

Eyster, Erik. 2019. Errors in strategic reasoning. In *Handbook of Behavioral Economics: Applications and Foundations 1*. Vol. 2 Elsevier pp. 187–259.

Fennema, Elizabeth, Penelope L Peterson, Thomas P Carpenter and Cheryl A Lubinski. 1990. "Teachers' attributions and beliefs about girls, boys, and mathematics." *Educational Studies in Mathematics* 21(1):55–69.

Flory, Jeffrey A, Andreas Leibbrandt and John A List. 2014. "Do competitive workplaces deter female workers? A large-scale natural field experiment on job entry decisions." *The Review of Economic Studies* 82(1):122–155.

Forgas, Joseph P and Hui Bing Tan. 2013. "Mood effects on selfishness versus fairness: affective influences on social decisions in the ultimatum game." *Social Cognition* 31(4):504.

Frank, Robert H. 1987. "If Homo Economicus Could Choose His Own Utility Function, Would He Want One with a Conscience?" *American Economic Review* 77(4):593–604.

Frank, Robert H. 1988. *Passions within reason: The strategic role of the emotions.* WW Norton & Co.

Geanakoplos, John, David Pearce and Ennio Stacchetti. 1989. "Psychological games and sequential rationality." *Games and economic Behavior* 1(1):60–79.

Gill, David and Victoria Prowse. 2016. "Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis." *Journal of Political Economy* 124(6):1619–1676.

Glover, Dylan, Amanda Pallais and William Pariente. 2017. "Discrimination as a self-fulfilling prophecy: Evidence from French grocery stores." *The Quarterly Journal of Economics* 132(3):1219–1260.

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

Gneezy, Uri and Alex Imas. 2014. "Materazzi effect and the strategic use of anger in competitive interactions." *Proceedings of the National Academy of Sciences* 111(4):1334–1337.

Gneezy, Uri, Christina Gravert, Silvia Saccardo and Franziska Tausch. 2017. "A must lie situation – avoiding giving negative feedback." *Games and Economic Behavior* 102:445–454.

Gurdal, Mehmet Y, Joshua B Miller and Aldo Rustichini. 2013. "Why blame?" *Journal of Political Economy* 121(6):1205–1247.

Hill, Miriam E and Martha Augoustinos. 1997. "Re-examining gender bias in achievement attributions." *Australian Journal of Psychology* 49(2):85–90.

Hirshleifer, Jack. 1987. "On the emotions as guarantors of threats and promises." *The Dark Side of the Force* pp. 198–219.

Ho, Teck-Hua and Catherine Yeung. 2014. "Giving feedback to clients." *Management Science* 60(8):1926–1944.

Ho, Teck-Hua and Xuanming Su. 2013. "A dynamic level-k model in sequential games." *Management Science* 59(2):452–469.

Hossain, Tanjim and Ryo Okui. 2013. "The binarized scoring rule." *Review of Economic Studies* 80(3):984–1001.

Jin, Ginger Zhe, Michael Luca and Daniel J Martin. 2018. Complex disclosure. Technical report National Bureau of Economic Research.

Jin, Ginger Zhe, Michael Luca and Daniel Martin. 2015. Is no news (perceived as) bad news? An experimental investigation of information disclosure. Technical report National Bureau of Economic Research.

Jones, E. E. and V. A. Harris. 1967. "The attribution of attitudes." *Journal of Experimental Social Psychology* 3:1–24.

Keller, Carmen. 2001. "Effect of teachers' stereotyping on students' stereotyping of mathematics as a male domain." *The Journal of Social Psychology* 141(2):165–173.

Koren, Miklós and Silvana Tenreyro. 2007. "Volatility and development." *The Quarterly Journal of Economics* 122(1):243–287.

Köszegi, Botond. 2006. "Ego utility, overconfidence, and task choice." *Journal of the European Economic Association* 4(4):673–707.

Kunda, Ziva. 1990. "The case for motivated reasoning." *Psychological bulletin* 108(3):480.

Landsman, Rachel. 2019. Topics in Labor and Experimental Economics PhD thesis University of Pittsburgh.

Leibbrandt, Andreas, Liang Choon Wang and Cordelia Foo. 2018. "Gender quotas, competitions, and peer review: Experimental evidence on the backlash against women." *Management Science* 64(8):3501–3516.

Lench, Heather C, Sarah A Flores and Shane W Bench. 2011. "Discrete emotions predict changes in cognition, judgment, experience, behavior, and physiology: a meta-analysis of experimental emotion elicitations." *Psychological Bulletin* 137(5):834.

Lerner, Jennifer S, Deborah A Small and George Loewenstein. 2004. "Heart strings and purse strings: Carryover effects of emotions on economic decisions." *Psychological Science* 15(5):337–341.

Lerner, Jennifer S, Roxana M Gonzalez, Deborah A Small and Baruch Fischhoff. 2003. "Effects of fear and anger on perceived risks of terrorism: A national field experiment." *Psychological science* 14(2):144–150.

Litvak, Paul M, Jennifer S Lerner, Larissa Z Tiedens and Katherine Shonk. 2010. Fuel in the fire: How anger impacts judgment and decision-making. In *International handbook of anger*. Springer pp. 287–310.

Loewenstein, George and Jennifer S Lerner. 2003. "The role of affect in decision making." *Handbook of affective science* 619(642):3.

López, Rafael. 2001. "On p-beauty contest integer games." *UPF Economics and Business Working Paper* (608).

Machiavelli, Niccolò. 2009. *Discourses on Livy*. University of Chicago Press.

McMahan, Ian D. 1982. "Expectancy of success on sex-linked tasks." *Sex Roles* 8(9):949–958.

Milkman, Katherine, Modupe Akinola and Dolly Chugh. 2013. "Discrimination in the Academy: A Field Experiment." *SSRN, Working Paper* .

Möbius, Markus M, Muriel Niederle, Paul Niehaus and Tanya S Rosenblat. 2014. Managing self-confidence: Theory and experimental evidence. Technical report National Bureau of Economic Research.

Moore, Don A and Paul J Healy. 2008. "The trouble with overconfidence." *Psychological Review* 115(2):502.

Moore, Don A, Samuel A Swift, Zachariah S Sharek and Francesca Gino. 2010. "Correspondence bias in performance evaluation: Why grade inflation works." *Personality and Social Psychology Bulletin* 36(6):843–852.

Munyo, Ignacio and Martín A Rossi. 2013. "Frustration, euphoria, and violent crime." *Journal of Economic Behavior & Organization* 89:136–142.

Nagel, Rosemarie. 1995. "Unraveling in guessing games: An experimental study." *The American Economic Review* 85(5):1313–1326.

Niederle, Muriel and Lise Vesterlund. 2007. "Do women shy away from competition? Do men compete too much?" *The Quarterly Journal of Economics* 122(3):1067–1101.

Oprea, Ryan and Sevgi Yuksel. 2020. "Social Exchange of Motivated Beliefs." *Working Paper* .

Passarelli, Francesco and Guido Tabellini. 2017. "Emotions and political unrest." *Journal of Political Economy* 125(3):903–946.

Phelps, Edmund S. 1972. "The statistical theory of racism and sexism." *American Economic Review* 62(4):659–661.

Proto, Eugenio, Aldo Rustichini and Andis Sofianos. 2019. "Intelligence, personality, and gains from cooperation in repeated interactions." *Journal of Political Economy* 127(3):1351–1390.

Pyszczynski, Tom and Jeff Greenberg. 1987. Toward an integration of cognitive and motivational perspectives on social inference: A biased hypothesis-testing model. In *Advances in Experimental Social Psychology.* Vol. 20 Elsevier pp. 297–340.

Räty, Hannu, Johanna Vänskä, Kati Kasanen and Riitta Kärkkäinen. 2002. "Parents' explanations of their child's performance in mathematics and reading: A replication and extension of Yee and Eccles." *Sex Roles* 46(3-4):121–128.

Rey-Biel, Pedro, Roman Sheremeta and Neslihan Uler. 2018. *When income depends on performance and luck: The effects of culture and information on giving.* Emerald Publishing Limited pp. 167–203.

Ross, Lee. 1977. The intuitive psychologist and his shortcomings: Distortions in the attribution process. In *Advances in experimental social psychology.* Vol. 10 Elsevier pp. 173–220.

Rotemberg, Julio J. 2005. "Customer anger at price increases, changes in the frequency of price adjustment and monetary policy." *Journal of Monetary Economics* 52(4):829–852.

Rottenberg, J, RD Ray, JJ Gross, JA Coan and JJB Allen. 2007. "The handbook of emotion elicitation and assessment." *JJB Allen & JA Coan (Eds.)* pp. 9–28.

Rubin, Jared and Roman Sheremeta. 2016. "Principal-agent settings with random shocks." *Management Science* 62(4):985–999.

Saccardo, Silvia, Aniela Pietrasz and Uri Gneezy. 2018. "On the size of the gender difference in competitiveness." *Management Science* 64(4):1541–1554.

Sánchez-Pagés, Santiago and Marc Vorsatz. 2009. "Enjoy the silence: An experiment on truth-telling." *Experimental Economics* 12(2):220–241.

Sarsons, Heather. 2017. "Recognition for group work: Gender differences in academia." *American Economic Review* 107(5):141–45.

Sarsons, Heather. 2019. "Interpreting signals in the labor market: evidence from medical referrals." *Working Paper* .

Schamborg, Sara, Ruth J Tully and Kevin D Browne. 2016. "The use of the State–Trait Anger Expression Inventory–II with forensic populations: A psychometric critique." *International Journal of Offender Therapy and Comparative Criminology* 60(11):1239–1256.

Schotter, Andrew and Isabel Trevino. 2014. "Belief elicitation in the laboratory." *Annual Review of Economics* 6(1):103–128.

Schwardmann, Peter and Joël Van der Weele. 2019. "Deception and self-deception." *Nature Human Behaviour* 3(10):1055–1061.

Selody, Karen. 2010. "Board independence and the gender pay gap for top executives." *Working paper* .

Selten, Reinhard. 1978. "The chain store paradox." *Theory and Decision* 9(2):127–159.

Serra-Garcia, Marta, Eric Van Damme and Jan Potters. 2011. "Hiding an inconvenient truth: Lies and vagueness." *Games and Economic Behavior* 73(1):244–261.

Siedlecka, Ewa and Thomas F Denson. 2019. "Experimental methods for inducing basic emotions: A qualitative review." *Emotion Review* 11(1):87–97.

Small, Deborah A and Jennifer S Lerner. 2008. "Emotional policy: Personal sadness and anger shape judgments about a welfare case." *Political Psychology* 29(2):149–168.

Solda, Alice, Changxia Ke, Lionel Page and William Von Hippel. 2019. "Strategically delusional." *Experimental Economics* pp. 1–28.

Stahl, Dale O. 1996. "Boundedly rational rule learning in a guessing game." *Games and Economic Behavior* 16(2):303–330.

Sternberg, Robert J., Elena L. Grigorenko and Donald A. Bundy. 2001. "The Predictive Value of IQ." *Merrill-Palmer Quarterly* 47(1):1–41.

Stipek, Deborah J and J Heidi Gralinski. 1991. "Gender differences in children's achievement-related beliefs and emotional responses to success and failure in mathematics." *Journal of Educational Psychology* 83(3):361.

Strack, Fritz, Norbert Schwarz and Elisabeth Gschneidinger. 1985. "Happiness and reminiscing: The role of time perspective, affect, and mode of thinking." *Journal of Personality and Social Psychology* 49(6):1460.

Tice, Dianne M, Ellen Bratslavsky and Roy F Baumeister. 2001. "Emotional distress regulation takes precedence over impulse control: If you feel bad, do it!" *Journal of Personality and Social Psychology* 80(1):53.

Tiedens, Larissa Z and Susan Linton. 2001. "Judgment under emotional certainty and uncertainty: The effects of specific emotions on information processing." *Journal of Personality and Social Psychology* 81(6):973.

Van Leeuwen, Boris, Charles N Noussair, Theo Offerman, Sigrid Suetens, Matthijs Van Veelen and Jeroen Van De Ven. 2017. "Predictably Angry Facial Cues Provide a Credible Signal of Destructive Behavior." *Management Science* .

Vohs, Kathleen D, Roy F Baumeister and George Loewenstein. 2007. *Do Emotions Help or Hurt Decisionmaking?: A Hedgefoxian Perspective.* Russell Sage Foundation.

