Employee-Side Discrimination: Beliefs and Preferences

- Evidence from an Information Experiment on Job-Seekers *†

Md Moshi Ul Alam ‡
University of Wisconsin-Madison
(*Job Market Paper*)

Sanket Roy
American Univeristy of Sharjah

Mehreen Mookerjee
Zayed University

November, 2020

Link to the most recent version

Abstract

If employees do not prefer to work for managers from a certain group (e.g. females), it could lead to lower promotion rates for members of that groups. We document novel evidence of employee discrimination, by studying which gender of managers, do employees prefer to work for. Additionally their choices are also driven by their beliefs on attributes that they may care about, but do not have information on (e.g. manager quality). Following the hypothetical choice methodology, we isolate employee preferences on manager's gender. To isolate employee beliefs on manager quality, we design a novel within-employee information experiment. In absence of information on manager quality, employees are indifferent between male and female managers. However, given information on manager quality, employees prefer to work for female managers. Using a structural model of job choice, we estimate employees are willing to give up 1.3-2.2% of average annual wages to work for female managers, on average. Hence in the absence of additional information on manager quality, employees believed female managers to be worse in quality. We corroborate this result of negative beliefs on female manager quality in an expost survey where we directly elicit employee beliefs. The results do not statistically differ by the gender of the employees. However, we do not find evidence of negative beliefs against female managers among science majoring female employees. The results suggest that glass ceilings for females at the managerial level, driven by discrimination by firm executives—who decide on promotion—could be potentially underestimated.

*For guidance, we thank Christopher Taber, Jeffrey Smith, Jesse Gregory and Matthew Wiswall. For helpful comments, we thank Peter Arcidiacono, John Kennan, Rasmus Lentz, Corina Mommaerts, Laura Schecter, Basit Zafar, seminar participants at the Department of Economics, Department of Agricultural and Applied Economics at University of Wisconsin-Madison and those of three anonymous referees who supervised the grants competition at the American University of Sharjah. For helpful discussions we thank Monica, Gary Baker, Natalie Duncombe, Amrita Kulka, Dennis McWeeny and Elan Segarra. We are grateful to Suraj Das, Iman Kalyan Ghosh, Kabir Rana, Shreya Seal and Aishani Sengupta for excellent research assistance in administering the survey. Errors, if any, are our own.

†This research is funded by the American University of Sharjah. Grant: FRG19-M-B41. IRB approval protocol number: 19-509.

‡Corresponding author: alam4@wisc.edu

1 Introduction

In this paper, we answer a novel question on whether employee discrimination on managers gender, could contribute towards generating glass-ceilings for females at the managerial levels. If employees do not prefer to work for females then they would need a wage premium to work for female managers—lowering the incentives of firm executives to hire female managers or promote females to managerial positions.[1] Additionally employee beliefs also play an important role–because when employees are searching for jobs, or switching teams, their choices are also driven by their beliefs on attributes that they may care about (e.g. manager quality), but do not have information on.[2] Existence of such preferences and beliefs would affect search behavior and the equilibrium match observed in the data. Consequently without taking into account this supply side selection, estimates of group level inequalities are potentially biased. With mounting evidence on employees sorting in jobs on non-pecuniary benefits (Dey & Flinn (2005), Wiswall & Zafar (2018), Mas & Pallais (2017), Sorkin (2018), Taber & Vejlin (2020)), it is likely that employees will sort on the gender of the manager if they have preferences and beliefs on the manager's gender and quality affecting their utility. Such sorting into jobs would be reflected in the wages through compensating differentials (Rosen 1986). We think of employee-side discrimination in the same spirit - as a form of selection, where individuals are willing to give up wages to work for their preferred managers in an otherwise identical job.[3]

In this paper we provide novel evidence on preferences and beliefs of employees leading them to discriminate on the manager's gender in choosing jobs. To ensure that demand side selection (preferences of employers) do not confound our results, we do not use data on realized job choices.[4] We design and conduct a survey following the hypothetical choice method-

[1] A reason why employees may have such preferences could be motivated by the example of in-group mentoring (Athey, Avery & Zemsky 2000) - if human capital accumulated on the job through mentoring is higher when both the employer (mentor) and employee (mentee) are from the same group (e.g. race, ethnicity or gender), then employees may have higher preference to work for in-group employers.

[2] This is analogous to employer discrimination where the employers may not only care about the gender or the race of the employee but also care about employee productivity on which they may not have perfect or complete information on.

[3] Becker (1971) conceptualized employee discrimination in the form of employees' dis-utility in working for specific group of employers, in the form of preferences.

[4] Employees could differ along many dimensions of job attributes that may care about (many being unobservable to the econometrician), which is very difficult to tease out from realized choice data. Data on realized job choices have their own advantages, especially to understand employer discrimination. This is because employers on average care about a consistent set of attributes in their employees.

ology (Blass, Lach & Manski (2010), Wiswall & Zafar (2018)) among job seeking students at a highly selective university in India. We present respondents with twenty hypothetical job choice scenarios and ask them to choose among three jobs in each choice scenario. We exogenously vary these jobs along realistic attributes (annual wages, flexible hours, manager's gender and manager quality) and cover the support of these attributes over the twenty scenarios.[5] Respondents are asked to choose their most preferred job and report wage compensations that make them indifferent between jobs, thus providing us a cardinal measure of their preferences. Extracting beliefs on manager quality through direct elicitation could be difficult, if we worry that respondents may be affected by social desirability bias. Hence within the hypothetical job choice scenarios we introduce a novel information experiment where we exogenously vary observability of a specific attribute on manager quality that employees care about—mentorship—motivated by our pilot surveys. Good mentors have large positive impacts on the progress of their mentees in academia (Blau, Currie, Croson & Ginther 2010)[6], in the military (Lyle & Smith 2014), in schools (Falk, Kosse & Pinger 2020)[7], and in jobs (Richard, Ismail, Bhuian & Taylor (2009), Bohannon (1985), Google's manager research[8]). Additionally, in presence of in-group mentoring, higher diversity at upper level jobs can help minorities break through the glass ceilings (Athey, Avery & Zemsky 2000).[9] In the same spirit, the American Economic Association 's official mentoring program CeMENT has senior women faculty mentoring junior women faculty.[10] We see this in our survey too where we find respondents care about mentorship more than competence and being pleasant to work with (Figure 5).

We quantify mentorship of a manger as a rating on a five point scale, motivated by the

[5]Each scenario could be thought as a market. Choice in a market provides individual demand in that market. Choices over multiple scenarios, by varying attributes over their support allows to trace out the individual demand curve. This generates a panel data of choices and compensations over the support of job attributes which provides identifying variation to estimate highly flexible models.

[6]The American Economic Association has a dedicated web-page on "Mentoring Reading Materials" as a part of the AEA initiative of Committee on the Status of Women in the Economics Profession (CSWEP). See https://www.aeaweb.org/about-aea/committees/cswep/mentoring/reading.

[7]Lyle & Smith (2014) using US military administrative data, provide empirical evidence on high quality mentorship to increase promotion probabilities for junior officers and its increasing with increased duration of mentoring.

