Jan 2020 No.452

Attribution Bias by Gender: Evidence from a Laboratory Experiment

James Fenske, Alessandro Castagnetti and Karmini Sharma

WORKING PAPER SERIES

Centre for Competitive Advantage in the Global Economy

Department of Economics

Attribution Bias by Gender: Evidence from a Laboratory Experiment*

James Fenske ⓡ
Alessandro Castagnetti ⓡ
Karmini Sharma ⓡ

December 25, 2019

Abstract

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. We find suggestive evidence 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.

*We thank Arun Advani, Robert Akerlof, Guillaume Fréchette, Sharun Mukand, Kirill Pogorelskiy, Jan Potters, Eugenio Proto, Roland Rathelot, Chris Roth, Andy Schotter, Daniel Sgroi, and seminar participants in Warwick for helpful comments. Financial support by CAGE, ESRC Grant Ref ES/L011719/1, is gratefully acknowledged. University of Warwick Economics Department IRB approval obtained on 23-07-2018.

Alessandro Castagnetti: University of Warwick, `S.castagnetti[at]warwick[dot]ac[dot]uk`
James Fenske: University of Warwick, `J.fenske[at]warwick[dot]ac[dot]uk`
Karmini Sharma: University of Warwick, `K.sharma.1[at]warwick[dot]ac[dot]uk`

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. In psychology, 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, while also attributing 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 for *others*, and has found mixed evidence on whether observers are more likely to attribute men's success 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). While this literature has not focused on cases where outcomes realized by the individual being evaluated affect the payoff of the individual making the evaluation, recent empirical evidence in economics suggests that attribution bias by gender may be at work 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 *et al.*, 2017). However, many other variables may be at work in these real-world environments that cannot be controlled for – these include agents' real contributions to outcomes, prior experience, and unobserved characteristics, among others. It is not possible to completely rule out factors other than attribution bias that might drive these differences in outcomes by gender.

In this paper, we present evidence from a laboratory experiment in order to test for the presence of attribution bias by gender in a controlled environment. A lab experiment provides a controlled setting in which other factors are unlikely to influence participants' behavior. In particular, we employ a principal-agent setup. Participants in this experiment are first randomly divided into two roles: principals and agents. In each round (out of 20), they are randomly matched into pairs. Agents perform a task for their principals in each round. Principals are rewarded based on the outcome of the agents' performance, while agents are paid by their principals after the outcome is revealed. Importantly, whether the outcome produced by the agent takes a high or low value is not a deterministic function of the agent's performance. It also depends on a random component. In each interaction, principals are shown information that allows them to identify the agent's gender. This piece of information is conveyed through agents' (nick-)names, and presented along with other demographic information in order 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. That is, following a high outcome, we test whether principals are more likely to attribute it to the agent's ability if male, while 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 sex.

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.

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)). 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 too have been shown to suffer from attribution bias (Deaux and Emswiller, 1974; Dweck *et al.*, 1978; Espinoza *et al.*, 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, and 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, and so isolates the dispositional and situational factors affecting the resulting 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 (Flory *et al.*, 2014).[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 *et al.*, 2017), self-beliefs (Beyer, 1998; Keller, 2001) and perceived performance (BenYishay *et al.*, 2020). 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 *et al.*, 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 *et al.*, 2018; Bordalo *et al.*, 2019; Carlana, 2019; Coffman, 2014; Coffman *et al.*, 2019; Milkman *et al.*, 2013). In particular, this literature finds that stereotypes about tasks lead to biased judgments 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.

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

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 design features all these pieces.

A detailed description of the experiment is presented below. First, we asked participants to fill out a demographic questionnaire. Second, we randomized participants into two roles: principals and agents. In each of the 20 rounds, agents and principals were matched into pairs using the stranger-matching protocol.[3] The agent then performed a task for the principal. The agent's performance influenced the resulting output, but not deterministically. This output determined the principal's earnings in that round. The principal then proceeded to pay the agent for his performance. In each round, we elicited agents' and principals' beliefs about the agent's contribution to the realized outcome. Finally, subjects were asked to complete a series of questions about the experimental task.

2.1 The Experiment

At the outset of the experiment, participants were asked to complete 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 into one of two roles: principals and agents. They were informed that these roles were fixed for the whole duration of 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.[4] 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 to complete a set of comprehension questions. Participants could not continue with the experiment until they had answered all questions correctly.

2.1.1 Description of a Round

In each round, principals and agents were matched into pairs following the stranger-matching protocol.[5] Participants were informed that, although they could earn money in each of the

[3]That is, at every round principals and agents were randomly rematched.