Vonk, Roos. 2002. "Self-serving interpretations of flattery: Why ingratiation works." *Journal of Personality and Social Psychology* 82(4):515.

Wang, Joseph Tao-yi, Michael Spezio and Colin F Camerer. 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.

Watson, David and Lee Anna Clark. 1999. "The PANAS-X: Manual for the positive and negative affect schedule-expanded form.".

Watson, David, Lee Anna Clark and Auke Tellegen. 1988. "Development and validation of brief measures of positive and negative affect: the PANAS scales." *Journal of Personality and Social Psychology* 54(6):1063.

Westermann, Rainer, Kordelia Spies, Günter Stahl and Friedrich W Hesse. 1996. "Relative effectiveness and validity of mood induction procedures: A meta-analysis." *European Journal of Social Psychology* 26(4):557–580.

Winter, Eyal. 2014. *Feeling smart: Why our emotions are more rational than we think*. PublicAffairs.

Winter, Eyal, Luciano Méndez-Naya and Ignacio García-Jurado. 2016. "Mental equilibrium and strategic emotions." *Management Science* 63(5):1302–1317.

Xiao, Erte and Daniel Houser. 2005. "Emotion expression in human punishment behavior." *Proceedings of the National Academy of Sciences* 102(20):7398–7401.

Yee, Doris K and Jacquelynne S Eccles. 1988. "Parent perceptions and attributions for children's math achievement." *Sex Roles* 19(5-6):317–333.

Zimmermann, Florian. 2020. "The Dynamics of Motivated Beliefs." *American Economic Review* 110(2):337–361.

Appendix A

A.1 Ranking Determination in the Non-Ego-relevant Conditions

In the non-ego-relevant treatments, the receivers' rankings were determined by drawing a random number from a distribution. I explained to subjects the specific distribution from which the number would be drawn. In a between-subject design I varied these distributions. There were three types of distributions as shown in Table A.1. Distributions I and II are positively and negatively skewed, respectively. In Distribution I rankings closer to the top are more likely to be drawn, whereas in Distribution II rankings closer to the bottom are more likely to be drawn. Distribution III is a uniform distribution where each rank is drawn will equal probability.

The different distributions were implemented to create exogenous variation in prior beliefs in the non-ego-relevant treatments.[1] Thus, these prior beliefs would be closer to those in the ego-relevant treatments.

[1] Experimental results indeed show that prior beliefs significantly differ by the allocated distribution in the non-ego-relevant treatments. It is assessed by running a regression of mean prior rank belief on the distribution type with robust standard errors. The reference category is Distribution I.

Table A.1: Distributions for the Determination of Not Ego-Relevant Rank

	Distribution I	Distribution II	Distribution III
Rank 1	10%	3%	10%
Rank 2	15%	4%	10%
Rank 3	20%	5%	10%
Rank 4	20%	8%	10%
Rank 5	15%	15%	10%
Rank 6	10%	15%	10%
Rank 7	4%	15%	10%
Rank 8	3%	15%	10%
Rank 9	2%	10%	10%
Rank 10	1%	10%	10%

Notes: Distribution I is positively skewed while Distribution II is negatively skewed. Distribution III is uniform.

A.2 Descriptive Statistics

Table A.2: Descriptive statistics

	Ego-relevant		Not Ego-relevant	
	Positive Payoff	Negative Payoff	Positive Payoff	Negative Payoff
Age (Mean)	21.062	21.851	21.135	20.442
	(2.384)	(3.417)	(2.124)	(1.841)
Female (Share)	0.605	0.493	0.635	0.596
	(0.492)	(0.504)	(0.486)	(0.495)
Student Status (Share)	0.901	0.761	0.923	0.962
	(0.300)	(0.430)	(0.269)	(0.194)
Risk Preferences (Mean)	6.407	6.746	6.577	6.519
	(1.672)	(1.735)	(1.742)	(2.072)
IQ score (Mean)	9.173	9.104	8.942	9.038
	(2.982)	(3.456)	(3.058)	(2.800)
N	81	67	52	52

Notes: the table shows descriptive statistics of the subjects who played in the role of receivers. Standard deviations are in parentheses.

A.3 Further Analyses

Here, I conduct alternative analyses to study treatment differences on actions in the game and posterior beliefs by taking into account *messages* and not *news* itself.

A.3.1 Actions in the Game

To perform these analyses, separately by positive and negative conditions, I run the following econometric specification:

$$\begin{aligned} Action_i = \beta_0 + \beta_1 message_i + \beta_2 treatment_i + \beta_3 rank\, prior_i \\ + \beta_4 std.\, prior_i + x_i'\beta_5 + \epsilon_i \end{aligned} \tag{A.1}$$

i is the receiver. $Action_i$ is the dependent variable and corresponds to receiver's i action in the game. $message_i$ is the message receiver's i received in the game. The $treatment_i$ variable is the treatment to which the receiver was randomly allocated. Then I use the following variables as controls. $rank\, prior_i$ is his mean rank prior belief and $std.\, prior_i$ is the standard deviation of his prior belief distribution. x_i is a vector of receiver's demographic variables (age, gender, risk preferences, and student status). I report robust standard errors in all specifications.

β_1 captures the effect of the message on receiver's action, while β_2 captures differences in actions across the two treatment. β_3 and β_4 capture the effects of the mean rank belief and the standard deviation of the prior belief distribution, respectively, on actions. β_5, is a vector of coefficients that capture the association between demographic variables and actions.

My main coefficient of interest is β_2. Following the conjectures of the paper, I expect β_2 to be negative. In other words, the effect of the ego-relevant treatments will reduce actions in the game.

The results of the estimation of Equation (A.1) are shown in Table A.3. In columns (1), (2) and (3), I report the results for the positive treatments, whereas in columns (4), (5), and (6) I report those for the negative treatments. As expected, messages have a strong effect on actions. A one-point higher message increases the action by 0.540 (p-value<0.000) in the positive condition and by 0.355 (p-value<0.000) in the negative condition. The treatment variable increases actions in the game in the positive condition and reduces them in the negative condition. This provides only limited support to the experimental hypotheses.[2]

A.3.2 Posterior Beliefs

I now conduct the same analysis as in Equation (A.1) but by using the posterior beliefs as dependent variable. That is:

$$\begin{aligned} Rank\, post_i = \beta_0 + \beta_1 message_i + \beta_2 treatment_i + \beta_3 rank\, prior_i \\ + \beta_4 std.\, prior_i + x_i'\beta_5 + \epsilon_i \end{aligned} \tag{A.2}$$

The coefficient of interest is again β_2. In Table A.4, I report the results. We can see that the estimated β_2 coefficient is almost zero (-0.028) in the positive treatment, while positive (0.262) in the negative condition, although they are not statistically different from zero.[3] In sum, there is little evidence that people process information self-servingly in the ego-relevant treatments.

I conduct further subsamples analyses on Equation (A.1) and Equation (A.2). First, I run it only for those individuals who received messages equal or lower to being in rank five (equal or higher to rank six) in the positive (negative) condition. Second, I run it only for those individuals that received a message lower (higher) than their mean prior belief in the positive (negative) conditions. I still find no

[2] Also, the effects are not statistically significant: p-values are equal to 0.322 in the positive and to 0.273 in the negative conditions.

[3] p-values equal to 0.883 and 0.180.

Table A.3: Regression Results for Actions by Ego-relevance of the State

	Positive Condition			Negative Condition		
	(1)	(2)	(3)	(4)	(5)	(6)
	Action	Action	Action	Action	Action	Action
message	0.548***	0.534***	0.540***	0.418***	0.366***	0.355***
	(0.076)	(0.076)	(0.077)	(0.076)	(0.071)	(0.073)
treatment	0.536*	0.358	0.345	-0.263	-0.417	-0.402
	(0.287)	(0.353)	(0.347)	(0.357)	(0.374)	(0.365)
rank prior		0.364***	0.374***		0.715***	0.781***
		(0.118)	(0.118)		(0.113)	(0.119)
std. prior		-0.238	-0.238		-0.149	-0.245
		(0.272)	(0.281)		(0.285)	(0.304)
constant	1.972***	0.705	0.827	2.906***	0.012	-0.342
	(0.382)	(0.931)	(1.772)	(0.500)	(0.857)	(1.444)
Demographics			✓			✓
R2	0.354	0.416	0.425	0.197	0.413	0.436
N	133	133	133	119	119	119

Notes: The table shows regression results of Equation (A.1). The regressions are estimated separately by positive and negative conditions. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.4: Regression Results for Posterior Beliefs by Ego-relevance of the State

	Positive Condition			Negative Condition		
	(1)	(2)	(3)	(4)	(5)	(6)
	Posterior	Posterior	Posterior	Posterior	Posterior	Posterior
message	0.307***	0.276***	0.272***	0.166**	0.098***	0.094**
	(0.063)	(0.054)	(0.056)	(0.064)	(0.037)	(0.038)
treatment	-0.040	-0.063	-0.028	0.254	0.270	0.262
	(0.224)	(0.190)	(0.187)	(0.269)	(0.192)	(0.194)
rank prior		0.676***	0.683***		0.802***	0.827***
		(0.068)	(0.068)		(0.055)	(0.060)
std. prior		0.084	0.109		0.237	0.198
		(0.161)	(0.157)		(0.153)	(0.156)
constant	3.818***	0.227	0.824	4.004***	-0.120	-0.513
	(0.311)	(0.520)	(1.006)	(0.411)	(0.388)	(0.609)
Demographics			✓			✓
R2	0.201	0.598	0.617	0.067	0.657	0.667
N	133	133	133	119	119	119

Notes: The table shows regression results of Equation (A.2). The regressions are estimated separately by positive and negative conditions. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

effect of the ego-relevant treatment on actions and posterior beliefs.[4]

[4]Results available upon request.

A.4 Experimental Instructions

Experimental Instructions translated from Spanish. The instructions are those for the positive ego-relevant treatment.

A.4.1 Welcome Page

Welcome to our experiment and thanks for your participation. The experiment will last about 30 minutes. During the experiment, you might have to wait while other participants reach the same part of the experiment that you are at. Please be patient.

This is a research project that aims at understanding how individuals make economic decisions. The study is conducted by [name researcher here] from the [name university research here] (email: [email researcher here]).

You will be paid S/. 5.00 for your participation. Moreover, you can earn extra money in the different parts of the experiment. In these parts of the experiment, you can earn up to S/. 20.00. Nevertheless, at the end of the experiment, the computer program will randomly (and with equal probability) select one of these parts. The chosen part will determine your extra payments in the experiment. In sum, you will definitely earn S/ 5.00 plus what you earn in one of the different parts of the experiment. It is in your best interest to pay attention to each part because any of these may be selected to count for your payments.

Please, read carefully the instructions for each part of the experiment. We will ask you questions about some of these parts after you have read the instructions. All your decisions in the experiment are anonymous. This means that other participants in the experiment (and even us, the researchers) will not be able to relate your decisions/choices with your name and surname.

Please notice that the experiment does not entail any sort of deception, as all experiments conducted in economics. This implies that you will be given truthful information regarding the instructions of the experiment and the experimental tasks. This research project has been granted ethical approval from the University of Warwick.

If you have questions during the experiment, you can contact us via whatsapp (laboratory phone number here) or via email (laboratory email here).

A.4.2 Instructions IQ test

In this part of the experiment, you will have to complete an IQ test. In particular, this test measures fluid intelligence. Many studies have found that there is a positive correlation between academic performance, performance at work, professional career progression, and even good health with IQ scores.

In this test you will have to solve 20 matrices. In particular, you will be shown the 20 matrices in two pages (10 matrices per page). For each page, you will have 5 minutes to solve the matrices and you can solve them in any order. You will be shown a clock on top of each page that shows the time that you have left.

Each matrix consists of a geometric composition of 9 pictures. One these pictures (the one at the bottom-right) is missing. You have to find the picture that is missing from one of 8 options that you will see below the matrix. Here you can see one example of a matrix (the answer to this question is option number 4).

[Raven matrix picture here]

If this part is selected to count for your extra payments in the experiment, the computer program will randomly select 3 answers. For each correct answer, you will earn S/. 5.00. This means that you can earn up to S/. 15.00.

Comprehension Questions

1. What type of test will you have to complete? Options: a) personality test; b) career test; c) aptitude test; d) intelligence test.

2. In total, how many matrices will you be asked to solve? Options: a) 5 matrices; a) 10 matrices; c) 15 matrices; d) 20 matrices.