[8]Under Google's manager research, their top-most criteria that determines what makes a manger great is "a good coach". For more details on Google's research on desirable manager attributes see https://rework.withgoogle.com/guides/managers-identify-what-makes-a-great-manager/steps/learn-about-googles-manager-research/

[9]In-group mentoring is associated with the increase in human capital due to mentoring being higher when the mentor and the mentee are of the same "type" (for example gender or ethnicity)

[10]For more details on CeMENT see https://www.aeaweb.org/about-aea/committees/cswep/programs/cement-mentoring-workshops/

new trend of rating of managers.[11] We describe it to the respondents as the average rating provided by the manager's current employees in an anonymous survey. The observability of the mentorship rating is varied in the following way. In the first ten scenarios individuals observe three jobs in each scenario with job attributes - annual wages, flexible hours and manager's name. Mentorship rating is mentioned but the data is shown to be not available. We call these first ten scenarios as incomplete scenarios throughout the rest of the paper, motivated by the fact that mentorship rating is not observable. In the last ten scenarios individuals observe three jobs in each scenario with all the above attributes. We call these last ten scenarios as complete scenarios.

The mechanism of individuals reporting their choices and compensations is as follows: in the incomplete scenarios, individuals observing all other attributes except for mentorship rating, form expectations on the rating. Thus their reported choices and compensations are a function of both their preferences and beliefs on how mentorship varies with the other characteristics. In the complete scenarios, since individuals observe all attributes, which as per instructions are the only attributes which differ across jobs within a given scenario, their reported choices and compensations are a function of only their preferences for those attributes. We instruct respondents in every scenario that jobs do not vary in attributes not mentioned in the survey and reported compensation only increases wages and changes nothing else about the job.[12] Thus variation in compensations within complete scenarios identify preferences. Variation in compensations between complete and incomplete scenarios because of the information experiment identify beliefs on mentorship. Observe that this information experiment is on all respondents allowing us to uncover the belief distributions.

We use our unique panel data of choices and compensations to estimate a structural model of job choice. We use this model to measure the preference and belief parameters in monetary value, as willingness to give up wages. Given our design, we get to observe the entire choice set — we not only observe what is chosen, but also observe what is not chosen. We also observe the compensations that make individuals indifferent between all choices in the choice set. This allows us to avoid any additional distributional assumptions on preference

[11] Comparably, Completed and Kunukunu are some of the notable start ups which provide ratings of managers and supervisors. They are provided in order for job seekers to have more information before starting or applying to new jobs. "Completed.com is a start-up with 25 million profiles of business people"- Bloomberg Law. *"Yelp for Bad Bosses: New Manager Rating Site Launched"* Vault and The Job Crowd are other notable mentions.

[12] This is one key advantage of using the hypothetical choice methodology over audit study field experiments.

and belief parameters. We use the variation in compensations and the information experiment to identify and estimate the distribution of preferences and beliefs of employees. Additionally, after respondents have completed all the job choice scenarios, our survey asks questions that directly elicit respondents' beliefs. We use this to corroborate the results obtained from the job choice model and discuss advantages of using the methods which we introduce in this paper.

We find that in the absence of information on manager mentorship — choices and compensations driven by *both preferences and beliefs* — employees are indifferent between male and female managers. However, with information on manager quality — choices and compensations driven by *only preferences* — employees prefer to work for female managers. On average, employees are willing to give up 1.7% of average annual wages to work for female managers. Hence, in absence of information on manager quality, employees believed female managers to be worse mentors. We estimate this statistical discrimination against female managers at 1.6% of average annual wages. We do not find evidence of negative beleifs on female manager mentorship among females majoring in Science. We corroborate this result of negative beliefs on female manager mentorship in an additional survey where we directly elicit employee beliefs, after respondents have completed the job choice and compensations in the twenty scenarios involving the information experiment. Our information experiment is at the employee level, i.e. the information treatment was given to every respondent. This allows us to uncover distributions of preferences and beliefs by estimating our job choice model for each individual. Around 62% of individuals prefer to work for female managers. Around 60% of individuals believed female managers to be worse mentors than male managers in the absence of information on manager mentorship.

Another important and novel contribution of our paper is quantifying demand for managers who are good mentors. We estimate individuals are willing to give up to 5.65% of average annual wages for one standard deviation increase in mentorship quality. We acknowledge that mentorship could be one among many attributes that employees care about in their managers. To push the extent to which mentorship flexibly captures manager quality, we extend our model by considering mentorship as a proxy for overall manager quality in the Appendix. This encapsulates other complicated cases which we discuss later in the paper. This additional flexibility trades-off with giving up point identification of preferences. In that setting only beliefs are point identified and results on preferences become suggestive.

Our paper contributes to multiple strands in the literature. The primary one is discrimination and the contribution is two-fold. First, to the best of our knowledge, this is the first paper known to us, to provide evidence on discrimination by job-seekers on manager's gender in choosing jobs, or in general, evidence on the labor supply side of discrimination.[13] The most important distinction with this literature is our survey puts individuals in hypothetical but realistic scenarios that they will face in the near future while choosing jobs in the labor market and further while switching teams within firms and between jobs.[14] Second, despite its vast expanse, the literature on discrimination deals with statistical and taste-based discrimination separately.[15] To the best of our knowledge, this paper is the first to allow for both beliefs driving statistical discrimination and preferences driving taste based discrimination and estimating them. The literature on statistical discrimination has focused mostly on tests of statistical discrimination (Altonji & Pierret (2001), Lange (2007), Agan & Starr (2018)).[16] Our framework and unique panel data on compensations allow us to not only test for statistical discrimination but also quantify it. We also explore and discuss weaker assumptions and identification of statistical discrimination in the Appendix, in more flexible models.

Our next contribution is methodological. The discrete choice literature using hypothetical choice methodology has focused on estimating preference parameters of models in the complete scenarios only i.e. choice scenarios where individuals have information on all attributes the researcher care for in any given object of choice, be it a product (Blass et al. 2010) or a job (Wiswall & Zafar 2018). Our methodological contribution is three-fold. First, we ask for wage compensations to derive indifferences between jobs, instead of asking for choice probabilities for each job. Asking for choice probabilities is standard in the literature, which provide a car-

[13]Employer side discrimination has been widely studied, starting with seminal papers by Becker (1957), who talks about prejudice and taste-based discrimination, and (Phelps 1972) and (Arrow 1972), who talk about statistical discrimination resulting from rational decision making conditional on information gaps. Recent work in the Experimental Economics literature has studied how individuals respond differently to criticism and advice from different genders (Abel (2019), Abel & Buchman (2020)).

[14]Other departures are in population of interest and in the model. Our sample are students who will be on the job market in a year thus belonging to the population of interest—job-seekers. Additionally these are job-seekers from a highly selective university, thus representative of the right tail of the ability distribution

[15] See Guryan & Charles (2013) for a discussion and review on types of discrimination. Alam (2019) discusses identification in models of both statistical discrimination and taste-based discrimination. For a detailed review on the literature of discrimination see Guryan & Charles (2013), Lang & Lehmann (2012) and Bertrand & Duflo (2017).