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

[5]We opted for the stranger-matching protocol since it avoids reputation building and related strategic concerns.

20 rounds, at the end of the experiment only one would be randomly selected to count for the final payments.

Agent's performance and output produced At the beginning of each round, the agent performs a task, consisting of a fixed number of questions. These are to be performed in 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 Once the 45 seconds have passed, the principal is shown the realization of the lottery, which constitutes 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 is fully aware of the mapping between the number of correctly answered questions and the probability of high output. The principal proceeds by choosing 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.[6] This separate 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 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 to indicate the number of questions that she thought that the agent had solved correctly. We incentivized this question by paying ₹50 if the answer was correct. In some sessions we also asked the same question, but did so while the agent was performing the task and so before the outcome of the lottery was realized. That is, we asked the principal to indicate her prior belief about the number of questions that the agent would solve correctly. Finally, we also asked the principal two unincentivized questions. We asked the principal to guess how many questions she thought

[6] At the time of the experiment, this amount corresponded to £3.92 (exchange rate as of July 2018: £1.00 = ₹89.21). We implemented this payment procedure following the same considerations as in Gurdal *et al.* (2013). In particular, 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 allow us to eliminate any efficiency motives (in terms of subjects versus experimenter considerations) that the principal might have.

that the agent attempted, and a hypothetical question on 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 also asked him his belief about the percentage of the ₹350 that he would receive from the principal.

2.1.2 Debriefing

Finally, after the two tasks have been 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 what task they thought as more difficult (out of the two) 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 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 through the means of nicknames. That is, in each round the computer software would assign a gender-congruent nickname to the agent. As the experiment took place in India, the realized nickname was randomly selected from a list of popular Indian names. Since we used only first names, they did not signal caste. All names were the most popular Hindu names. For instance, female names included "Akansha", "Neha" and "Priya", while for male names these included, among others, "Amit", "Ashish" and "Nitin".[7]

The use of nicknames instead of a direct statement of the agent's gender was implemented to mask the fact that our research question was gender-related and, therefore, to prevent potential distortions due to demand effects and social desirability concerns. Moreover, we opted for nicknames as opposed to 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.

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

2.2.1 The Tasks and Output

In each session agents performed two different tasks. However, since we vary these tasks across sessions, we have a total of four tasks: a math task, a Raven task, an effort task, and a memory task. We now provide a description of each of these tasks.

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. In particular for our experiment, we implemented the matrices from the Raven Advanced Progressive Matrices (APM). This test is commonly used to measure fluid intelligence (Carpenter *et al.*, 1990).

The effort task For this 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 task was a working memory exercise. In particular, 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 in the following 20 seconds.

2.2.2 Lotteries and Output

As indicated previously, each correct answer in a given task 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.[8] 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

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

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 vary the level of the high output to see whether principals' payments and beliefs are 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.

2.3 Attribution Bias by Gender

Our experiment is designed to analyze whether principals make biased attributions regarding the performance of the agents. In particular, to capture attribution bias, we designed an environment in which outputs represent noisy signals of the agent's performance. That is, the output produced in each round is a function of the number of questions answered correctly by the agent, but also luck. The principal therefore has to base her payment on the basis of the lottery's outcome.

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 male, while she will attribute it to luck if the agent is female. Similarly, following a low output, the principal will attribute it to misfortune if the agent is male while to performance if the agent is a female. Similarly, this difference in the principal's beliefs by gender would affect the way principals make their payments to agents. We therefore implement the following tests:

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

2. We test whether the principal's payments are higher for male agents as compared to 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 as compared to female agents.

3 Experimental Results

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 a total of 84 subjects and conducted 5 sessions that lasted around 75 minutes each. The participants earned on average ₹510, which includes the show-up fee of ₹250. We programmed the experiment with oTree (Chen *et al.*, 2016).

We begin by examining agents' performance and their beliefs in Section 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' performance in Section 3.3. We thus analyse whether principals make biased attributions and payments depending on the gender of their matched agents. We present our econometric specifications in Section 3.4, and the results of these estimations in Section 3.5. In Section 3.6 we discuss alternative factors that might be driving our results: the salience of gender information, principals' prior beliefs about the agents' performance, 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 A.1, A.2, and A.3 in Appendix A.

3.2 Agents

3.2.1 Agents' Performance 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 answers correct. If we look at performance broken down by gender in Table 1, we find no difference in performance by gender across tasks.[9]

[9] 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 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, agents' beliefs about their own performance 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 believe they have solved an excess of roughly 15% questions correctly.

Table 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 on the gender of the agent. Standard errors are clustered at the individual level.