3. In total, how much time will you have to solve the test? Options: a) 3 minutes; b) 7 minutes; c) 10 minutes; d) 15 minutes.

4. How much money would you earn if you answer correctly 2 out of the 3 randomly chosen questions that determine your payments in this part of the experiment? Options: a) S/ 0.00; b) S/ 5.00; c) S/ 10.00; d) S/ 20.00.

A.4.3 Instructions Prior Beliefs (Receiver)

For this second part of the experiment, you have been grouped with 9 other participants (randomly chosen) who have previously taken part in this experiment and have completed the same IQ test that you have just completed. From the group of 10 people, the computer has constructed a ranking based on the IQ scores. The person who scored the highest is in the first place, the one who obtained the second highest score is the second, and so on. If two or more people got the same score, then the computer will randomly determine who of these ranks higher.

We want to know what you believe your rank is. In particular, we want to know what is the probability with which you think you occupy each of these ten possible ranks. In other words, we will ask you: "What is the probability with which you believe you occupy the first place (ranking equal to one)"; "What is the probability with which you believe you occupy the second place (ranking equal to two)"; ...; "What is the probability with which you believe you occupy the tenth place (ranking equal to ten)".

You can only insert whole numbers (0, 1, ..., 100). Please realize that the sum of these 10 probabilities has to be 100%.

Your earnings in this part of the experiment

If this second part of the experiment is randomly chosen to count for your payments in the experiment, this is how the payments will be determined.

There are two possible prizes: S/20.00 and S/. 0.00. The closer your answers will be to your true rank, the higher the probability for you to earn S/ 20.00. The payment method is such that it incentivizes you to answer what you really believe your rank is. For more information about the elicitation mechanism, please click on the button "Payment method" that you can find at the bottom of this page.

Although the method looks complicated, its implications are very simple. For you to maximize the chances of earning the S/. 20.00, your answers to these questions should correspond to what you really believe the probability of occupying each rank is.

A.4.4 Instructions Sender-Receiver Game (Receiver)

Now you will play a game that consists of two players: Player A and Player B. You have been randomly assigned to be Player B.

In the game, Player A will be informed about your rank in the IQ test you previously completed. Remember that your rank is based on your performance relative to 9 other people. Player A is not one of these 9 other people.

After this, Player A will send you a message. The message will read: "You are placed XXX in the ranking". Player A will complete the sentence with one of the 10 possible ranks. Realize that the message she sends might or might not be your actual rank in the IQ test.

After you read the message Player A sent you, you will have to take an action in the game about your rank in the test. The figure below shows the game dynamics.

[Sender-Receiver game picture here]

The earnings in the game will depend on both your own true rank in the intelligence test and your action in the game. The following figure shows the different combinations of payoffs in the game for both players.

[Payoff table here]

The first column of the table indicates the true rank of player B in the IQ test, while the first row indicates the action that Player B takes. In the table, the numbers show the earnings of each player in the game: the numbers in red show Player A's payoff, while the numbers in blue show the earnings of Player B.

From the table, you can realize that Player's A earns more money as Player B takes actions that correspond to him being in higher ranks. For example, Player's A highest earnings, S/. 15, are when Player B plays an action that corresponds to him being of rank 1. On the other hand, Player B maximizes his earnings in the game when his action in the game corresponds to his own true rank.

Example: Player A is informed that Player's B rank in the IQ test is five (fifth in the ranking). Player A then sends a message to Player B. Player B, once he reads the message, he plays the action in the game that corresponds to him being in the second place. The cell that determines the earnings in the game is then the following: row "Ranking B=5" (player's B true rank), column "B=2" (player's B action in the game). Player A earns S/. 14.00, while Player B earns S/. 12.00.

If this third part of the experiment is randomly chosen, then your payments in the experiment will be determined by your earnings in this game.

Comprehension Questions

1. What determines the ranking upon which the game is based? Options: a) Player's A ranking that is determined by her performance in an IQ test; b) Player's A ranking that is randomly determined; c) Player's B ranking that is determined by his performance in an IQ test; d) Player's B ranking that is randomly determined.

2. Who sends the message in the game and what it consists of? Options: a) Player A sends a message about Player's A rank; b) Player A sends a message about Player's B rank; c) Player B sends a message about Player's A rank; d) Player B sends a message about Player's B rank.

3. How much money would Player A earn if Player B plays action "six" (in the table, player's A payoffs are shown in red). Options: a) S/. 3.00; b) S/. 8.00; c) S/. 13.00; d) S/. 15.00.

4. If Player's B ranking in the test is "fourth" and he plays action "six", how much money would Player B earn? (in the table, player's B payoffs are shown in blue). Options: a) S/. 3.00; b) S/. 8.00; c) S/. 13.00; d) S/. 15.00.

A.4.5 Instructions Posterior Beliefs (Receiver)

Again, we will ask you what you really believe your rank in the IQ test is. You can change the probabilities that you previously wrote as you wish. You may take or may not take into account the game you just played and the message that Player A sent you.

If this fourth part of the experiment is randomly chosen to count for payments, we will use the same method as before. There are two possible prizes: S/20.00 and S/. 0.00. The closer your answers will be to your true rank, the higher the probability for you to earn S/ 20.00. The payment method is such that it incentivizes you to answer what you really believe your rank is.

You can only insert whole numbers (0, 1, ..., 100). Please realize that the sum of these 10 probabilities has to be 100%.

Appendix B

B.1 Summary Statistics

Table B.1: Summary statistics of our sample (1)

	Principals	Agents
Female	74%	55%
	(0.45)	(0.50)
Age	21.64	22.02
	(1.10)	(1.18)
Degree of study		
Commerce	0.36	0.36
	(0.49)	(0.48)
Economics	0.62	0.62
	(0.49)	(0.49)
Geography	0.00	0.00
	(0.00)	(0.00)
Sociology	0.02	0.02
	(0.02)	(0.15)
Other/Prefer not say	0.00	0.07
	(0.00)	(0.26)
Year of study		
1st year MA	0.07	0.07
	(0.26)	(0.26)
2nd year MA	0.83	0.86
	(0.38)	(0.35)
MPhil	0.00	0.00
	(0.00)	(0.00)
PhD	0.00	0.00
	(0.00)	(0.00)
Other/Prefer not say	0.10	0.07
	(0.30)	(0.26)
Language		
English	0.02	0.00
	(0.15)	(0.00)
N	42	42

Notes: Table shows descriptive statistics (in means) of the experimental dataset. Standard deviations are in parentheses. Female is the share of female participants. Age is the reported age of the participant. Degree of study: 1=Sociology, 2=Commerce, 3=Geography, 4=Economics, 5=Other. Year of study: 1=First year master degree, 2=Second year master degree, 3=Master of philosophy (mphil), 4=PhD, 5=Other. Language: 1=English, 2=Other.

Table B.2: Summary statistics of our sample (2)

	Principals	Agents
Religion		
Muslim	0.07	0.05
	(0.26)	(0.22)
Hindu	0.88	0.88
	(0.33)	(0.33)
Sikh	0.00	0.02
	(0.00)	(0.15)
Christian	0.02	0.05
	(0.15)	(0.22)
Buddhist	0.00	0.00
	(0.00)	(0.00)
Parsi	0.00	0.00
	(0.00)	(0.00)
Other/Prefer not say	0.02	0.00
	(0.15)	(0.00)
Caste		
Scheduled caste	0.07	0.12
	(0.26)	(0.33)
Scheduled tribe	0.00	0.00
	(0.00)	(0.00)
Other backward castes	0.26	0.285
	(0.45)	(0.46)
General	0.67	0.595
	(0.48)	(0.50)
Other/Prefer not say	0.00	0.00
	(0.00)	(0.00)
N	42	42

Notes: Table shows descriptive statistics (in means) of the experimental dataset. Standard deviations are in parentheses. Religion: 1=Muslim, 2=Hindu, 3=Sikh, 4=Christian, 5=Buddhist, 6=Parsi, 7=Other, 8=Prefer not say. Caste: 1=Scheduled caste, 2=Scheduled tribe, 3=Other backward castes, 4=General, 5=Other, 6=Prefer not say.

Table B.3: Summary statistics of variables in main econometric specification

	Mean	Standard Deviation
High Outcome	0.45	0.50
Female Agent	0.55	0.50
Female Agent × High Outcome	0.25	0.43
Female Principal	0.73	0.44
Female Principal × High Outcome	0.32	0.47
Same Gender	0.55	0.50
N	804	804

Notes: Table shows descriptive statistics of the corresponding variables.

B.2 Mean Payments to Each Party by Realized Outcome

In the paper we have analysed principals' payment decisions to their matched agents following low and high outcomes. Here, we now present a visual representation of average payments to each party following both low and high outcomes (see Figure B.1). This figure shows that, going from a low to a high outcome, principals' payments to their matched agent increase (from ₹135.15 to ₹243.35) whereas payments decrease to both the other randomly drawn agent (from ₹120.37 to ₹75.99) and to the experimenter (from ₹94.48 to ₹30.66).

Figure B.1: Mean payments to each party by realized outcome

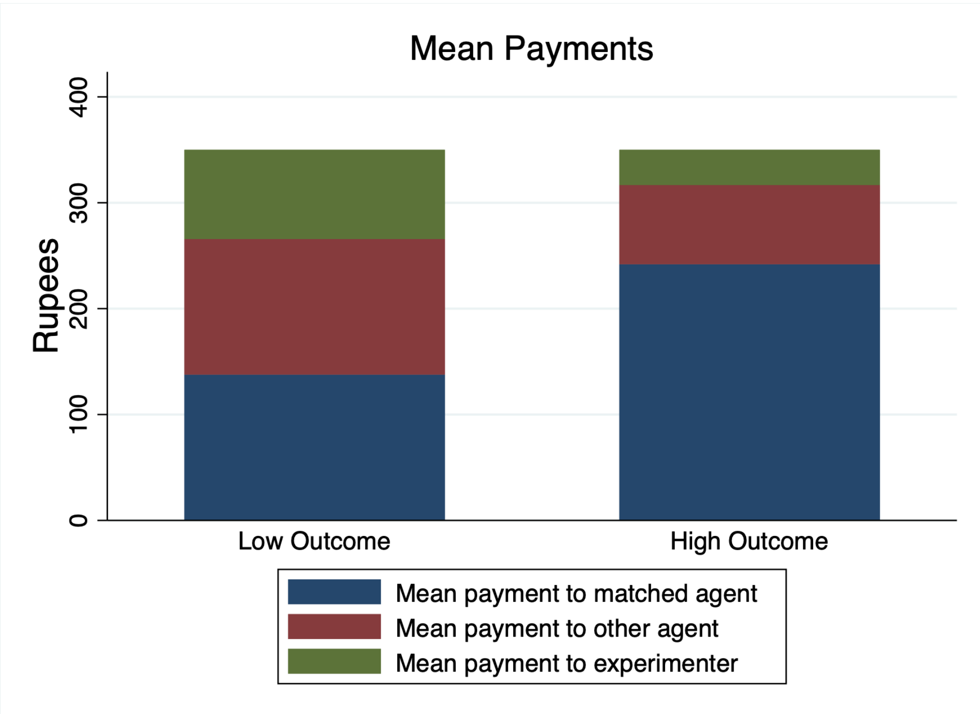

If we look at mean payments by taking into account the gender of the matched agent, we find very similar patterns (Figure B.2). Indeed, while agents (irrespective of their gender) are rewarded for high outcomes, this comes at the cost of lower payments to both the other randomly matched agent and the experimenter.