[16]Bohren, Imas & Rosenberg (2019) distinguish between discrimination resulting from correct beliefs (which they associate with statistical discrimination) and those from incorrect beliefs. We do not take that route because we do not have access to data on the population distribution of mentorship which could act as a benchmark to examine the extent to which beliefs are biased or not.

dinal measure of preference—the analogue of which is the compensation we have asked for. In the model section, we explain why depart from the literature, by delineating the advantages of using compensations over choice probabilities from a cognitive load on the respondents standpoint, and consistency with how the model works in conjunction to the instructions given to the respondents. Second, we develop and estimate a model of job choice which can be implemented with both the complete and the incomplete scenarios to identify both preferences and beliefs using our information experiment. The simplifying feature of our design and hence our models is that it exploits variations in compensating differentials which make individuals indifferent between jobs.[17] This allows us to directly estimate and interpret preference and belief parameters as measures of willingness to give up wages. Third, our question of interest involves identification of beliefs similar to Adams-Prassl & Andrew (2019). However we differ by indirectly eliciting beliefs with our novel information experiment, thus providing a novel method of belief elicitation in settings where the researcher may worry that individuals may not report truthfully because of social desirability bias. This requires incomplete choice scenarios and thus contributes to the design of surveys in the literature which has focused on complete scenarios. We also contribute to the growing literature of online surveys and experiments involving information treatments to study beliefs (Wiswall & Zafar (2015), Bordalo, Coffman, Gennaioli & Shleifer (2016), Alesina, Miano & Stantcheva (2018), Alesina & Stantcheva (2020), Kuziemko, Norton, Saez & Stantcheva (2015), Charité, Fisman & Kuziemko (2015) Boneva & Rauh (2018)).

Our paper also contributes to the literature on the effects of good mentors. The impact of high quality mentors on the progress of mentees is well studied (Blau, Currie, Croson & Ginther (2010), Lyle & Smith (2014), Falk, Kosse & Pinger (2020)) with consistent evidence on impacts being positive in academia, corporate sector, military and high schools. Our paper provides the first estimates on the demand for high quality mentors as a measure of wages job-seekers are willing to give up to work for managers who are better mentors.

The paper is organized as follows: Section 2 describes the administration of the survey,

[17] If we had followed the literature and asked for choice probabilities, we would need to use an additional estimation step for the relative variances of the error terms in the complete and incomplete scenarios. In that case we would have had to identify the relative variances in the error terms could by exploiting the variation between the variances of the compensations in the complete and incomplete scenarios. In our design, by using compensations, the value of a dollar remains a dollar, in both complete and incomplete scenarios, we do not need to do that extra step and can simply estimate preferences and beliefs in monetary terms by normalizing the preference parameter on wages to one in the employee utility function.

the survey instrument, the information experiment and the sample descriptives. We highlight key aspects of the design. It ends with discussion of patterns in the raw data and testing for in-group preferences of female employees using a difference-in-differences estimation strategy. Section 3 describes the job choice model spanning complete and incomplete scenarios and how their preferences and beliefs play a role in the choices they make and compensations they report. Using variation in the panel data on compensating differentials which make individuals indifferent across jobs, section 4 shows identification of the parameters of interest under assumptions which parallel the instructions given to survey respondents. Section 5 takes the model to the data and discusses estimation details and results. Section 6, presents heterogeneity results by sub-groups. Section 7 uses data from the expost direct belief elicitation part of our survey to corroborate the results on beliefs obtained from the job choice model. Section 8, presents the empirical distribution of beliefs and preferences constructed by estimating the model for each individual. This provides descriptives on proportions of individuals with different preferences and beliefs. Section 9 discusses implications of this paper and potential avenues of future research. Section 10 concludes.

2 Data

We collected data in the second week of April 2020 over 4 days, using an online survey administered to students of a highly selective public university in the state of West Bengal in India. Students eligible to participate in the survey were only those who were away from the job market by at most two years. In 2020, total number of enrolled students are 11,064. Out of them, 6,283 are enrolled in Undergraduate programs, 2,588 in Masters program and remaining 1,193 are enrolled in MPhil and PhD programs. Access to internet was not a concern for students studying in a premier University in Kolkata. However, we paid special attention in designing the survey to be mobile friendly, appreciating the fact that even though small but a significant proportion of the target population may not have access to a computer. Upon completion of the survey, participants were paid INR 500 ($\approx$ \$24)[18] through their preferred mode of online payment.[19]

[18] Purchasing power parity equivalent

[19] 1 USD in 2019 purchasing power parity is equivalent to 21.07 INR. Source: https://data.oecd.org/conversion/purchasing-power-parities-ppp.htm

See Appendix A.2 for further details on the administration and implementation of the online survey.

2.1 The survey and the information experiment

The survey started with instructions to the respondents and had three broad sections - (1) twenty choice and compensation *scenarios*, (2) direct belief elicitation and (3) demographic questions. The structure is schematically represented in Figure 1. Below we describe the design and purpose of each section in details.

2.1.1 Instructions

The first part of the survey included definitions of the attributes of jobs shown to individuals in each job as shown in Appendix Figure 1. Next individuals were given instructions to the survey as shown in Appendix Figure 2. There were two key instructions. First, individuals were to assume that the jobs do not vary in any attribute not mentioned in the scenario. Second, individuals when asked to report the minimum increase in wages in unchosen jobs required to make them indifferent to their chosen job, were to assume that this increase in wages do not change anything else about the job. After instructions individuals were shown two example scenarios to get them acquainted before starting the main survey.

2.1.2 Choice and compensation scenarios

Each *scenario* had two components - a choice among three hypothetical jobs, followed by asking for compensations that will make respondents indifferent between jobs. The job attributes we exogenously varied were annual wages, availability of flexible hours, manager's name and manager's mentorship rating. We varied the attributes subject to no job in each scenario was strictly dominating. A total of 20 such scenarios were administered in the survey to substantially vary the job attributes thus generating variation in reported compensations which we use to identify and estimate our parameters of interest.[20] We embedded the information experiment as follows — for every respondent, although mentorship rating was mentioned as an

[20] Ideally we would want to vary job attributes along the full range of attributes, and thus would end up asking a large number of questions. However that would come at a cognitive cost to the respondents. hence to strike a balance between cognitive load and substantial variation in job attributes we chose to administer 20 scenarios. This was mostly informed by our pilot surveys.