3.2.2 Agents' Beliefs about Outcomes and Expected Payments

Beliefs about realized outcomes and expected principals' payments At the top of Table 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 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.

In sum, we find that, while performance differs by task, it does not significantly differ by agent 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 not statistically significantly so.

Table 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	0.0.55	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.

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).

3.3.1 Wages and Principals' Beliefs

Wages In Figure 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 1: Principals' payment decisions by realized outcome

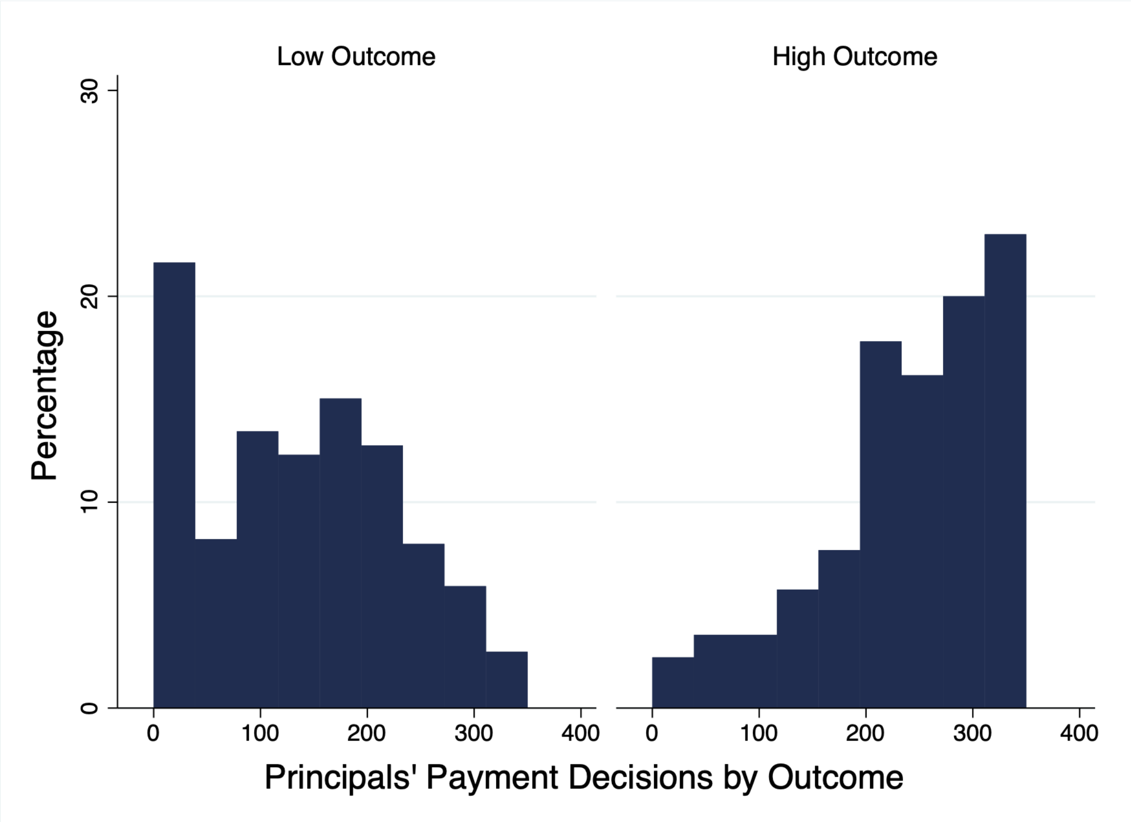

Notes: the histograms show the distribution of principals' payment decisions by realized outcome.

Beliefs Figure 2 shows that principals' beliefs about their agents' performance 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 as compared to when it is low. A Mann-Whitney test shows the the distribution of principals' beliefs differ significantly by the realized outcome (p-value<0.00).