Figure B.2: Mean payments to each party by realized outcome and gender of the agent

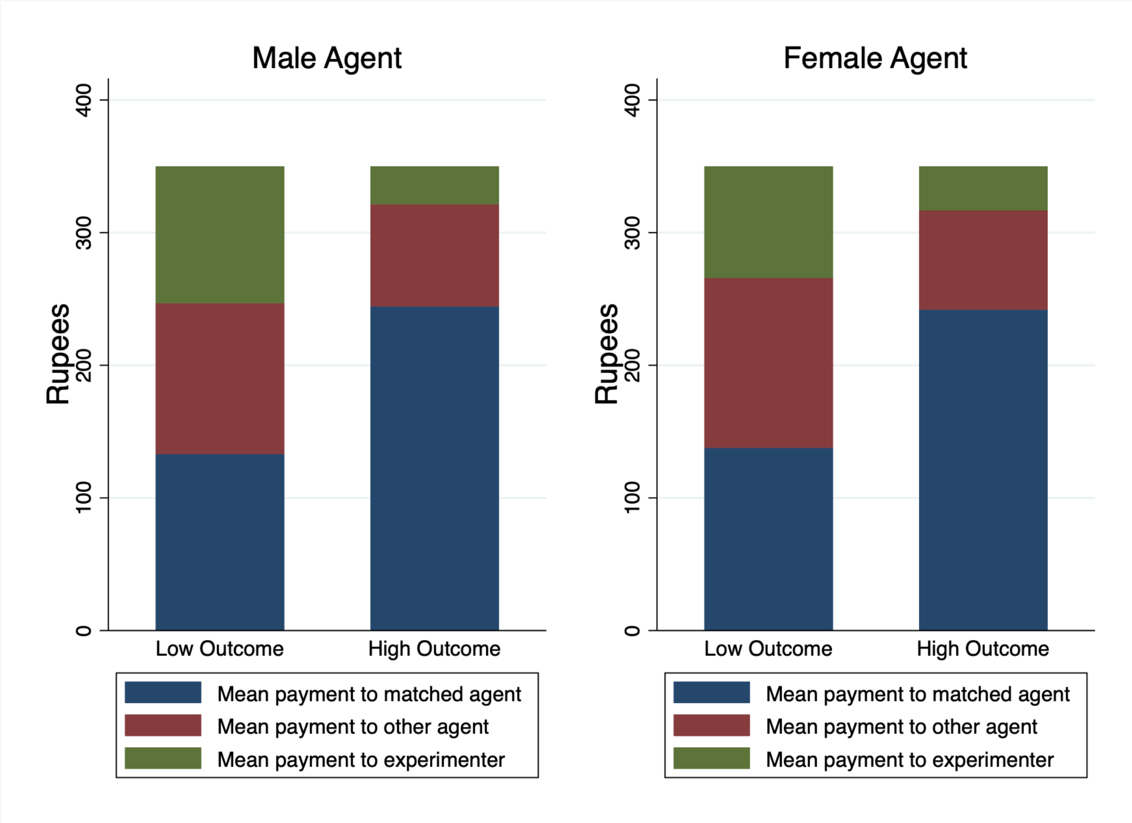

B.3 Robustness Checks for Principals' Payment Decisions

Table B.4: Regression results with principal fixed effects

	(1)	(2)
High Outcome	110.95***	0.22***
	(33.53)	(0.027)
Female Agent	9.19	0.02
	(9.06)	(0.02)
Female Agent × High Outcome	-6.78	-0.26
	(12.37)	(0.03)
Principal FE	✓	✓
Task Controls	✓	✓
R-Squared	0.62	0.48
N	804	804

Dependent variable in column 1 is principals payments and in 2 is principals beliefs. Task controls are dummy variables for each task. Standard errors are clustered at the principal level and principal fixed effects are included in both specifications.

Table B.5: Regression results for principals' payments with session fixed effects

	(1)	(2)	(3)	(4)
High Outcome	101.64***	102.27***	114.19***	112.27***
	(14.80)	(18.17)	(33.43)	(33.39)
Female Agent		3.27	2.74	8.61
		(9.08)	(8.69)	(9.30)
Female Agent × High Outcome		-1.09	-0.17	-0.96
		(13.31)	(12.43)	(12.50)
Female Principal			-34.96	-34.49
			(36.85)	(36.53)
Female Principal × High Outcome			-16.55	-14.53
			(35.06)	(34.95)
Same Gender				-11.79*
				(6.45)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.37	0.37	0.39	0.39
N	804	804	804	804

Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Task controls are dummy variables for each task. Session fixed effects are included in all specifications. Standard errors are clustered at the principal level.

Table B.6: Regression results for principals' payments with round fixed effects

	(1)	(2)	(3)	(4)
High Outcome	101.25***	99.94***	111.83***	109.31***
	(14.05)	(17.39)	(32.71)	(32.76)
Female Agent		0.92	0.31	8.64
		(9.60)	(9.40)	(9.67)
Female Agent × High Outcome		2.39	3.31	2.08
		(13.45)	(12.67)	(12.85)
Female Principal			-33.23	-32.75
			(41.41)	(40.53)
Female Principal × High Outcome			-16.17	-13.29
			(34.60)	(34.57)
Same Gender				-16.45**
				(6.86)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.34	0.34	0.35	0.35
N	804	804	804	804

Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Task controls are dummy variables for each task. Round fixed effects are included in all specifications. Standard errors are clustered at the principal level.

Table B.7: Regression results for principals' payments without controls

	(1)	(2)	(3)	(4)
High Outcome	108.21***	104.21***	116.58***	113.76***
	(14.72)	(18.76)	(33.33)	(33.61)
Female Agent		-4.81	-1.34	7.22
		(11.45)	(11.79)	(11.49)
Female Agent × High Outcome		7.33	3.70	2.35
		(16.98)	(16.16)	(16.28)
Female Principal			-27.71	-29.09
			(29.27)	(29.01)
Female Principal × High Outcome			-16.22	-12.93
			(36.14)	(36.23)
Same Gender				-17.18**
				(6.34)
R-Squared	0.25	0.25	0.27	0.28
N	804	804	804	804

No controls are added to the regressions. Standard errors are clustered at the principal level.

Table B.8: Regression results for principals' payments for the first ten rounds only

	(1)	(2)	(3)	(4)
High Outcome	100.94***	101.99***	117.64***	111.16***
	(15.54)	(18.85)	(38.08)	(39.91)
Female Agent		0.52	-0.23	12.34
		(15.60)	(16.29)	(20.80)
Female Agent × High Outcome		-1.93	-1.19	-1.13
		(20.34)	(20.77)	(21.11)
Female Principal			-22.24	-19.78
			(47.32)	(46.91)
Female Principal × High Outcome			-21.33	-15.34
			(41.91)	(44.04)
Same Gender				-24.09*
				(14.24)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.31	0.31	0.32	0.33
N	420	420	420	420

Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Task controls are dummy variables for each task. Results for only the initial 10 rounds are shown. Standard errors are clustered at the principal level.

Table B.9: Regression results for principals' payments after removing the first five rounds

	(1)	(2)	(3)	(4)
High Outcome	109.94***	117.91***	129.97***	128.10***
	(15.05)	(18.59)	(32.08)	(31.87)
Female Agent		7.11	5.49	15.25
		(11.03)	(10.72)	(10.36)
Female Agent × High Outcome		-14.39	-11.97	-13.21
		(14.31)	(13.42)	(13.51)
Female Principal			-36.68	-35.99
			(41.01)	(39.37)
Female Principal × High Outcome			-17.68	-15.89
			(34.19)	(33.81)
Same Gender				-20.11**
				(7.41)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.37	0.37	0.38	0.38
N	594	594	594	594

Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Task controls are dummy variables for each task. First five rounds were removed for regressions above. Standard errors are clustered at the principal level.

Table B.10: Regression results for alternative definition of dependent variable

	($\geq$ 50)	($\geq$100)	($\geq$150)	($\geq$200)	($\geq$250)	($\geq$300)
High Outcome	0.21***	0.32***	0.37***	0.38***	0.33***	0.21***
	(0.06)	(0.07)	(0.07)	(0.07)	(0.07)	(0.06)
Female Agent	-0.01	0.02	0.00	-0.01	0.04	0.01
	(0.03)	(0.03)	(0.04)	(0.04)	(0.04)	(0.02)
Female Agent × High Outcome	0.01	-0.03	0.03	0.05	-0.01	-0.04
	(0.04)	(0.04)	(0.06)	(0.07)	(0.06)	(0.04)
Mean	0.81	0.72	0.62	0.36	0.24	0.12
Demographics	✓	✓	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓	✓	✓
R-Squared	0.24	0.25	0.26	0.27	0.26	0.27
N	804	804	804	804	804	804

Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Task controls are dummy variables for each task. Robust standard errors are reported.

Table B.11: Regression results for principals' payments for different tasks

	(Maths)	(Ravens)	(Memory)	(Effort)
High Outcome	97.21***	101.38***	72.94**	160.58***
	(23.15)	(26.68)	(30.53)	(20.42)
Female Agent	-0.79	2.72	-13.01	16.10
	(21.43)	(13.08)	(15.50)	(17.57)
Female Agent × High Outcome	0.24	7.39	46.14*	-15.45
	(25.83)	(23.88)	(25.17)	(14.78)
Demographics	✓	✓	✓	✓
R-Squared	0.29	0.29	0.55	0.85
N	312	252	150	90

Each column depicts results for a regression of principal payments for the particular task mentioned in the heading. Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Standard errors are clustered at the principal level.

B.4 Robustness Checks for Principals' Beliefs

Table B.12: Regression results for principals' beliefs with session fixed effects

	(1)	(2)	(3)	(4)
High Outcome	0.21***	0.22***	0.14***	0.14***
	(0.02)	(0.03)	(0.03)	(0.03)
Female Agent		0.02	0.03	0.04**
		(0.02)	(0.02)	(0.02)
Female Agent × High Outcome		-0.02	-0.02	-0.03
		(0.03)	(0.03)	(0.03)
Female Principal			-0.14**	-0.14**
			(0.05)	(0.05)
Female Principal × High Outcome			0.11**	0.11**
			(0.04)	(0.04)
Same Gender				-0.03
				(0.02)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.33	0.34	0.34	0.35
N	804	804	804	804

Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Task controls are dummy variables for each task. Session fixed effects are included in all specifications. Standard errors are clustered at the principal level.

Table B.13: Regression results for principals' beliefs with round fixed effects

	(1)	(2)	(3)	(4)
High Outcome	0.22***	0.23***	0.16***	0.16***
	(0.03)	(0.03)	(0.03)	(0.03)
Female Agent		0.02	0.03	0.04*
		(0.03)	(0.03)	(0.02)
Female Agent × High Outcome		-0.02	-0.03	-0.03
		(0.03)	(0.03)	(0.03)
Female Principal			-0.14**	-0.14**
			(0.05)	(0.05)
Female Principal × High Outcome			0.10**	0.10**
			(0.05)	(0.05)
Same Gender				-0.02
				(0.02)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.33	0.34	0.35	0.35
N	804	804	804	804

Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Task controls are dummy variables for each task. Session fixed effects are included in all specifications. Standard errors are clustered at the principal level.

Table B.14: Regression results for principals' beliefs without controls

	(1)	(2)	(3)	(4)
High Outcome	0.17***	0.18***	0.12***	0.12***
	(0.03)	(0.04)	(0.04)	(0.04)
Female Agent		0.02	0.02	0.03
		(0.02)	(0.03)	(0.03)
Female Agent × High Outcome		-0.02	-0.02	-0.02
		(0.04)	(0.04)	(0.04)
Female Principal			-0.01	-0.01
			(0.06)	(0.06)
Female Principal × High Outcome			0.08*	0.08*
			(0.05)	(0.05)
Same Gender				-0.02
				(0.02)
R-Squared	0.25	0.25	0.25	0.25
N	804	804	804	804

No controls are included in the regressions. Standard errors are clustered at the principal level.

Table B.15: Regression results for principals' beliefs for the first ten rounds only

	(1)	(2)	(3)	(4)
High Outcome	0.23***	0.25***	0.21***	0.20***
	(0.04)	(0.04)	(0.04)	(0.04)
Female Agent		0.04	0.04	0.06
		(0.03)	(0.03)	(0.03)
Female Agent × High Outcome		-0.03	-0.04	-0.04
		(0.05)	(0.05)	(0.05)
Female Principal			-0.15*	-0.15**
			(0.07)	(0.07)
Female Principal × High Outcome			0.06	0.07
			(0.06)	(0.06)
Same Gender				-0.03
				(0.02)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.31	0.31	0.33	0.33
N	420	420	420	420

Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Task controls are dummy variables for each task. Results for only the initial 10 rounds are shown. Standard errors are clustered at the principal level.

Table B.16: Regression results for principals' beliefs after removing the first five rounds

	(1)	(2)	(3)	(4)
High Outcome	0.23***	0.22***	0.15***	0.15***
	(0.02)	(0.03)	(0.03)	(0.03)
Female Agent		0.01	0.02	0.02
		(0.02)	(0.02)	(0.02)
Female Agent × High Outcome		0.00	-0.00	-0.00
		(0.03)	(0.03)	(0.03)
Female Principal			-0.14**	-0.14**
			(0.05)	(0.05)
Female Principal × High Outcome			0.11	0.11***
			(0.04)	(0.04)
Same Gender				-0.02
				(0.02)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.37	0.36	0.38	0.38
N	594	594	594	594

Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Task controls are dummy variables for each task. Results above are shown after removing first five rounds of the sessions. Standard errors are clustered at the principal level.