Figure 1: Schematic representation of survey flow

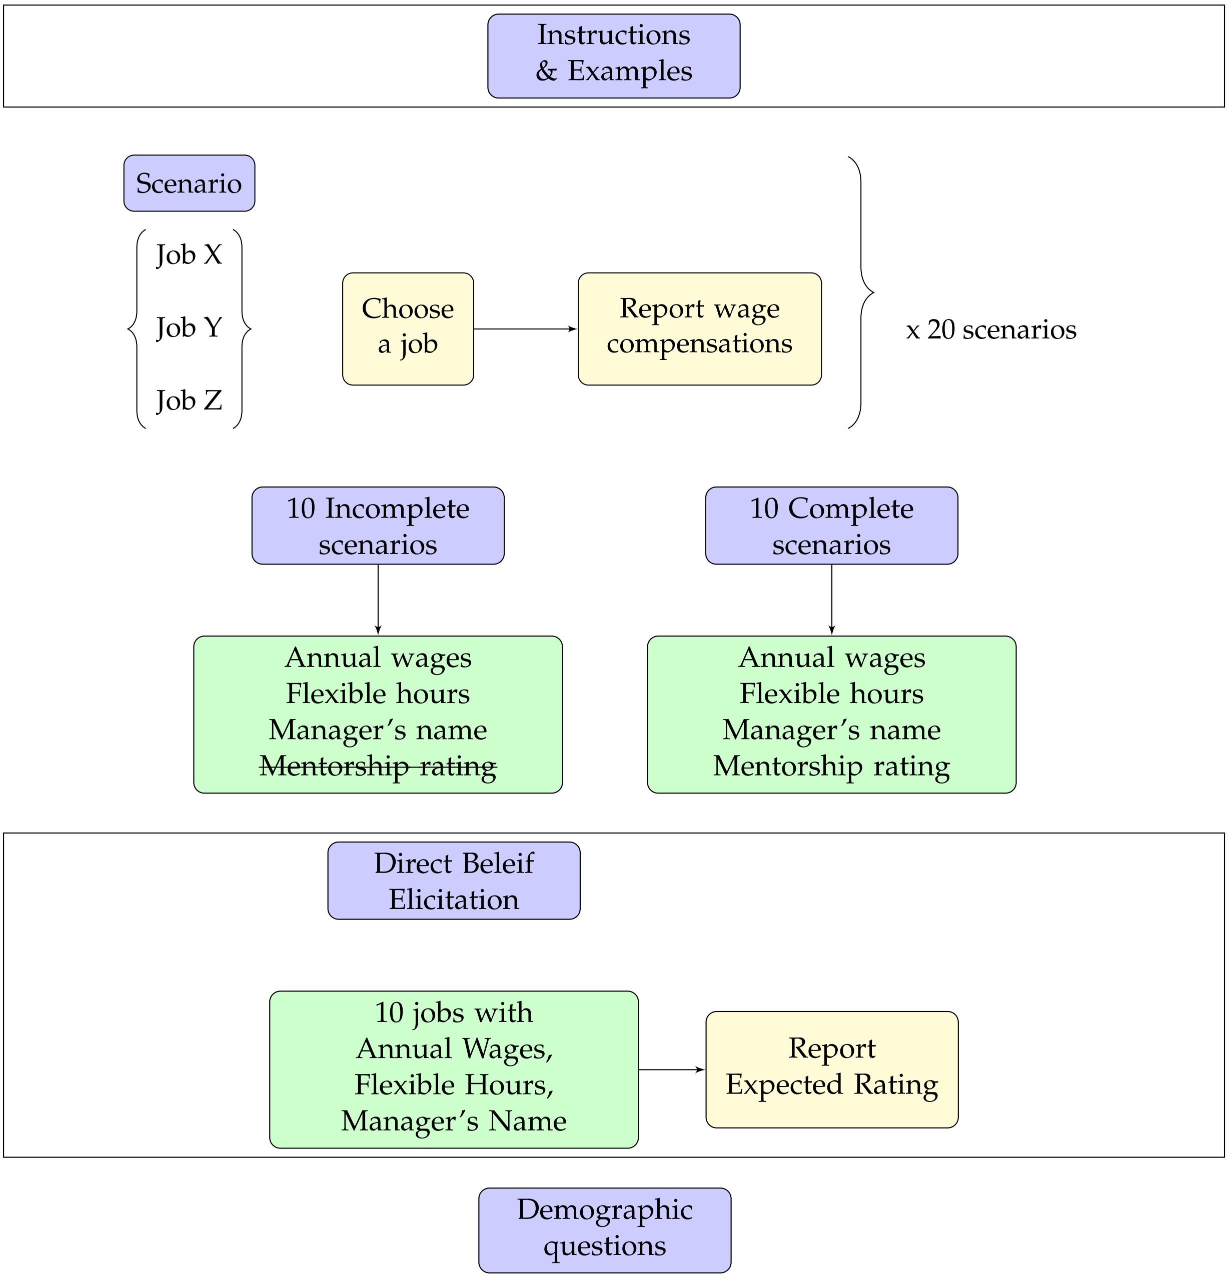

Notes: Every scenario consists of three different jobs X,Y and Z. Individuals choose their most preferred job. For the jobs that they did not choose, individuals are asked how much minimum increase in wages do they need in those jobs for them to choose them instead. There were 20 such scenarios. In the first 10 scenarios, individuals did not observe the mentorship rating of the manager however in the last 10 they did, along with other attributes.

attribute in all 20 scenarios, the first 10 scenarios (incomplete scenarios) did not have data on manager's mentorship while the last 10 scenarios (complete scenarios) did. The information experiment was designed in this way, so that the variation in compensations reported between complete and incomplete scenarios will identify beliefs on manager mentorship. Given the instructions, the reported choices and compensations in the incomplete scenarios are a function of both beliefs and preferences and those in complete scenarios are a function of only preferences. Examples of a complete and an incomplete scenario are shown in Table 1.

Table 1: Complete and Incomplete Scenario Examples

Incomplete Scenario example

Rating of each manager could be different, but the data is unavailable. Anything else that you don't see here, is the SAME across all jobs. Please select your most preferred job.

Attributes	**JOB X**	**JOB Y**	**JOB Z**
Annual Wages	6	6.6	6.4
Flexible hours	yes	no	yes
Manager	Anirban	Shrinita	Arup
Manager rating	N/A	N/A	N/A

Job X

Job Y

Job Z

Complete Scenario example

Anything else that you don't see here, is the SAME across all jobs. Manager rating is on a scale of 1-5. 1:Poor; 2:Fair; 3:Good; 4:Very good; 5:Excellent. Please select your most preferred job.

Attributes	**JOB X**	**JOB Y**	**JOB Z**
Annual Wages	8.2	8	8.6
Flexible hours	yes	no	yes
Manager	Mohan	Mohit	Mahima
Manager rating	3.70	4.00	3.15

Job X

Job Y

Job Z

Notes: Incomplete Scenario Example: Job X and Z have male managers. Job Y has a female manager
Complete Scenario Example: Job X and Y have male managers and Job Z has female manager.
Overall across 10 incomplete scenarios, five scenarios had two jobs with male managers and one job with a female managers and the rest five had the opposite. Same was the case for jobs under the 10 complete scenarios. Thus across all jobs half of them had male managers and half ahd female managers.

The compensation question was programmed in such a way that it would ask only about the jobs which were not chosen in the choice question. For each job that was not chosen, respondents were asked how much minimum increase in wages would they need in those jobs, so that they would choose them instead. This data provides us with wage compensations that make the respondents indifferent between jobs. Individuals could report this on a slider scale which ranged between 0 and 2 lakhs INR ($\approx$ 0 USD to 2857 USD). Individuals were told that if they needed more than 2 lakhs, they could max out the slider and another page would automatically appear asking them how much more would they need.

Suppose the respondent chose Job X. The survey was designed to recognize that and automate to ask compensations in Job Y and Z for them to be chosen instead.

You chose Job X.

If you were to negotiate your wage, how much MINIMUM INCREASE IN WAGES would you need, in each of the other jobs for you to choose it instead of Job Z?

The scale here ranges from 0 to 2 lakhs.

Attributes	**JOB X**	**JOB Y**	**JOB Z**
Annual Wages	8.2	8	8.6
Flexible hours	yes	no	yes
Manager	Mohan	Mohit	Mahima
Manager rating	3.70	4.00	3.15

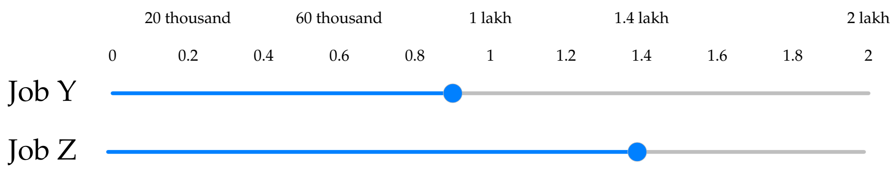

In the 20 scenarios there were 60 jobs, with half male and half female managers evenly distributed across the complete and the incomplete scenarios. Manager names used in the survey were common names directly indicative of gender. We paid careful attention that the names did not vary by social class. This circumvents any potential concern that could arise because of differences in perceived gender roles between social classes. Around half of them had flexible hours and the other half did not. Average annual wages were 7 lakh INR ($\approx$ \$ 39,444 in PPP)[21] and the average rating of managers in the complete scenarios was 3.41. Table

[21] This was the average wages of jobs that were offered to past cohorts. The variance in wages in the jobs we show

2 shows summary statistics of attributes shown over the 60 jobs and table 3 shows a balance of attributes across male and female managers.