Figure 2: Principals' beliefs by realized outcome

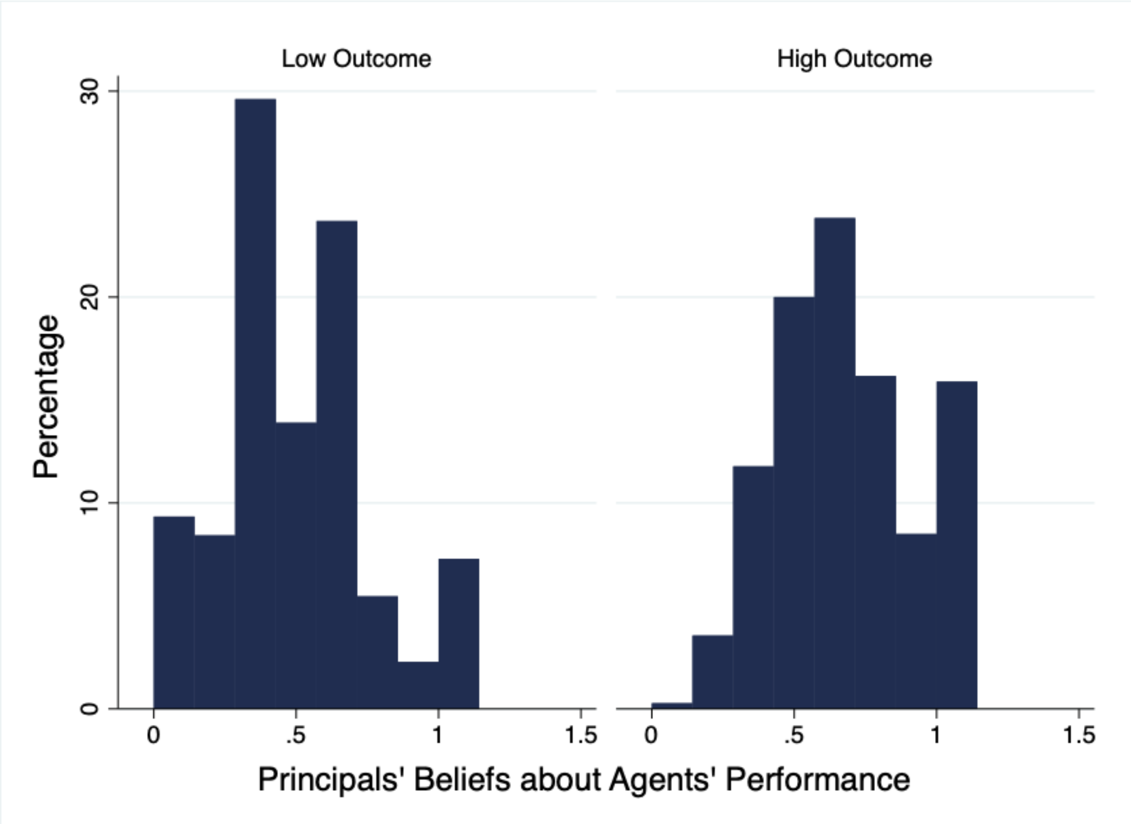

Notes: the histograms show the distributions of principals' beliefs about the matched agents' performance by realized outcome.

3.3.2 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 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).[10]

[10] In appendix B, 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 3: Principals' wages by realized outcome and agents' gender