Table B.17: Regression results for principals' beliefs for different tasks

	(Maths)	(Ravens)	(Memory)	(Effort)
High Outcome	0.17***	0.33***	0.16**	0.23***
	(0.05)	(0.06)	(0.04)	(0.05)
Female Agent	-0.00	0.05	0.04	-0.05
	(0.04)	(0.04)	(0.03)	(0.06)
Female Agent × High Outcome	0.01	0.00	-0.04	0.03
	(0.05)	(0.06)	(0.06)	(0.07)
Demographics	✓	✓	✓	✓
R-Squared	0.22	0.37	0.42	0.63
N	312	252	150	90

Each column depicts results for a regression of principal beliefs for the particular task mentioned in the heading. Demographic variables include: principal's age, religion, caste, main language, state, education level, and field of study. Standard errors are clustered at the principal level.

B.5 Principals' Payment Decisions and Beliefs by Agents' Age

Table B.18: Regression results for principal's payments

	(1)	(2)	(3)	(4)
High Outcome	102.18***	101.18***	96.45***	96.31***
	(14.35)	(17.42)	(20.32)	(20.31)
Age Agent		-4.32	-4.94	-4.45
		(9.97)	(9.52)	(9.36)
Age Agent × High Outcome		1.51	-1.06	-1.00
		(12.98)	(12.90)	(12.96)
Age Principal			21.42	21.53
			(26.17)	(26.18)
Age Principal × High Outcome			14.61	14.65
			(29.87)	(29.91)
Same Age				1.68
				(8.56)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.34	0.34	0.35	0.35
N	804	804	804	804

Demographic variables include: principal's gender, religion, caste, main language, state, education level, and field of study. Standard errors are clustered at the principal level. The agent's and principal's age variables are dummy variables. The "Age Agent" variable is equal to 1 if the agent's age is above or equal to the median agents' age and 0 otherwise. The "Age Principal" variable is equal to 1 if the principal's age is above to the median principals' age and 0 otherwise.

Table B.19: Regression results for principal's beliefs

	(1)	(2)	(3)	(4)
High Outcome	0.21***	0.21***	0.22***	0.22***
	(0.02)	(0.03)	(0.03)	(0.03)
Age Agent	-0.02	-0.02	-0.02	-0.02
		(0.02)	(0.03)	(0.03)
Age Agent × Outcome		0.00	0.00	0.00
		(0.03)	(0.03)	(0.03)
Age Principal			0.03	0.03
			(0.07)	(0.07)
Age Principal × Outcome			0.00	0.00
			(0.06)	(0.06)
Same Age				-0.01
				(0.02)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
R-Squared	0.30	0.31	0.31	0.31
N	804	804	804	804

Demographic variables include: principal's gender, religion, caste, main language, state, education level, and field of study. Standard errors are clustered at the principal level. The agent's and principal's age variables are dummy variables. The "Age Agent" variable is equal to 1 if the agent's age is above or equal to the median agents' age and 0 otherwise. The "Age Principal" variable is equal to 1 if the principal's age is above to the median principals' age and 0 otherwise.

Appendix C

C.1 Summary Statistics

Table C.1: Descriptive statistics – demographic variables

Experiment 1	Mean	Std. Dev.	Min	Max	N
Demographics					
Female (Share)	58.48%	0.494	0	1	171
Age (Mean)	21.205	3.352	18	43	171
Culture					
Anglo Cultures	0.263	0.442	0	1	171
Confucian Asia	0.123	0.329	0	0	171
Eastern Europe	0.082	0.275	0	1	171
Germanic Europe	0.023	0.152	0	1	171
Latin America	0.006	0.076	0	1	171
Latin Europe	0.047	0.212	0	1	171
Nordic Europe	0.012	0.108	0	1	171
Southern Asia	0.392	0.490	0	1	171
Other	0.053	0.224	0	1	171
Experiment 2	Mean	Std. Dev.	Min	Max	N
Demographics					
Female (Share)	52.78%	0.501	0	1	180
Age (Mean)	20.456	2.813	18	33	180
Culture					
Anglo Cultures	0.261	0.440	0	1	180
Confucian Asia	0.094	0.293	0	0	180
Eastern Europe	0.117	0.322	0	1	180
Germanic Europe	0.022	0.148	0	1	180
Latin America	0.011	0.105	0	1	180
Latin Europe	0.072	0.260	0	1	180
Nordic Europe	0.006	0.075	0	1	180
Southern Asia	0.406	0.492	0	1	180
Other	0.011	0.105	0	1	180

Table C.2: Descriptive statistics – education and other variables

Experiment 1	Mean	Std. Dev.	Min	Max	N
Education					
High School Final Mark (Normalized)	85.87%	0.112	0.3	1	169
Degree Quantitative (Share)	66.86%	0.472	0	1	169
Year of Study					
1st Year	0.456	0.500	0	1	171
Other	0.532	0.500	0	1	171
Not a Student	0.012	0.108	0	1	171
Other					
Risk Preferences (Mean)	5.647	1.943	0	10	171
Experience in the Game (Share)	4.09%	0.199	0	1	171
Experiment 2	Mean	Std. Dev.	Min	Max	N
Education					
High School Final Mark (Normalized)	86.26%	0.107	0.4	1	179
Degree Quantitative (Share)	72.22%	0.449	0	1	180
Year of Study					
1st Year	0.444	0.498	0	1	180
Other	0.533	0.500	0	1	180
Not a Student	0.022	0.148	0	1	171
Other					
Risk Preferences (Mean)	6.011	1.894	1	10	180
Experienced in the Game (Share)	3.33%	0.180	0	1	180

Table C.3: Descriptive statistics by condition – demographic variables

Experiment 1	Anger	Control	Difference	p-value
Demographics				
Female (Share)	63.64%	53.01%	10.62%	0.161
Age (Mean)	21.034	21.386	-0.351	0.495
Culture				
Anglo Cultures	0.284	0.241	0.043	0.525
Confucian Asia	0.136	0.108	0.028	0.581
Eastern Europe	0.080	0.084	-0.005	0.910
Germanic Europe	0.023	0.024	-0.001	0.953
Latin America	0.000	0.012	-0.012	0.304
Latin Europe	0.045	0.048	-0.003	0.933
Nordic Europe	0.000	0.024	-0.024	0.145
Southern Asia	0.398	0.386	0.012	0.871
Other	0.034	0.072	-0.038	0.266
Experiment 2	Sadness	Control	Difference	p-value
Demographics				
Female (Share)	54.12%	51.55%	2.54%	0.735
Age (Mean)	20.718	20.221	0.497	0.238
Culture				
Anglo Cultures	0.247	0.274	-0.027	0.687
Confucian Asia	0.153	0.042	0.111	0.011
Eastern Europe	0.082	0.147	-0.065	0.177
Germanic Europe	0.024	0.021	0.002	0.911
Latin America	0.012	0.011	0.001	0.937
Latin Europe	0.094	0.053	0.041	0.286
Nordic Europe	0.000	0.011	-0.011	0.346
Southern Asia	0.376	0.432	-0.055	0.455
Other	0.012	0.011	0.001	0.937

Table C.4: Descriptive statistics by condition – education and other variables

Experiment 1	Anger	Control	Difference	p-value
Education				
High School Final Mark (Normalized)	84.517%	87.306%	-2.879%	0.107
Degree Quantitative (Share)	69.318%	64.198%	5.121%	0.483
Year of Study				
1st Year	43.181%	48.182%	-0.050%	0.514
Other	56.181%	49.398%	7.421%	0.334
Not a Student	0	2.410%	-2.410%	0.145
Other				
Risk Preferences (Mean)	5.489	5.795	-0.307	0.304
Experienced in the Game (Share)	3.409%	4.819%	-0.141%	0.664
Experiment 2	Sadness	Control	Difference	p-value
Education				
High School Final Mark (Normalized)	86.024%	86.474%	-0.450%	0.779
Degree Quantitative (Share)	74.118%	70.526%	3.591%	0.594
Year of Study				
1st Year	43.529%	45.263%	-1.733%	0.817
Other	55.294%	51.579%	3.715%	0.620
Not a Student	1.176%	3.158%	-1.981%	0.371
Other				
Risk Preferences (Mean)	6.176	5.863	0.313	0.269
Experience in the Game (Share)	3.529%	3.158%	0.372%	0.891