Table 2: Summary statistics of Job attributes

Variable	**Mean**	**Std. Dev.**	**N**
Female Manager	0.5	0.504	60
Flexible hours	0.533	0.503	60
Annual wages	7.11	1.476	60
Rating	3.41	0.495	30

Table 3: Job Attributes Across Male and Female Managers

Attribute	Male manager	Female Manager	Difference	Std. Diff
Rating	3.373	3.447	0.073	0.103
	(0.519)	(0.485)	(0.183)	
Flexible hours	0.567	0.500	-0.067	-0.093
	(0.504)	(0.509)	(0.131)	
Annual wages	7.080	7.140	0.060	0.029
	(1.529)	(1.445)	(0.384)	

2.1.3 Direct Belief Elicitation

This section of the survey directly elicited beliefs on manager mentorship. This was designed to corroborate results from taking the data of the choice and compensation scenarios to the job choice model. In this section, individuals were presented with 10 jobs with manager's name, annual wages and flexible hours. Individuals were given a slider scale from zero to five and were asked to report their expected ratings in each job. This is depicted in Appendix Figure 3.

2.1.4 Demographic questions

In this section we asked demographic questions to the respondents on their broad are of study (Arts, Science or Engineering), family income, parental education and occupation. Then individuals were asked questions which were specifically designed to infer whether individuals followed instructions or not. The survey ended with mode of online payment.

were not too high. This mitigates any concerns that some jobs could be interpreted as entry-level jobs and some as senior level jobs. The attributes were varied such that no job in any scenario turned out to be strictly dominant.

2.2 Key Highlights of the Design

We have already explained how the design of the information experiment in conjunction with the instructions given to the respondents allow us to isolate beliefs and preferences of employees. In this section we explain some of the other selected highlights of the design which help us to identify and flexibly estimate the parameters of our interest.

1. Observe the entire choice set and compensating differentials.

   - This is one of the key advantages of our design. We observe which jobs are chosen and which jobs are not chosen. Additionally we also observe the compensations that make individuals indifferent between all choices in the choice set. This allows us to avoid making any distributional assumptions on the preference or the belief parameters.

2. The information treatment is given to every individual.

   - For each individual we observe a sequence of choices made under incomplete scenarios and then a sequence of choices made over complete scenarios. This allows to uncover distributions of preferences and beliefs.

3. Data on compensating differentials which make individuals indifferent between jobs.

   - This allows us to directly estimate and interpret parameters as measure of willingness to pay or willingness to give up wages.[22]

4. Multiple scenarios of job choice and compensations for each individual.

   - We do this to cover as much of the support of the job attributes that we vary and their possible combinations. It is not feasible to make individuals choose among a large number jobs within each scenario because it would induce very high cognitive load. Hence we do this over a panel of scenarios making them choose and provide compensations over three jobs per scenarios. This generates a job demand panel data of choices and compensations over jobs which considerably vary over their support. Also, this provides us with sufficient variation to identify and estimate highly flexible models.

5. Only first names of managers

[22] This is also possible with data on choice probabilities, but requires an additional step to 'convert' the estimates to willingness to pay measures. See Wiswall & Zafar (2018).

- We used first names of managers making sure that the names do not vary by social class. In the Indian context last names can identify castes and religion. Since we only wanted to vary the dimension on gender we did not show any last names. This is unlike in the USA wherein first names could be associated with a race and a gender (Bertrand & Mullainathan 2004).

2.3 Selection and Description of Sample

The total number of participants in our survey was 604 out of which 591 completed the survey. Out of these we dropped 11 sample points who could either not be verified as students, or who completed the survey in less than 15 minutes, or both. The first percentile of duration to survey completion was at 13.89 minutes. The median time to survey completion was 51.37 minutes. Our final sample size used in all following estimation is 580.

Table 4 reports sample descriptives. 41.72% of our sample consists of female students and the remaining 58.28% were male students.[23] 44% were enrolled in a department in the Arts faculty, 33.28% in Engineering and the remaining 22.07% were in Science. Female students were predominantly from Arts (67%). Males were predominantly from Engineering (49%) and Science (33%).

2.4 Patterns in the raw data

In this subsection we explore how individuals choices and reported compensations varied across the complete and the incomplete scenarios between jobs with male and female managers. This section is important before we give structure to rationalize the data in a framework involving preferences and beliefs. Table 5 reports the percent of jobs chosen which had male managers and the percent of jobs chosen with female managers, in both the complete and the incomplete scenarios.

The first observation is that the percentage of jobs chosen with managers of both genders does not differ that much between male and female respondents. Furthermore we notice that in the absence of information on manager's mentorship (incomplete scenarios) the percent of jobs chosen with female managers were not that different from those with a male manager.

[23]In the survey we asked for the biological sex of the individual. We did not ask for the gender with which students would like to associate themselves, other than their biological sex.

Table 4: Sample Demographics

	Gender		
	Male	Female	Total
580 respondents	58.3	41.7	100
Area of Study:			
Arts	25.1	71.9	44.7
Engineering	49.4	10.7	33.3
Science	25.4	17.4	22.1
Family Income:			
Less than 2 lakhs (Less than $ 9,492)	37.6	24.4	32.1
2 lakhs to 5 lakhs ($9,492 to $23,730)	26.9	24.0	25.7
5 lakhs to 10 lakhs ($23,730 to $47,460)	21.3	30.2	25.0
10 to 20 lakhs ($47,460 to $ 94,921)	11.5	15.3	13.1
Above 20 lakhs (Above $ 94,921)	2.7	6.2	4.1
Mother's Education:			
Below High School	17.8	11.2	15.0
High School	32.8	18.6	26.9
Bachelor's	38.5	46.3	41.7
Master's	8.0	16.9	11.7
Above Master's	3.0	7.0	4.7
Father's Education:			
Below High School	9.2	7.4	8.4
High School	21.3	11.2	17.1
Bachelor's	51.5	55.4	53.1
Master's	13.9	17.8	15.5
Above Master's	4.1	8.3	5.9
Mother's Occupation:			
Government	10.9	15.3	12.8
Home-Maker	70.1	63.2	67.2
Not Applicable	4.4	3.3	4.0
Private Sector	4.7	8.7	6.4
Self-employed	9.8	9.5	9.7
Father's Occupation:			
Government	30.8	33.1	31.7
Home-Maker	3.3	0.8	2.2
Not Applicable	16.3	11.2	14.1
Private Sector	16.9	20.7	18.4
Self-employed	32.8	34.3	33.4

Notes: All variables are categorical. Numbers represent percentages. Variables on income categories are in INR and have their corresponding purchasing power parity adjusted equivalent USD below wach category

Table 5: Percent of chosen jobs with male and female managers, by scenarios

Respondent	*All*		*Female*		*Male*	
Manager / *Scenario*	Male	Female	Male	Female	Male	Female
Incomplete	48.2	51.8	48.1	52.9	48.3	51.7
Complete	38.9	61.1	38.8	61.2	38.9	61.1

However, in the complete scenarios, upon information revelation of manager mentorship, we observe that the percentage of chosen jobs with female managers is 61.1% which is around 20 percentage points higher than those with male managers.