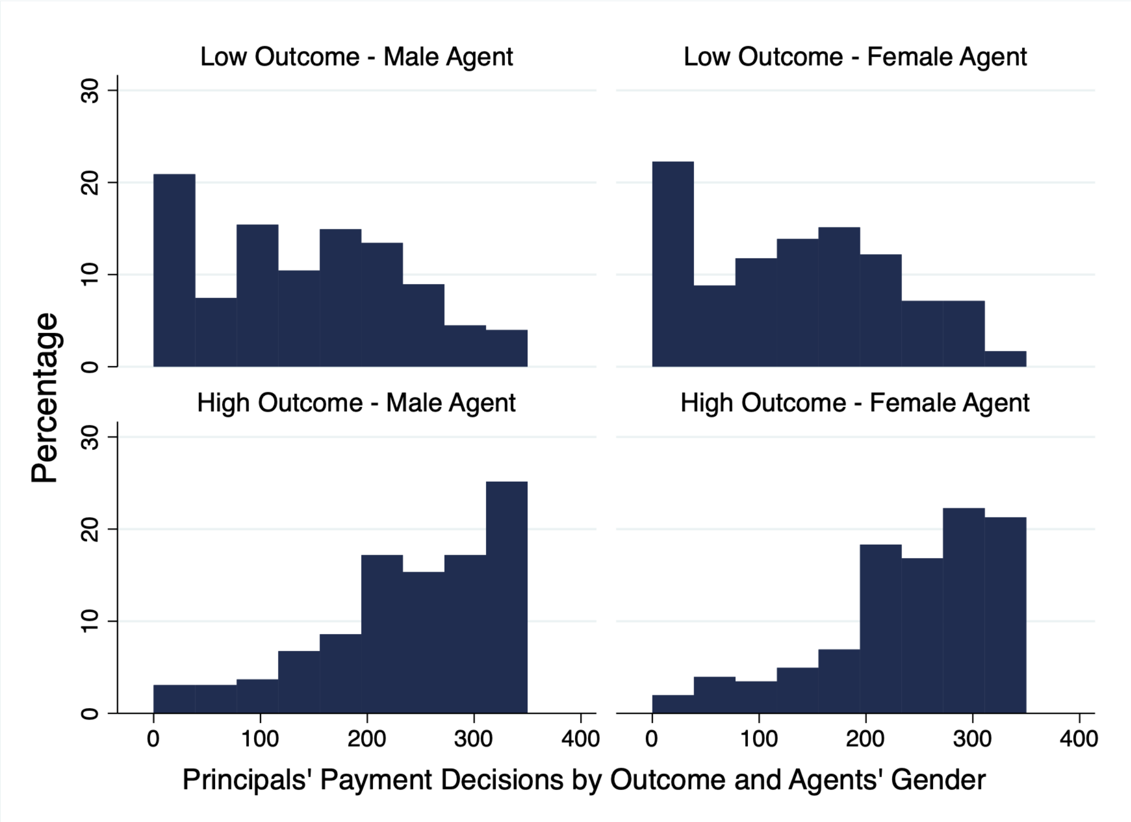

Notes: the histograms show the distribution of principals' payment decisions by realized outcome and the gender of the matched agents.

Beliefs by gender of the agent The results for beliefs match those for wages. Figure 4 shows that, while beliefs about the number of questions solved correctly are heavily influenced by the realized outcome, they do not shift depending on 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, respectively).

Figure 4: Principals' beliefs by realized outcome and agents' gender

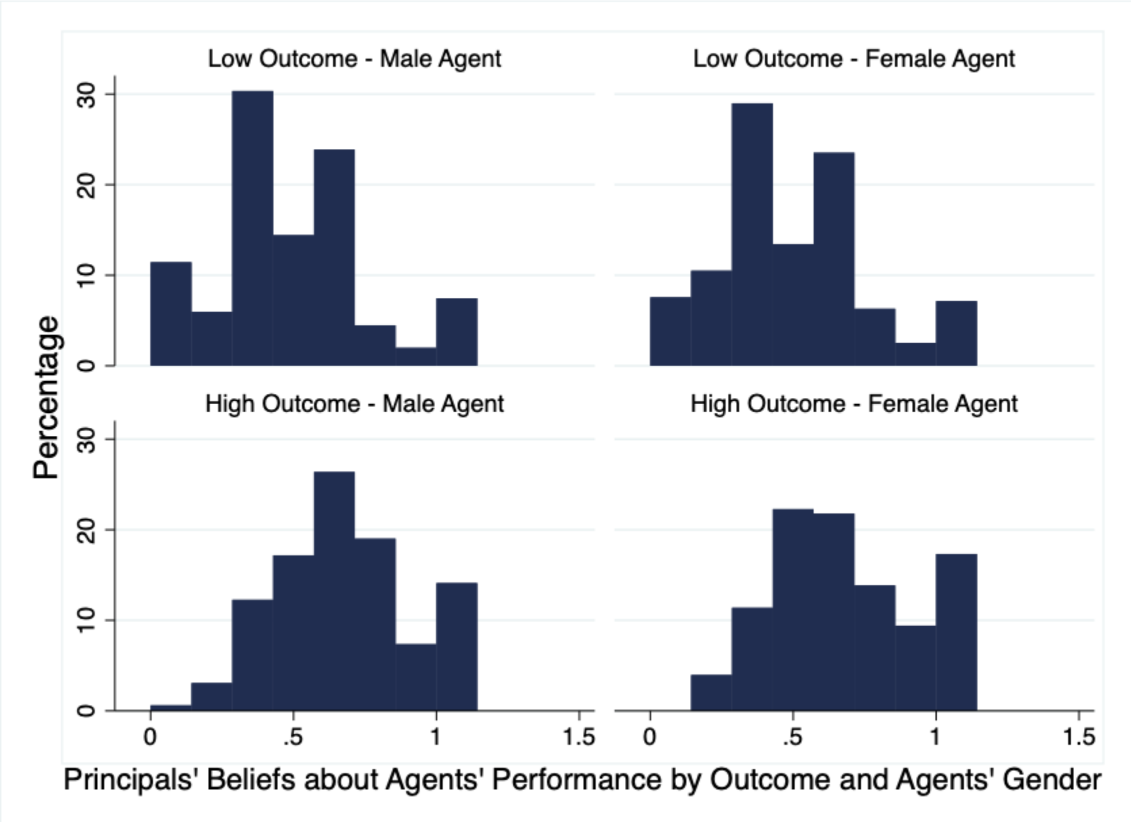

Notes: the histograms show the distributions of principals' beliefs about the matched agents' performance by realized outcome and gender of the matched agents.