C.2 Screenshots of the Experiment

C.2.1 Emotional Induction

Figure C.1: General Instructions

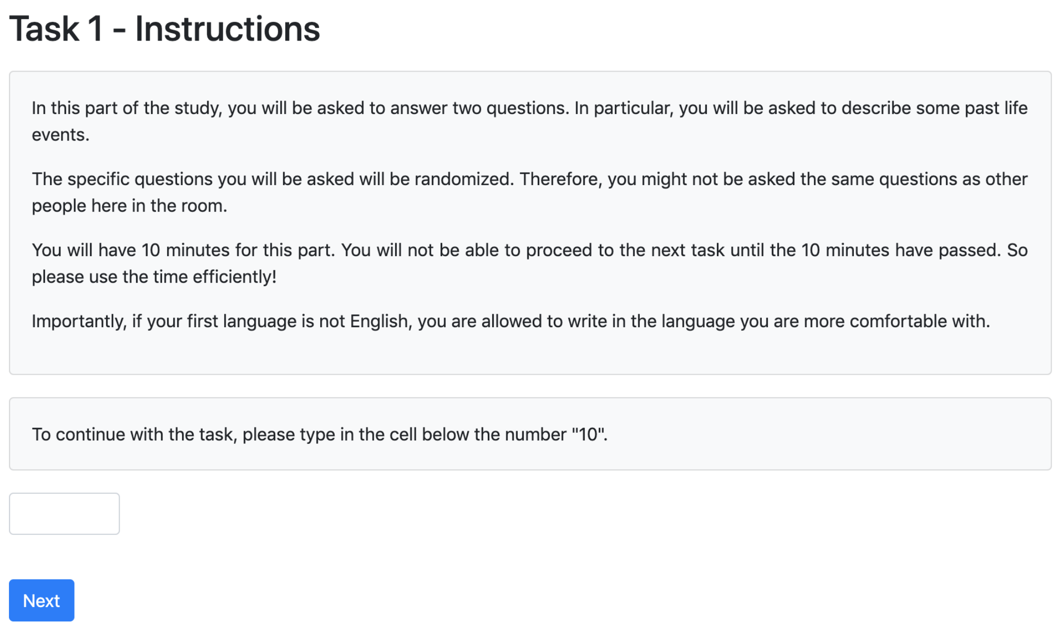

Anger Treatment

Figure C.2: Anger Induction – Question 1

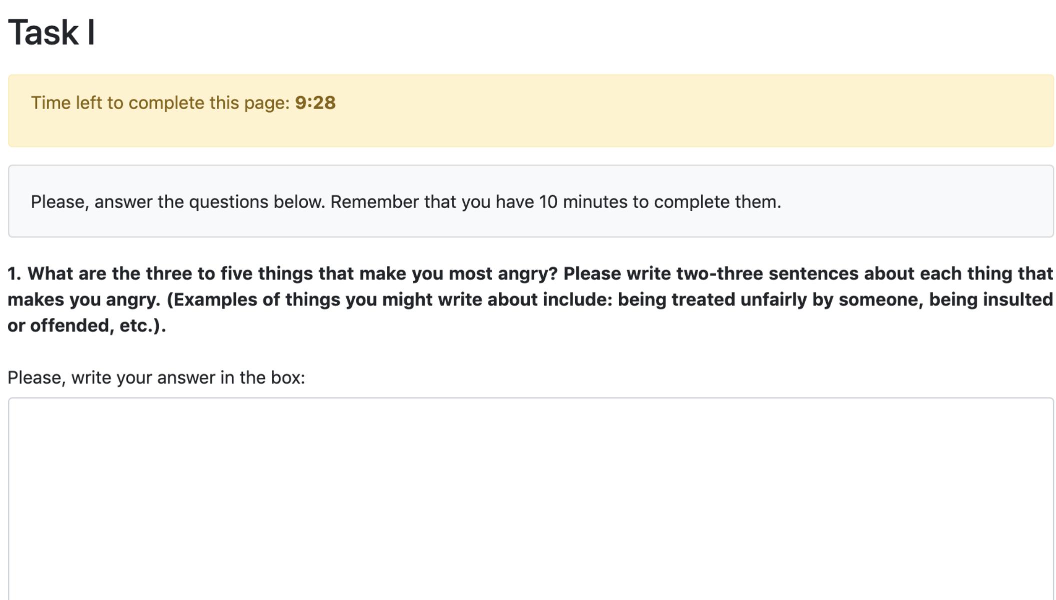

Figure C.3: Anger Induction – Question 2

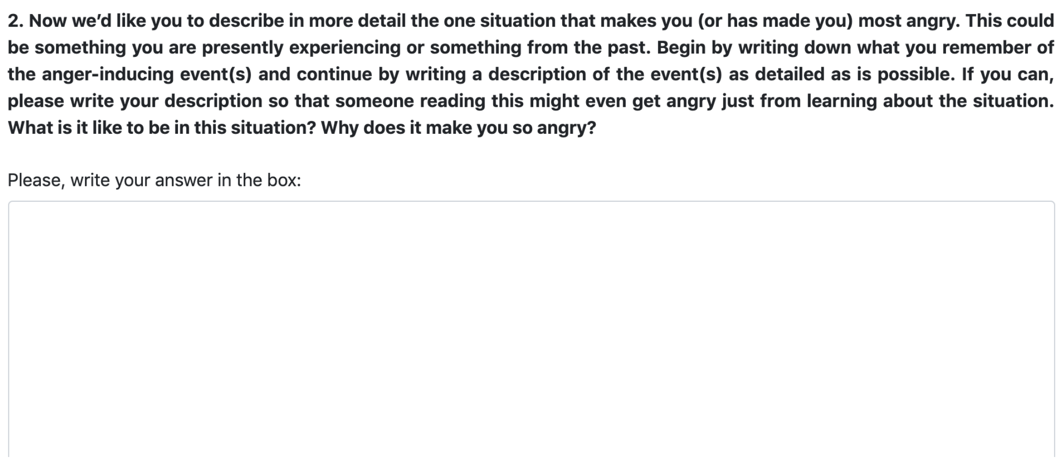
2. Now we'd like you to describe in more detail the one situation that makes you (or has made you) most angry. This could be something you are presently experiencing or something from the past. Begin by writing down what you remember of the anger-inducing event(s) and continue by writing a description of the event(s) as detailed as is possible. If you can, please write your description so that someone reading this might even get angry just from learning about the situation. What is it like to be in this situation? Why does it make you so angry?

Please, write your answer in the box:

Sadness Treatment

Figure C.4: Sadness Induction – Question 1

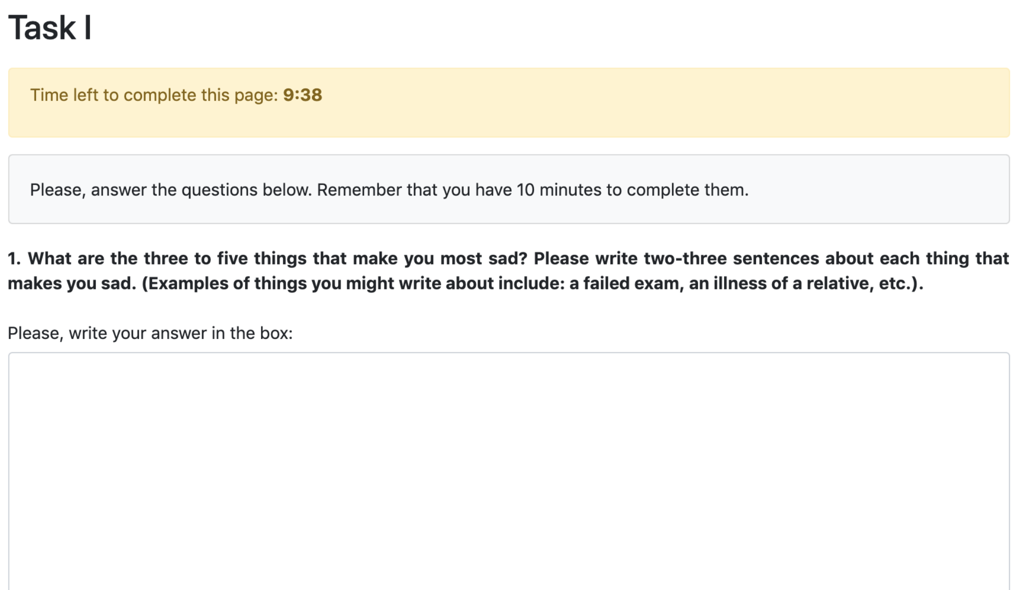
Task I

Time left to complete this page: **9:38**

Please, answer the questions below. Remember that you have 10 minutes to complete them.

1. What are the three to five things that make you most sad? Please write two-three sentences about each thing that makes you sad. (Examples of things you might write about include: a failed exam, an illness of a relative, etc.).

Please, write your answer in the box:

Figure C.5: Sadness Induction – Question 2

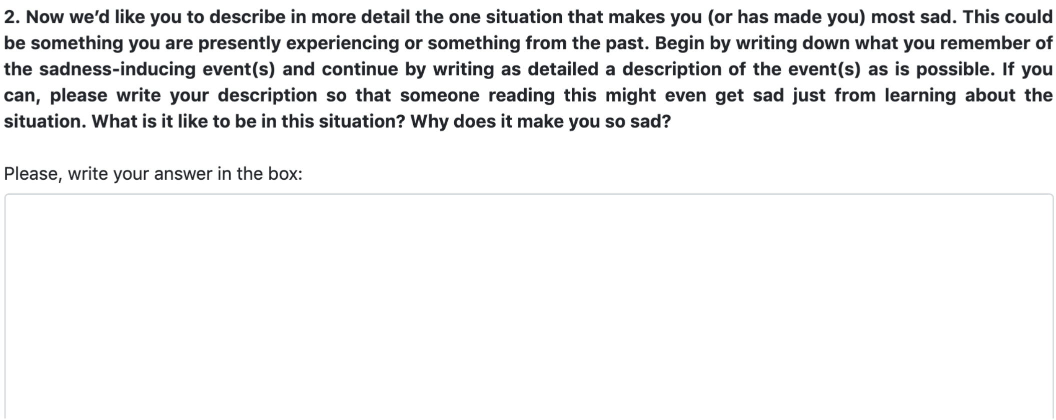
2. Now we'd like you to describe in more detail the one situation that makes you (or has made you) most sad. This could be something you are presently experiencing or something from the past. Begin by writing down what you remember of the sadness-inducing event(s) and continue by writing as detailed a description of the event(s) as is possible. If you can, please write your description so that someone reading this might even get sad just from learning about the situation. What is it like to be in this situation? Why does it make you so sad?

Please, write your answer in the box:

Control Treatment

Figure C.6: No Emotion Induction – Question 1

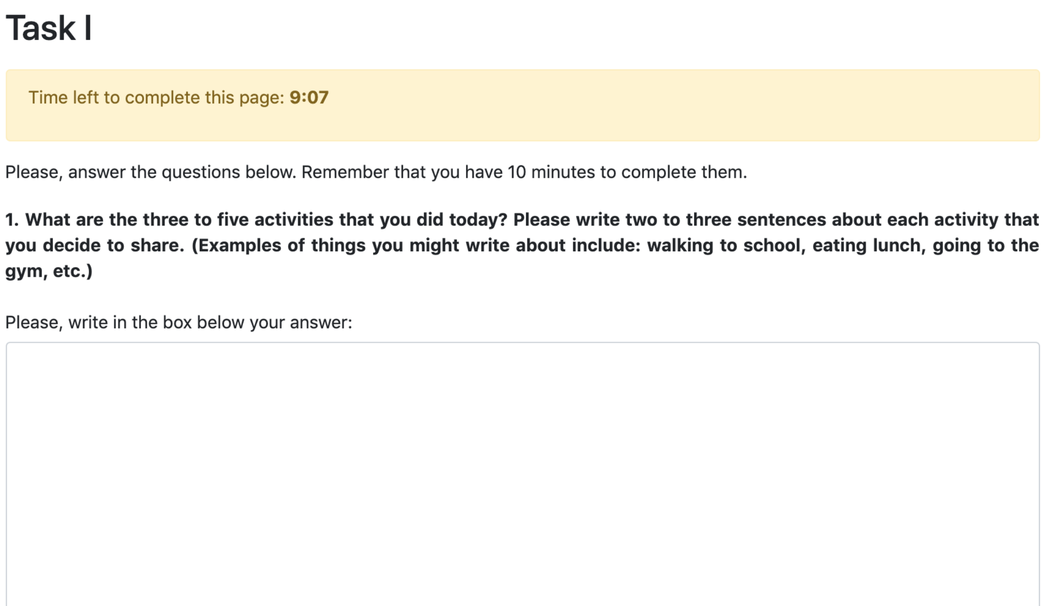
Task I

Time left to complete this page: **9:07**

Please, answer the questions below. Remember that you have 10 minutes to complete them.

1. What are the three to five activities that you did today? Please write two to three sentences about each activity that you decide to share. (Examples of things you might write about include: walking to school, eating lunch, going to the gym, etc.)

Please, write in the box below your answer:

Figure C.7: No Emotion Induction – Question 2

2. Now we'd like you to describe in more detail the way you typically spend your evenings. Begin by writing down a description of your activities and then figure out how much time you devote to each activity. Examples of things you might describe include eating dinner, studying for a particular exam, hanging out with friends, watching TV, etc. If you can, please write your description so that someone reading this might be able to reconstruct the way in which you, specifically, spend your evenings.

Please, write in the box below your answer:

C.2.2 The p-beauty Contest Game

Game Instructions

Figure C.8: p-Beauty Contest Game Instructions

Task 2 - Instructions

We will now describe to you the second part of the experiment. This will be a decision making task.

This part is made up of 10 rounds. You will be anonymously matched into groups of 3 participants. You will never get to know your group members' identity. You will stay in the same group for all 10 rounds.

In each round, you and your other 2 group members will separately choose a whole number between 0 and 100 (0, 100 or any whole number in between is allowed). The group member whose chosen number is closest to 70% of the average of all 3 chosen numbers will earn £10.00 and the other 2 group members will earn nothing in that round. If more than one group member chooses a number which is closest to 70% of the average of all 3 chosen numbers, the £10.00 will be split equally among the group members who chose the closest number (or numbers).

Importantly, at the end of the experiment you will be paid according to your earnings in one randomly drawn round.

At the end of each round, you will discover: (i) the numbers chosen by all your group members; (ii) the average of all 3 chosen numbers; (iii) what 70% of the average of all 3 chosen numbers was; and (iv) how much you earned in that round.

In the next screen you will also find a reminder of the rules of the task (shown at the bottom of this page too).

Please raise your hand if you have any questions.

To continue with the task, please type in the cell below the number "20".

Next

Game Play

Figure C.9: p-Beauty Contest Game Play

Round 1 - Your Choice

Please, now type below your chosen number from 0 to 100.

Next

Reminder: Instructions of Task 2

Your Group

You are matched with two other participants. The three of you will be playing the task together for 10 rounds.

Your Action in the task

In each round, each one of you will have to choose a whole number between 0 and 100.

Payoffs

One round will be randomly chosen for payment.

For the randomly chosen round, the group member whose number is closest to 70% of the average of all 3 chosen numbers will be paid £10.00 and the other 2 group members will be paid nothing.

If more than one group member chooses a number which is closest to 70% of the average of all 3 chosen numbers, the £10.00 will be split equally among the group members who chose the closest number (or numbers).

Figure C.10: p-Beauty Contest Game Feedback

Round 1 - Results

Below you can see the three numbers guessed in your group:

[10, 30, 34]

70% of the average of these numbers is 17.27.

And the closest guess was 10.

Your guess was 10.

Therefore, you have made the closest guess. If this round is randomly selected for payment, then your payoff in this task is £10.00.