Table 6: Average compensations in unchosen jobs, by gender of manager across scenarios

	Male manager	Female Manager	Difference	St Diff
Incomplete scenarios	0.957	1.020	-0.063***	-0.046
	(0.965)	(0.998)	(0.018)	
Complete scenarios	1.100	1.038	0.061***	0.042
	(1.058)	(0.977)	(0.019)	

Notes: Units are in 1 lakh (hundred thousand) INR

Table 6 reports the average compensations reported in unchosen jobs with male and female managers, and the difference between them along with the associated standard error and standardized difference. The table reports these numbers separately for the complete and the incomplete scenarios. In the absence of information on manager mentorship, we observe that individuals on average report higher compensations to choose jobs with female managers than jobs with male managers, by 6.3 thousand INR ($\approx$ \$ 300). However this result flips in the complete scenarios. Individuals when provided information on the mentorship rating, on average demand 6.1 thousand INR ($\approx$ \$ 290) more in unchosen jobs with male managers. Both differences are statistically significant at the 99% level. We should maintain caution in interpreting these numbers because this compares compensations among the set of jobs that they did not choose. Nevertheless, these numbers provide useful information on the set of unchosen jobs. In a world where individuals would have ranked jobs, these numbers suggest that in the complete scenarios, jobs with female managers would have been more likely to be ranked 2 but in the incomplete scenarios, would have been ranked last. The more informative way to understand the compensation data would be to compare attributes of the chosen and unchosen jobs, given the compensations. This is what we do in the job choice model.

Before we delve into the model, we provide evidence for a natural question that arises in these contexts, on in-group preferences- in particular whether female respondents are more likely to choose jobs with female managers.

2.5 In-group preferences of females

In this section we discuss whether females relative to males, are more likely to choose jobs with female managers. This can be answered with data from the complete scenarios using a difference in differences estimation strategy. We do not include the incomplete scenario data here, because we do not want to deal with omitted variable bias that will arise from how individuals form beliefs on mentorship which is unobserved to them in the incomplete scenarios.

Individuals are indexed by $i = 1, .., N$ and jobs by $j = 1, .., J$. Define $Choice_{ijs}$ as a binary variable taking value 1 if individual i in scenario s chose job j and 0 otherwise.

$$\begin{aligned} Choice_{ijs} = \delta_0 + \delta_1 \underbrace{\mathbb{I}(g_i = f)}_{\text{female}} \overbrace{\mathbb{I}(MG_{j(s)} = f)}^{\text{female manager}} + \delta_2 \underbrace{\mathbb{I}(g_i = f)}_{\text{female}} + \delta_3 \overbrace{\mathbb{I}(MG_{j(s)} = f)}^{\text{female manager}} \\ + Attributes'_{j(s)}\gamma_1 + Demographics'_i\gamma_2 + \lambda_s + e_{ijs} \end{aligned} \tag{1}$$

Respondent i's gender is denoted by g_i and gender of the manager in job j of scenario s is denoted by $MG_{j(s)}$. $Attributes_{j(s)}$ is the vector of job attributes associated with job j in scenario s, other than gender of the manager i.e., wages, flexible hours and mentorship of the manager. $Demographics_i$ is a vector of individual level demographics described in the previous section. The specification controls for scenario fixed effects (λ_s). This is necessary because a choice is made within a scenario thus estimation needs to use the variation in attributes within each scenario. We estimate this difference-in-differences equation with a logit model. Table 7 shows the marginal effect estimates of equation (1). We find no evidence that females are more likely to choose jobs with female managers than males. Observe that we see this in the raw data as well, where we find no difference in choices or in compensations across male and female respondents in both the complete and the incomplete scenarios. We will see this again when we present estimates of preferences and beliefs in the job choice model, by the gender of the respondent. As one would expect, higher wages, availability of flexible hours and better mentors are associated with higher likelihood of the jobs being chosen.

We now move on to describe and estimate a job choice model to unwrap this evidence of jobs with female managers being chosen more on average, through the lens of preferences of individuals and how beliefs operate in the absence of information on the manager's rating.

Table 7: Difference-in-Differences Estimates from the Complete Scenarios Data

VARIABLES	Margins
Female	-0.001
	(0.010)
Female manager	0.081***
	(0.009)
Female X Female Manager	0.001
	(0.014)
Annual wages	0.347***
	(0.019)
Rating	0.484***
	(0.009)
Flexible hours	0.262***
	(0.009)
Scenarios FE	yes
Observations	17,400

Notes: The estimates show the marginal effects of each of the attributes in a difference and difference specification. Standard errors bootstrapped at the individual level with 1000 replications. 30 data points/individual for 580 individuals in the complete scenarios

3 Model

In order to understand the mechanisms behind the choices made and compensations reported, we use a simple job choice model. We explain how the model works in the complete scenarios and then in the incomplete scenarios, through individual preferences and beliefs.

Individuals are indexed by $i \in \{1, \ldots, N\}$ and jobs are indexed by $j \in \{1, \ldots, J\}$. Let X_j denote the K-dimensional vector of attributes of job j, over which individuals have preferences. The utility of an individual i, with preference parameter vector $\beta_i \in \mathbb{R}^K$ from job j is given by

$$U_{ij} = u_i(X_j; \beta_i) + \epsilon_{ij}$$

ϵ_{ij} are all unobservables that affect the utility of individual i from job j.

To keep things simple, we use a linearly separable model. As we will explain in the identification section, this functional form has no bearing on the identification of the parameter or its distribution, except for additive separability of the error terms ϵ_{ij}.[24] This is because we have panel data of choices and compensations for 20 scenarios with large variation in job attributes,

[24]In the Appendix, we explore a more generic utility function and show identification of parameters of interest.

which can identify parameters of more complex models.

Individuals have preferences over the set of attributes $X = \{G, W, H, R\}$. Here G, W, H, R denote the indicator for a male manager, annual wages, availability of flexible hours and mentorship rating of the manager respectively.[25] In the complete scenarios respondents observe X for each job. In the incomplete scenario respondents observe $\tilde{X}$ where $\tilde{X} \equiv X \setminus R$. For ease of exposition, we will first explain how the individual makes decisions in the complete scenarios followed by the incomplete scenarios. Utility of an individual i, with preference parameter vector β_i from job j with attributes X_j is given by

$$U_{ij} = X_j' \beta_i + \epsilon_{ij} \tag{2}$$

3.1 Complete Scenarios

In the complete scenarios individuals observe all attributes in set X for each job. The expected utility of individual i from job j conditional on its observable attributes in the complete scenarios is given by

$$\mathbb{E}_i[U_{ij}|X_j] = \sum_{x \in X} \beta_i^x x_j + \mathbb{E}_i(\epsilon_{ij}|X_j) \tag{3}$$

We assume that all individuals know their preferences and hence do not take expectations over them. As explained above, this draws a clear parallel with asking for choices instead of choice probabilities.