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

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:

$$Y_{ij} = \beta_0 + \beta_1 Z_{ijr} + \beta_2 Female_i + \beta_3 Female_i \times Z_{ijr} + x'_{ijr}\beta_4 + \epsilon_{ij} \tag{1}$$

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 plays any role in the payments made to the agent. We estimate this using the following econometric specification:

$$\begin{aligned} Y_{ij} = \beta_0 + \beta_1 Z_{ijr} + \beta_2 Female_i + \beta_3 Female_i \times Z_{ijr} + \\ \beta_4 Female_j + \beta_5 Female_j \times Z_{ijr} + x'_{ijr}\beta_6 + \epsilon_{ij} \end{aligned} \tag{2}$$

$Female_j$ 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:

$$\begin{aligned} Y_{ij} = \beta_0 + \beta_1 Z_{ijr} + \beta_2 Female_i + \beta_3 Female_i \times Z_{ijr} + \beta_4 Female_j \\ + \beta_5 Female_j \times Z_{ijr} + \beta_6 SameGender_{ij} + x'_{ijr}\beta_7 + \epsilon_{ij} \end{aligned} \tag{3}$$

β_6 here captures whether being matched to an agent of the same gender leads to any differential effect on payments made by principals.

3.5 Econometric Results

The results for principals' payment decisions (mean ₹184 and s.d.107.15) are shown in Table 4. As is apparent from columns 1 to 4, the outcome of the lottery is important in determining the payment made to the agent.[11] 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 both insignificant and small. The coefficient on the female agent

[11]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.

Table 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.

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).

If we perform the same regressions for beliefs (mean 0.54 and s.d. 0.26), we find the same patterns. Table 5 shows that principals' beliefs are significantly shaped 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 increase 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 4, though this difference is not significant.

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 statistically insignificant

Table 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.

and larger payment made to female agents equivalent to 1.3% of the mean. Egan *et al.* (2017), by contrast, finds 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. Nevertheless, 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 insignificant 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 in column 4 of Table 4. In particular, principals 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 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 supporting evidence of gender-biased attributions.

3.6 Robustness Checks

3.6.1 Session and Round Fixed Effects

In the appendix, we show that our results on principals' payments are robust to including session fixed effects (see Table C.4) and round fixed effects (see Table C.5). Running regressions without controls shows that the main result on principals' payments is robust to not including controls (see Table C.6). 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 C.7) and round fixed effects (see Table C.8). Results without controls are similar in magnitude and significance, as shown in Table C.9.

3.6.2 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, we report similar regressions restricting the sample to the first ten rounds in Tables C.10 and C.11. The results for attribution bias by gender hold as for the rest of the sample. 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 statistically insignificant. For the case of payments, the coefficient on same gender becomes larger by ₹7 than the estimate obtained using the whole sample. In Table C.11, 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.

3.6.3 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 C.12 and Table C.13, 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 statistically insignificant and 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.

3.6.4 Alternative Dependent Variables

To evaluate whether attribution bias manifests in changes in the shape of the payment distribution, we have estimated Equation (1) using alternative definitions of the dependent variable. We thus estimate the following econometric specification:

$$Y_{ijx} = \beta_0 + \beta_1 Z_{ijr} + \beta_2 Female_i + \beta_3 Female_j + x'_{ijr}\beta_3 + \epsilon_{ij} \tag{4}$$

This equation is the same as the one in Equation (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 C.14. For all these different cutoffs, we find similar results in that the gender of the agent does not play any role in driving payments.

3.6.5 Results for Different Tasks

We show in Table C.15 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. 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 10% level for the memory task, though the coefficient is now positive. Similarly, in Table C.16 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.

4 Possible Threats

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 3.3.1 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.

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,[12] none guessed it was about gender. The most common guesses included 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".[13]

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.

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.

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.

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

[13]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.

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 C.17 and C.18).

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.

References

Abeler, J., Falk, A., Goette, L. and Huffman, D. (2011). Reference points and effort provision. *American Economic Review*, **101** (2), 470–92.

Alan, S., Ertac, S. and Mumcu, I. (2018). Gender stereotypes in the classroom and effects on achievement. *Review of Economics and Statistics*, **100** (5), 876–890.

Azmat, G. and Petrongolo, B. (2014). Gender and the labor market: What have we learned from field and lab experiments? *Labour Economics*, **30**, 32–40.

Becker, G. S. (2010). *The Economics of Discrimination*. University of Chicago press.