Next

C.3 Further Text Analyses

C.3.1 General Affect

Figure C.11: General affect in the texts

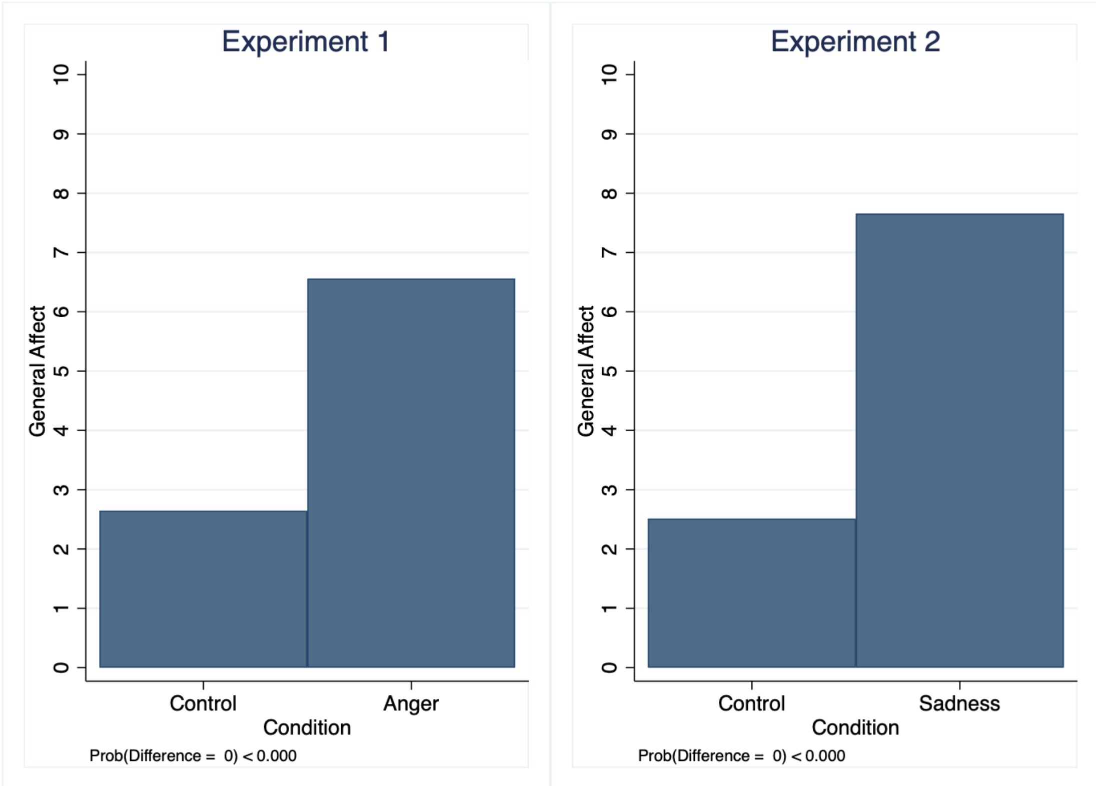

Notes: 1) The bars report the average "affect" in subjects' written words for the different inductions. 2) Analysis based on the LIWC2015 (Linguistic Inquiry and Word Count) dictionary (Pennebaker, 2015). 3) The notes report the results of the corresponding Mann-Whitney test.

C.3.2 Negative Affect

Figure C.12: Negative affect in the texts

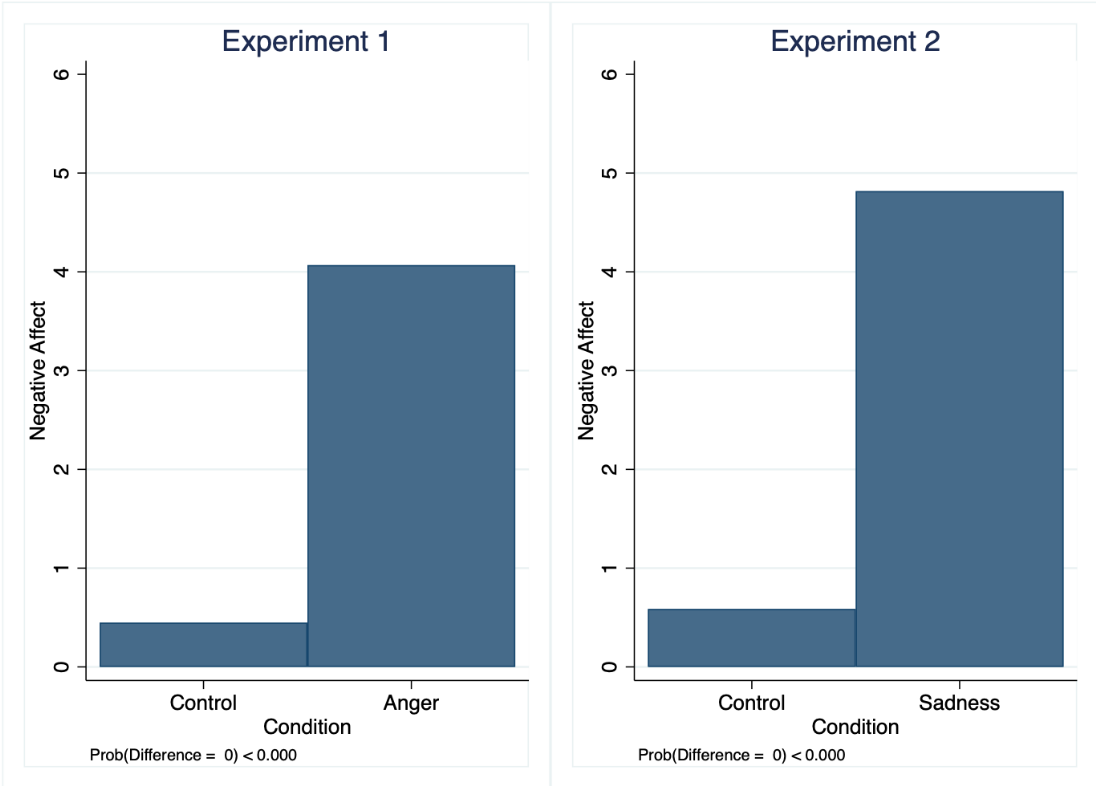

Notes: 1) The bars report the average negative affect in subjects' written words for the different inductions. 2) Analysis based on the LIWC2015 (Linguistic Inquiry and Word Count) dictionary (Pennebaker, 2015). 3) The notes report the results of the corresponding Mann-Whitney test.

C.3.3 Another Negative Emotion: Anxiety

Figure C.13: Anxiety in the texts

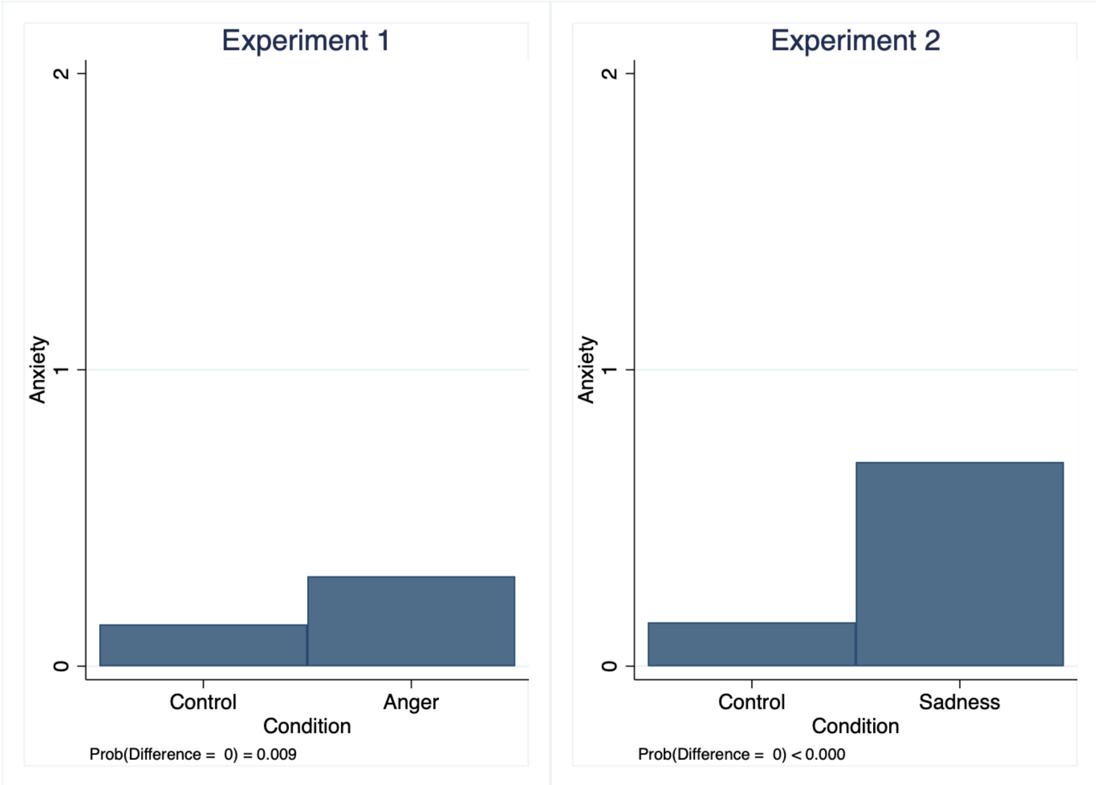

Notes: 1) The bars report the average anxiety in subjects' written words for the different inductions. 2) Analysis based on the LIWC2015 (Linguistic Inquiry and Word Count) dictionary (Pennebaker, 2015). 3) The notes report the results of the corresponding Mann-Whitney test.

C.4 Further Analyses

C.4.1 Emotional Self-assessment

Figure C.14 shows the effect of the induction using the questionnaire about the self-reported induction effectiveness. There is a significant difference in the levels of anger reported in the anger treatment compared to the control condition ($\Delta = 2.942$, p-value< 0.000) (top-left panel). This is also true for sadness but this difference ($\Delta = 1.525$, p-value< 0.000) is significantly lower (bottom-left panel). As expected, the opposite is true in the sadness experiment (Experiment 2). Here subjects report a significant difference in sadness (bottom-right panel) compared to the control condition ($\Delta = 3.791$, p-value<0.000). They also report a significant difference in reported anger ($\Delta = 2.149$, p-value<0.000) (top-right panel) but this latter difference is significantly lower compared to the former.

Figure C.14: Self-reported Anger and Sadness felt in the induction

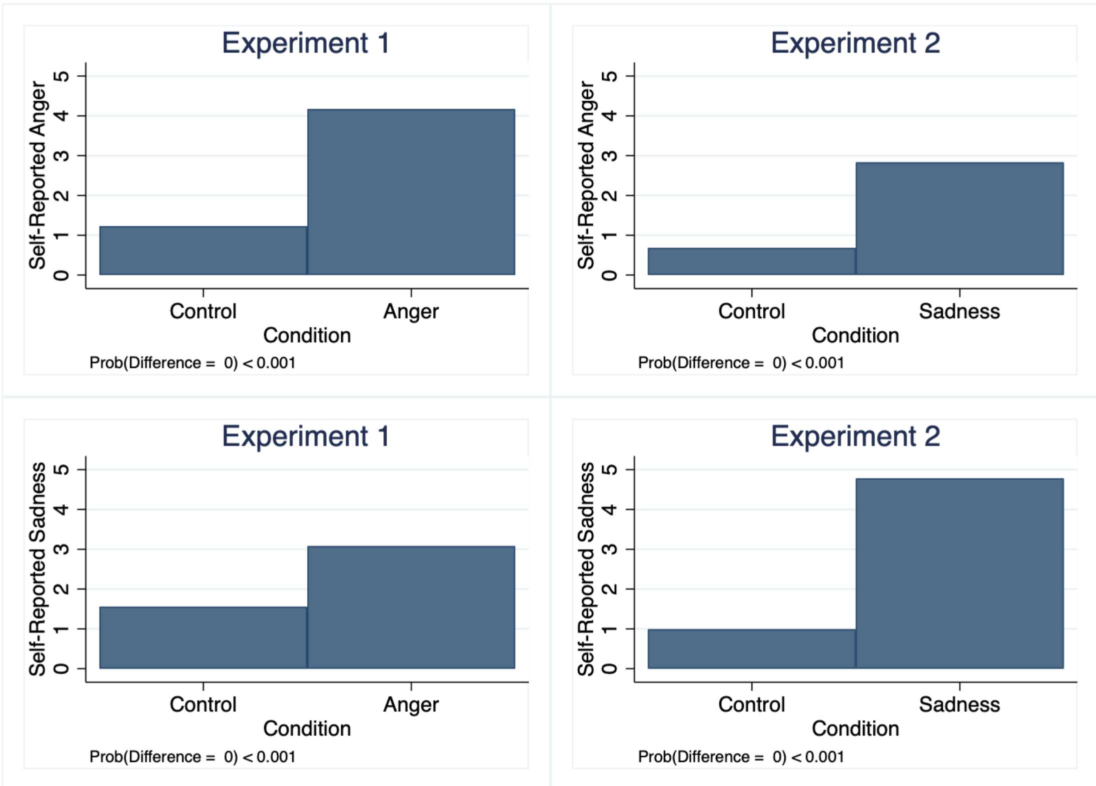

Notes: 1) The bars report the average difference of anger or sadness felt at the end of the sessions. 2) Questions are: "Please indicate the greatest amount of anger (sadness) you experienced while writing about the past life events"; and are coded from 0 (low) to 8 (high). 3) The notes report the results of the corresponding Mann-Whitney test.