3.2 Incomplete scenarios:

In the incomplete scenarios rating of the manager is mentioned but the data is shown to be unavailable to the individuals. Hence, individuals form expectations over it in reporting their choices and compensations. Denote the set of observable attributes in job j as $\tilde{X}_j \equiv X_j \setminus \{R_j\}$ in the incomplete scenarios. The expected utility of individual i from job j conditional on its

[25]In this model an individual directly cares about the mentorship of the manager. We acknowledge that this may not be the one single manager attribute that individuals care about in their managers, for which in the Appendix we write and discuss identification of a more generic model in which the mentorship rating of the manager is taken as a signal of overall manager quality.

observable attributes $\tilde{X}_j$ in incomplete scenarios is given by

$$\mathbb{E}_i[U_{ij}|\tilde{X}_j] = \sum_{x \in \tilde{X}} \beta_i^x x_j + \mathbb{E}_i(\beta_i^R R_j + \epsilon_{ij}|\tilde{X}_j) \tag{4}$$

To proceed we parameterize beliefs conditional on observables. Again to keep things simple we will specify the process of beliefs linearly. We revisit this specification later and consider alternatives when we interpret how beliefs about manager's ratings differ by the manager's gender.

$$\mathbb{E}_i(R_j|\tilde{X}_j) = \sum_{x \in \tilde{X}} \alpha_i^x x_j + \mathbb{E}_i(\eta_{ij}|\tilde{X}_j) \tag{5}$$

Observe that $\alpha_i^G = \mathbb{E}_i(R_j|G_j = male, W_j, H_j) - \mathbb{E}_i(R_j|G_j = female, W_j, H_j)$ represents how much on average individual i believes a male manager's mentorship rating differs from that of a female manager. Simplifying, the expected utilities in the incomplete scenarios are given by

$$\mathbb{E}_i[U_{ij}|\tilde{X}_j] = \sum_{x \in \tilde{X}} (\beta_i^x + \beta_i^R \alpha_i^x) x_j + \mathbb{E}_i(\epsilon_{ij} + \beta_i^R \eta_{ij}|\tilde{X}_j) \tag{6}$$

3.2.1 Discussion on model: Choices and Compensations versus Choice probabilities

The literature in hypothetical choice methodology has focused primarily on asking for the respondent's probability of choosing each alternative within each scenario. Asking for probabilities make sense when individuals are asked about future choices with revelation of relevant information in between. The difficulty is that interpretation of the role of unobservables in this setting, could be made in one of two possible ways. Asking for probabilities imply resolution of resolvable uncertainty (Blass, Lach & Manski 2010). The uncertainty could arise from individuals learning about their preferences over time (Delavande & Manski 2015).[26] But it is at odds with how the model operates, where it is taken as given that individuals know their preferences and the unobservables are interpreted as preference shocks. Thus it is not clear what individuals are understanding when they are being instructed to consider the attributes mentioned in the survey to be the only attributes along which the jobs vary.

To maintain clarity on the role of unobservables, we ask for one job choice among three options in each scenario, followed by a compensation to make individuals indifferent between

[26] In Delavande & Manski (2015) voting behavior today could potentially be different from voting behavior in future, if the individual learns about their preferences over time till prior to the time of vote.

jobs. This provides a clear parallel between the instructions provided and how the unobservables are modeled. Given the instructions provided in our setting - unobservables do not vary across jobs within scenarios, conditional on observables. Additionally asking for choices is more straightforward than probabilities and thus more appealing with respect to reducing cognitive load on respondents.[27]

4 Identification

In this section we show how our experimental panel data, on choices and compensations identify preference and belief parameters of our model of job choice by exploiting variation in reported compensations within and between the complete and the incomplete scenarios.[28] As shown in Appendix Figure 2, individuals were instructed to assume that,

> **Assumption 1:** All attributes which are not mentioned in the survey are the same for all jobs.
>
> **Assumption 2:** Reported compensation only increases wages and changes nothing else about the job.

Observe that instruction 1 is an assumption between jobs, while instruction 2 is for within jobs. The purpose of these instructions was to ensure that there is no selection on attributes which are not mentioned in the survey. Wiswall & Zafar (2018) delineate the importance of this assumption (1) in contrast to audit and correspondence studies, where there is no control on employers to not assume anything else conditional on observables.[29] For both scenarios, assumption (2) implies that the compensation only increases the wage and does not change the conditional expectation of the unobservables. Note that for the incomplete scenarios, it applies to the conditional expectation of the rating of the manager as well.[30] We use the data

[27]We administered pilots where we asked for probabilities of choosing each job. Debriefs revealed individuals spent longer times in understanding or calculating probabilities. Time data revealed individuals spent a long time on the pages which included a primer in probabilities and probability examples, thus adding to our concern of cognitive load.

[28]In the Appendix we write down a more flexible model where the rating variable is used as a signal for overall manager quality and show identification in that setting.

[29]For more on audit studies in detecting discrimination see Heckman (1998) and related papers within

[30]Note that for the incomplete scenarios, the rating variable was mentioned but there were no rating available for each manager. In order to make sure that neither do we prime individuals to think that rating is indeed different across managers, nor do we make them assume that rating is the same across all managers we used the following wording : *"Rating of each manager could be different, but the data is not available*, as shown in Table 1.

on compensations to equate expected utilities in the complete and the incomplete scenarios. These two assumptions which form clear parallel to the instructions given to the respondents, form the basis of our identification.

4.1 Preferences

In this section we show identification of the preference parameters β^x for all $x \in \{G, H, R\}$. The parameter of interest is β^G which is the preference on male managers. We use assumptions (1) and (2) to identify preferences using variation in compensations within the complete scenarios.

Implication of Assumption (1):

The instructions imply that unobservables across different jobs in conditional expectations are the same within each scenario. For every individual i and every job $j \neq k$ within each complete scenario,

$$\mathbb{E}_i(\epsilon_{ij}|X_j) = \mathbb{E}_i(\epsilon_{ik}|X_k)$$

Implication of Assumption (2):

In the complete scenarios, for each job k, individuals observe the vector of attributes $X_k = \{G_k, H_k, W_k, R_k\}$. Suppose individual i chose job j and then provided a compensation of Δ_{ijk} in order to choose job k instead. The instructions imply that for unchosen jobs like job k,

$$\mathbb{E}_i(\epsilon_{ik}|X_k, \Delta_{ijk}) = \mathbb{E}_i(\epsilon_{ik}|X_k) \tag{7}$$

Given the above, equating expected utilities between job j and job k with provided compensation of Δ_{ijk} and normalizing $\beta_i^W = 1$, we have,

$$\begin{aligned} \mathbb{E}_i(U_{ij}|X_j) &= \mathbb{E}_i(U_{ik}|X_k, \Delta_{ijk}) \\ \Delta_{ijk} &= \sum_{x \in X} \beta_i^x (x_j - x_k) \end{aligned} \tag{8}$$

Thus preference parameters β_i^x are identified using variation in reported compensations in the complete scenarios under assumptions (1) and (2) for all $x \in X$

4.2 Beliefs

Now we turn to the incomplete scenarios in conjunction with the complete scenarios and show identification of the belief parameter α, exploiting the variation between the reported compensations between the complete and the incomplete scenarios.