BenYishay, A., Jones, M., Kondylis, F. and Mobarak, A. M. (2020). Gender gaps in technology diffusion. *Journal of Development Economics*, **143**, 102380.

Bertrand, M. (2011). New perspectives on gender. In *Handbook of Labor Economics*, vol. 4, Elsevier, pp. 1543–1590.

— and Duflo, E. (2017). Field experiments on discrimination. In *Handbook of Economic Field Experiments*, vol. 1, Elsevier, pp. 309–393.

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

Bordalo, P., Coffman, K., Gennaioli, N. and Shleifer, A. (2019). Beliefs about gender. *American Economic Review*, **109** (3), 739–73.

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

Carpenter, P. A., Just, M. A. and Shell, P. (1990). What one intelligence test measures: A theoretical account of the processing in the raven progressive matrices test. *Psychological review*, **97** (3), 404.

Chen, D. L., Schonger, M. and Wickens, C. (2016). oTree — an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, **9**, 88–97.

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

—, Collis, M. and Kulkarni, L. (2019). Stereotypes and belief updating. *Working paper*.

Croson, R. and Gneezy, U. (2009). Gender differences in preferences. *Journal of Economic Literature*, **47** (2), 448–74.

Deaux, K. and Emswiller, T. (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.

Dweck, C. S., Davidson, W., Nelson, S. and Enna, B. (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, M. L., Matvos, G. and Seru, A. (2017). When harry fired sally: The double standard in punishing misconduct. *National Bureau of Economic Research Working Paper No. 23242*.

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

Fennema, E., Peterson, P. L., Carpenter, T. P. and Lubinski, C. A. (1990). Teachers' attributions and beliefs about girls, boys, and mathematics. *Educational Studies in Mathematics*, **21** (1), 55–69.

Flory, J. A., Leibbrandt, A. and List, J. A. (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.

Glover, D., Pallais, A. and Pariente, W. (2017). Discrimination as a self-fulfilling prophecy: Evidence from french grocery stores. *The Quarterly Journal of Economics*, **132** (3), 1219–1260.

Gurdal, M. Y., Miller, J. B. and Rustichini, A. (2013). Why blame? *Journal of Political Economy*, **121** (6), 1205–1247.

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

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

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

Landsman, R. (2019). *Topics in Labor and Experimental Economics*. Ph.D. thesis, University of Pittsburgh.

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

Milkman, K., Akinola, M. and Chugh, D. (2013). Discrimination in the academy: A field experiment. *SSRN, Working Paper.*

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

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

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

Räty, H., Vänskä, J., Kasanen, K. and Kärkkäinen, R. (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.

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

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

— (2019). Interpreting signals in the labor market: evidence from medical referrals. *Working Paper.*

Selody, K. (2010). Board independence and the gender pay gap for top executives. *Working paper.*

Stipek, D. J. and Gralinski, J. H. (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.

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

A Appendix A: Summary Statistics

Table A.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 A.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 A.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 Appendix B