C.4.2 General Affect at the Beginning of the Experiments

Figure C.15: General positive and negative affect at the outset of the experiment

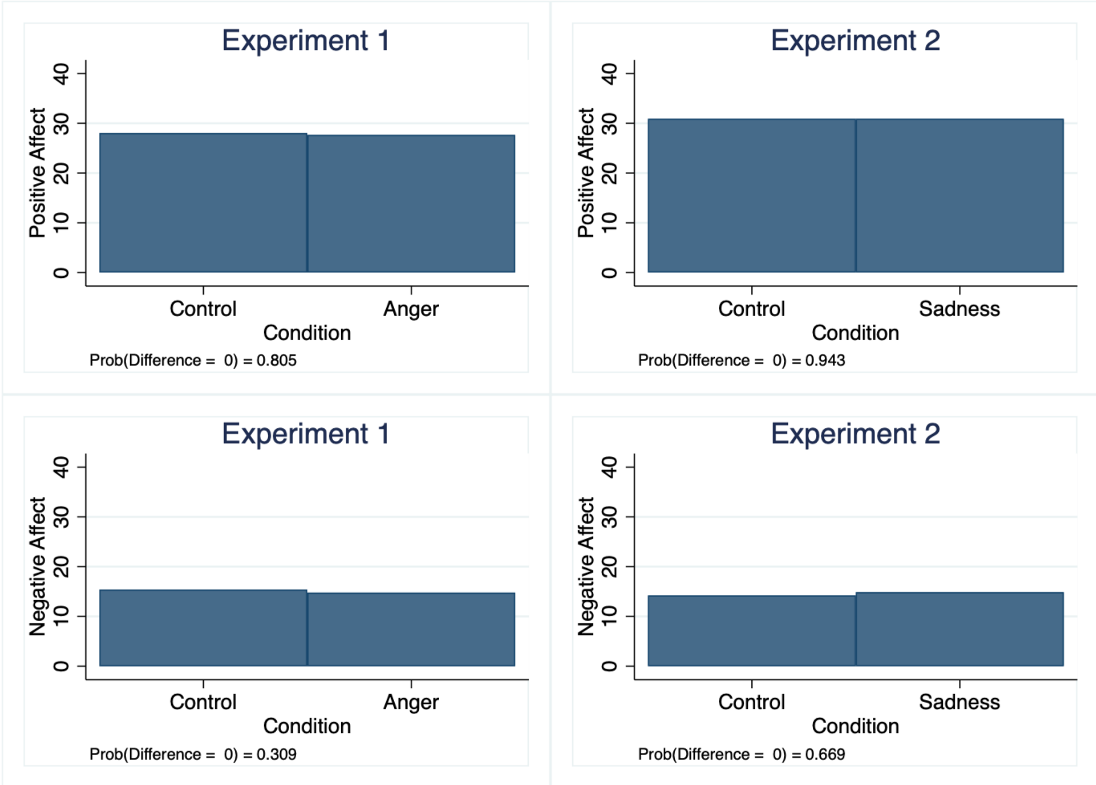

Notes: 1) The bars report the total positive (negative) affect experienced at the outset of the session. 2) Questions are takes from the PANAS questionnaire and ask: "Please, indicate the extent you are feeling this way right now", in terms of 20 scales or emotional states. Each item is rated on a 5-point scale of 1 (not at all) to 5 (very much). 3) Positive affect: Active, Alert, Attentive, Determined, Enthusiastic, Excited,Inspired, Interested, Proud, and Strong. Negative affect: Afraid, Ashamed, Distressed, Guilty, Hostile, Irritable, Jittery, Nervous, Scared, and Upset. 4) The total positive (negative) affect score is the sum of the scores in each positive (negative) emotion. 5) The notes report the results of the corresponding Mann-Whitney test.

C.4.3 The Effect of Anger and Sadness on Guesses in the Two Experiments

The mean unconditional guess across all rounds is 23.68 (s.d. 0.525) in Experiment 1 and 24.20 (s.d. 0.507) in Experiment 2. As Figure C.16 shows, treated subjects in Experiment 1 guessed on average higher numbers than those in the control. This holds true in almost every round and on aggregate (left panels). In Experiment 2, by contrast, no clear pattern emerges. Guesses by treatment are similar across rounds and on aggregate (right panels).

Figure C.16: The effect of anger and sadness on the average guess

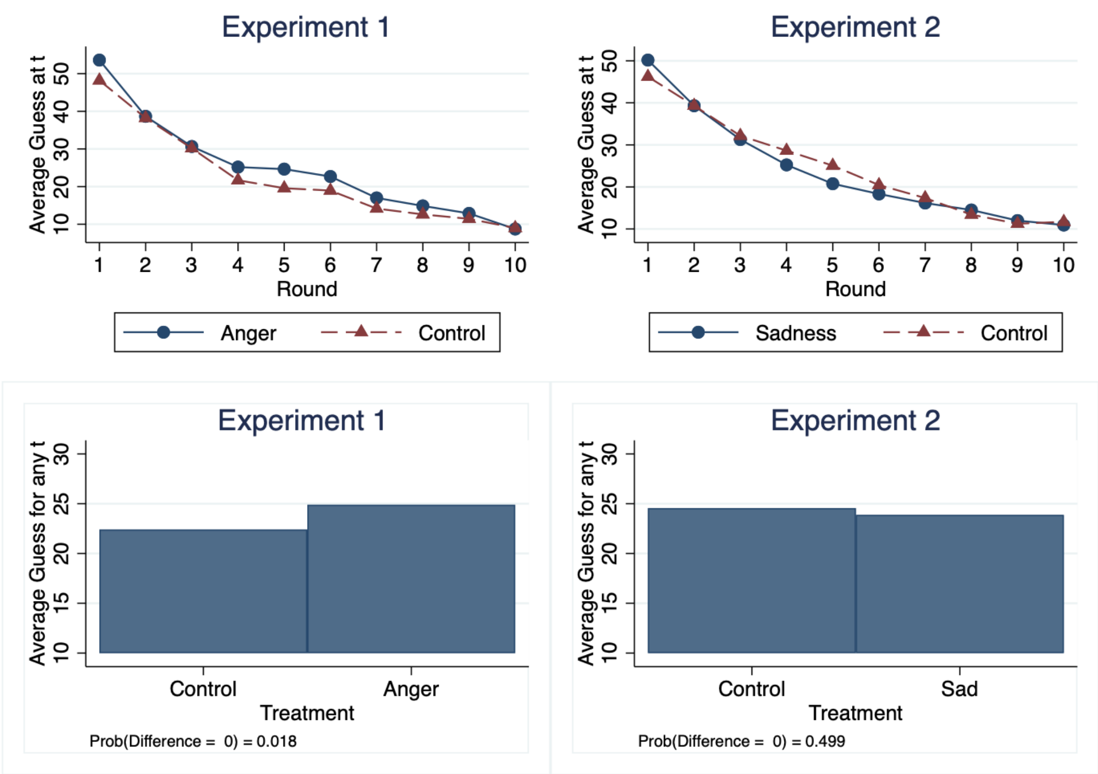

Notes: 1) The lines in the top panels report the average guess for each round of play by condition and experiment. 2) The bottom panels report the average guess across all rounds by condition and experiment. 3) The notes report the results of the corresponding t-test.

In the first round of Experiment 1, the average guess among those who undertook the anger exercise is 53.60, compared to 48.20 among those in the control group. This difference is significantly different from zero at the 10 percent significance level (p-value=0.098).[1] In Experiment 2, this difference is again larger among those in the

[1]In Subsections 3.2 and 3.2, we compute t-tests to assess average differences in guesses and

treatment group compared to those in the control group (50.17 vs. 46.21, Δ=3.96), although not significantly so (p-value=0.242).

We find that the average guess across all rounds among those who experienced the anger treatment is 24.89, while it is 22.40 among those in the control group. The difference (Δ=2.49) is significant (p-value=0.018). In contrast, in Experiment 2 the average guess among those in the sadness treatment is lower than the average guess among those in the control group (23.87 vs. 24.56), although the difference is not statistically significant (p-value=0.499).

The previous analysis on average guesses across all rounds does not consider that guesses are influenced by past behavior. Therefore, in order to take into account previous game play and group fixed effects we estimate the following model:

$$Guess_{i,j,t} = \beta_0 + \beta_1 Treatment_i + \beta_2 AverageGuess_{j,t-1} + \beta_3 t + \gamma_j + \epsilon_{i,j,t}; \quad \text{(C.1)}$$

where i indicates the subject in group j, while t is the round of play. Our dependent variable is the guess in the game, $Guess_{i,j,t}$. Our independent variable of interest is $Treatment_i$, which is a dummy variable indicating the emotion treatment individual i received in one of the two experiments. Control variables include: $AverageGuess_{j,t-1}$ that is the average guess in the previous round, t is the round of play, γ_j is the group-level effect, while $\epsilon_{i,j,t}$ is the error term.

We estimate Equation (C.1) by using an OLS model with group fixed effects. We cluster standard errors at the group level. The results are reported in Table C.5. Column (1) reports the results for Experiment 1; Column (2) reports those for Experiment 2; Column (3) reports the combination of the two. The anger treatment has a positive and significant effect on guesses. The guesses of subjects who experienced the anger treatment are more than two units higher on average, compared to subjects' guesses in the control (p-value=0.009). The sadness treatment has an insignificant negative effect on guesses.

payoffs across conditions and experiments.

Table C.5: The effect of the treatment on guesses in both experiments

	Experiment 1 Guess b/se	Experiment 2 Guess b/se	Experiments 1 & 2 Guess b/se
Anger Treatment	2.439**		2.439***
	(0.921)		(0.917)
Sadness Treatment		–0.626	–0.626
		(1.021)	(1.017)
Average Guess at $t-1$	0.092***	0.106***	0.098***
	(0.019)	(0.020)	(0.014)
Round	–2.114***	–2.028***	–2.078***
	(0.304)	(0.307)	(0.216)
Group Exp FE	Yes	Yes	Yes
N	1539	1620	3159
Individuals	171	180	351
R2	0.465	0.445	0.455

Notes: 1) OLS estimator; 2) Standard errors (shown in parentheses) are clustered at the group level. 3) * p-value<0.1, ** p-value<0.05, *** p-value<0.01.

C.4.4 Further Econometric Analysis

Table C.6: The effect of anger and sadness on response times in the two experiments

	Experiment 1 Response Time b/se	Experiment 2 Response Time b/se	Experiments 1 & 2 Response Time b/se
Anger Treatment	0.339		0.339
	(0.469)		(0.467)
Sadness Treatment		-0.027	-0.027
		(0.462)	(0.460)
Average Guess at $t-1$	0.005	-0.010	-0.002
	(0.005)	(0.009)	(0.005)
Round	0.044	-0.106	-0.024
	(0.077)	(0.121)	(0.724)
Group Exp FE	Yes	Yes	Yes
N	1539	1620	3159
Individuals	171	180	351
R2	0.086	0.159	0.136

Notes: 1) OLS estimator. 2) Standard errors (shown in parentheses) are clustered at the group level. 3) * p-value<0.1, ** p-value<0.05, *** p-value<0.01.

C.4.5 Different Hypothesis about the Distributional Form in the Structural Analysis

Table C.7: Estimated Level-k types by condition: alternative distribution

	Experiment 1		Experiment 2	
	Anger	Control	Sadness	Control
Poisson Distribution				
Level 0	0.257	0.200	0.197	0.232
Level 1	0.660	0.650	0.735	0.666
Level 2	0.000	0.071	0.000	0.034
Level 3	0.084	0.080	0.067	0.062
Level 4	0.000	0.000	0.001	0.006
Log likelihood	-3,202	-2,942	-3,365	-3,056
Average Level-k	0.910	1.030	0.942	0.946