B.1 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

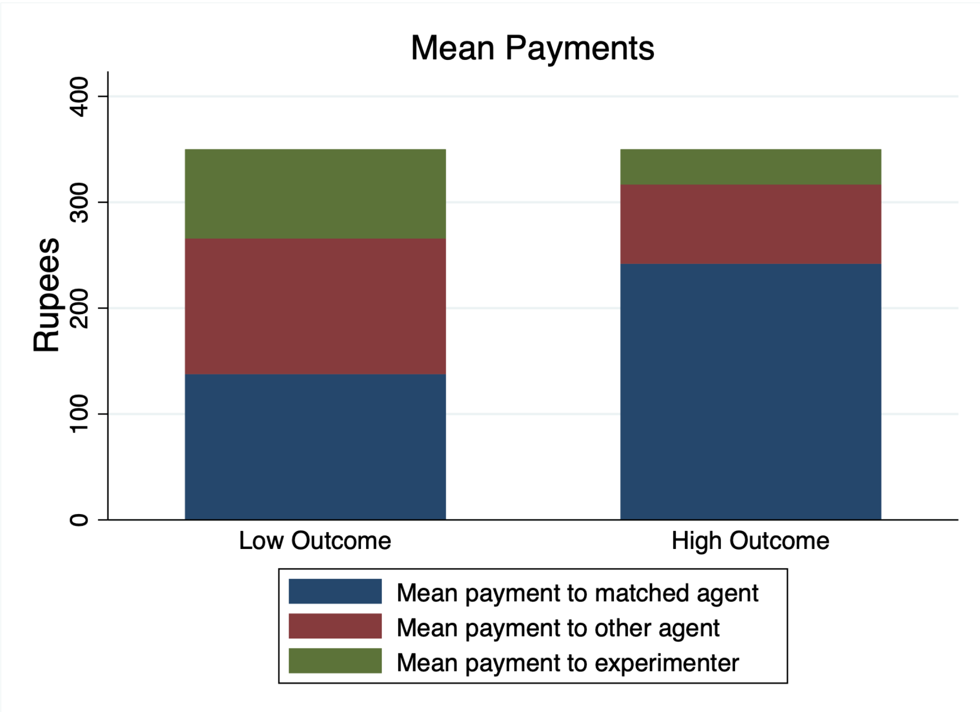

B.2 Mean Payments to Each Party by Realized Outcome and Gender of the Agent

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

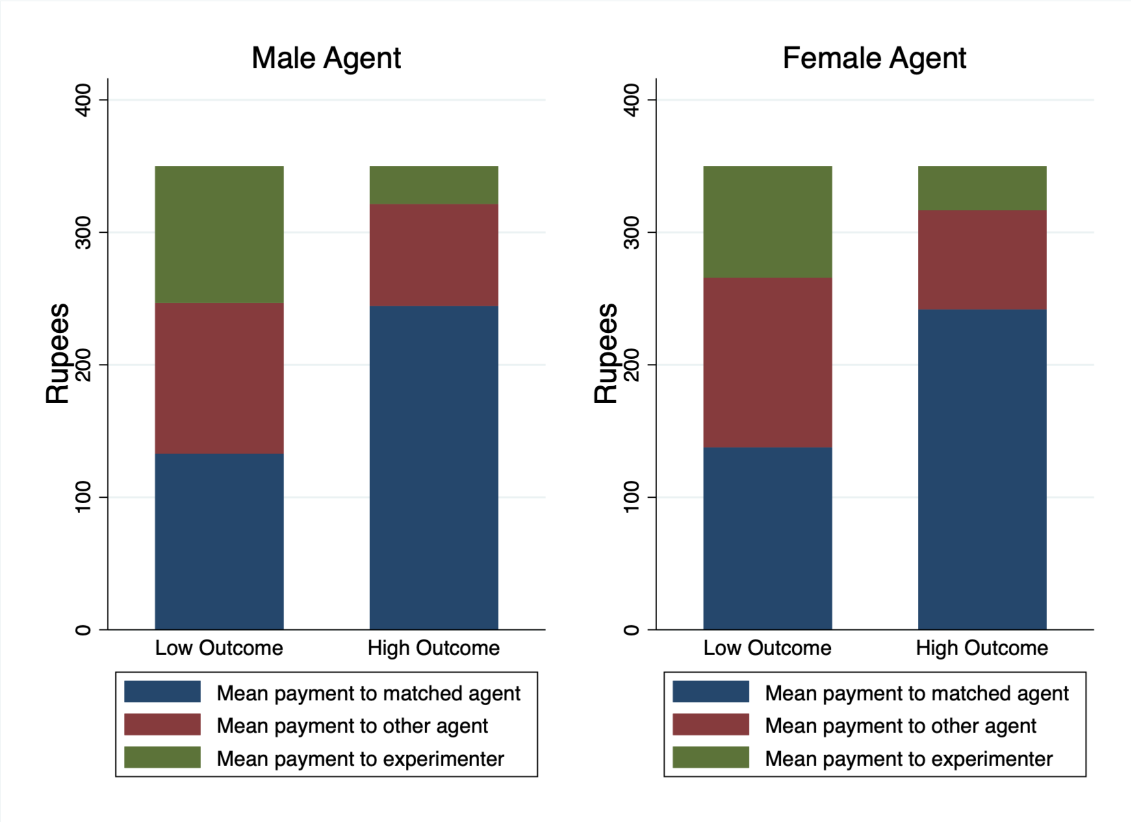

C Appendix C

C.1 Robustness Checks for Principals' Payment Decisions and Beliefs

Table C.4: 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 C.5: 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 C.6: 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 C.7: 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 C.8: 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 C.9: 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)
Demographics	✓	✓	✓	✓
Task Controls	✓	✓	✓	✓
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 C.10: 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 C.11: 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 C.12: 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 C.13: 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 C.14: Regression results for alternative definition of dependent variable

	(≥ 50)	(≥100)	(≥150)	(≥200)	(≥250)	(≥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 C.15: 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.

Table C.16: 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.

C.2 Principals' Payment Decisions and Beliefs by Agents' Age

Table C.17: 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 C.18: 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.