Implication of Assumption (1):

The instruction implies that unobservables affecting utilities and beliefs, across different jobs in conditional expectations are the same within each scenario. For every individual i and every job $j \neq k$ within each incomplete scenario,

$$\begin{aligned}\mathbb{E}_i(\epsilon_{ij}|\tilde{X}_j) &= \mathbb{E}_i(\epsilon_{ik}|\tilde{X}_k)\\ \mathbb{E}_i(\eta_{ij}|\tilde{X}_j) &= \mathbb{E}_i(\eta_{ik}|\tilde{X}_k)\end{aligned}$$

Implication of Assumption (2):

In the incomplete scenarios, for each job k, individuals observe the vector of attributes $\tilde{X}_k = \{G_k, H_k, W_k\}$. Suppose individual i chose job j and provided a compensation of $\tilde{\Delta}_{ijk}$ in order to choose job k instead. All that the compensation does is that it increases the wages in job k by $\tilde{\Delta}_{ijk}$. The implication of Assumption (2) is that it has no effect on the conditional expectation of managers' ratings or on the conditional expectation of the unobservables affecting utility. That is,

$$\begin{aligned}&\mathbb{E}_i(R_k|\tilde{X}_k, \tilde{\Delta}_{ijk}) = \mathbb{E}_i(R_k|\tilde{X}_k)\\ \text{and}\quad &\mathbb{E}_i(\epsilon_{ik}|\tilde{X}_k, \tilde{\Delta}_{ijk}) = \mathbb{E}_i(\epsilon_{ik}|\tilde{X}_k)\end{aligned} \tag{9}$$

Thus the expected utility from job k taking into account the compensation of $\tilde{\Delta}_{ijk}$ is :

$$\begin{aligned}&\mathbb{E}_i(U_{ik}|\tilde{X}_k, \tilde{\Delta}_{ijk})\\ &= \beta_i^G G_k + \beta_i^H H_k + \beta_i^W (W_k + \tilde{\Delta}_{ijk}) + \beta_i^R \mathbb{E}_i(R_k|\tilde{X}_k, \tilde{\Delta}_{ijk}) + \mathbb{E}_i(\epsilon_{ik}|\tilde{X}_k, \tilde{\Delta}_{ijk})\\ &= \beta_i^G G_k + \beta_i^H H_k + \beta_i^W (W_k + \Delta_{ijk}) + \beta_i^R \mathbb{E}_i(R_k|\tilde{X}_k) + \mathbb{E}_i(\epsilon_{ik}|\tilde{X}_k)\\ &= \sum_{x \in \tilde{X}} (\beta_i^x + \beta_i^R \alpha_i^x) x_k + \beta_i^W \tilde{\Delta}_{ijk} + \mathbb{E}_i(\epsilon_{ik} + \beta_i^R \eta_{ik}|\tilde{X}_k)\end{aligned} \tag{10}$$

given that beliefs about ratings are modeled as in (5).

Individuals are assumed to know their preference parameters and hence do not take expectations over them. Equating expected utilities between job j and job k with provided compensation of $\tilde{\Delta}_{ijk}$ and as before normalizing $\beta_i^W = 1$ under A1, we have

$$\begin{aligned} \mathbb{E}_i(U_{ij}|\tilde{X}_j) &= \mathbb{E}_i(U_{ik}|\tilde{X}_k, \tilde{\Delta}_{ijk}) \\ \tilde{\Delta}_{ijk} &= \sum_{x \in \tilde{X}} \underbrace{(\beta_i^x + \beta_i^R \alpha_i^x)}_{\tilde{\beta}_i^x}(x_j - x_k) \end{aligned} \tag{11}$$

Now identification of α is straightforward. We know that, for each attribute $x \in \tilde{X}$

$$\tilde{\beta}_i^x = \beta_i^x + \beta_i^R \alpha_i^x \tag{12}$$

Observe that $\tilde{\beta}^x$ is comprised of two terms: the preference parameter for attribute x and how much x affects the belief about manager's rating, given that the individual cares about manager's rating.

Thus given identification of β_i^x and $\tilde{\beta}_i^x$, we have for all $x \in \tilde{X}$

$$\alpha_i^x = \frac{\tilde{\beta}_i^x - \beta_i^x}{\beta_i^R} \tag{13}$$

Thus α_i^x are identified $\forall x \in \tilde{X}$.

A key observation on the intuition of two sets of equations of compensations in the complete and in the incomplete scenarios is how do the differences in attributes between jobs j and k map into the compensating differentials of each individual i to make them indifferent between job j and k. In the complete scenarios, the compensations are a function of how do jobs j and k vary in their attributes, weighted by how much individual i cares about each attribute. Contrasting it to the incomplete scenarios, they are weighted by how much individuals care about each attribute and how much they belief each attribute is correlated with mentorship and how it differs between jobs j and k in each of the incomplete scenarios.

5 Estimation

We use variation in the reported compensations to estimate the preference and belief parameters of our model. The compensations are reported with two different but independent mea-

Figure 2: Reported compensations in unchosen jobs

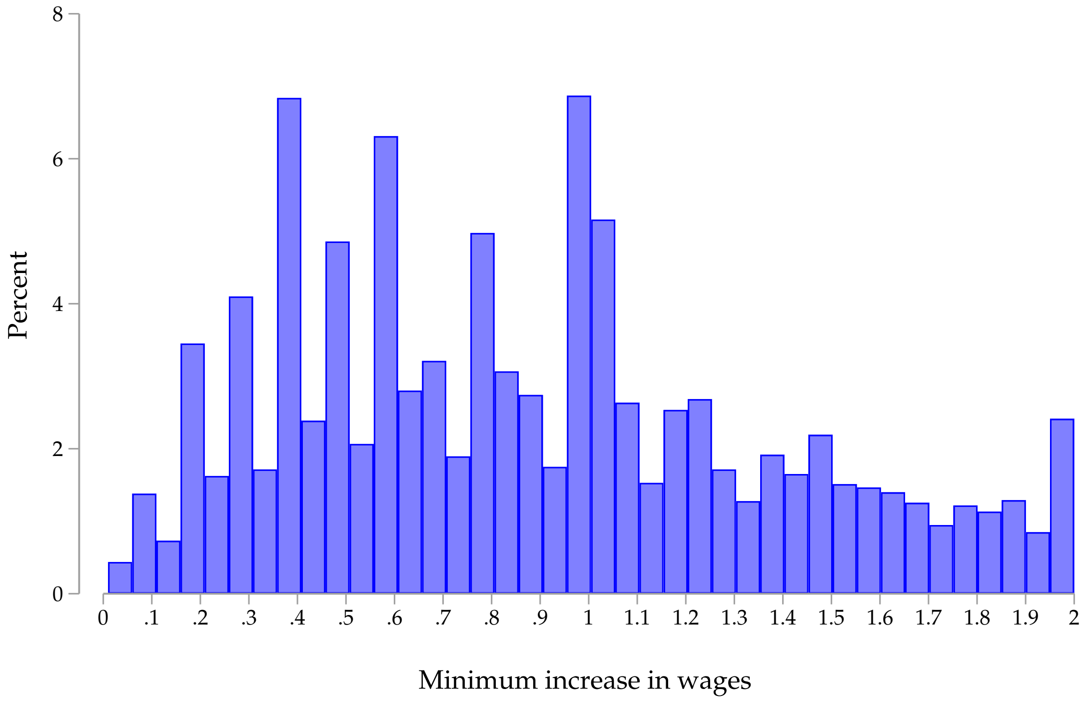

Notes: The increase in wages are in the units of 1 lakh (hundred thousand) INR. The figure is plotted for values only in between 0 and 2 lakhs.

surement errors. One measurement error results from reporting compensations in multiples of five as shown in Figure 2. This is also observed in the surveys asking for choice probabilities (Blass, Lach & Manski 2010). The second type of measurement error arises from the design of the survey. Individuals had a slider to report compensations and sliding the slider could cause some random measurement error, even though individuals could see the exact amount as they slid the slider. Both types of measurement errors are classical in nature and will only inflate standard errors. Consistency of the estimates are not affected. Denote Δ^*_{ijk} and $\tilde{\Delta}^*_{ijk}$ as the latent compensations and e_i and $\tilde{e}_i$ as the composite classical measurement errors in the complete and the incomplete scenarios respectively. Thus we have the following set of estimating equations for each individual i and each pair of jobs j and k within every scenarios.

$$\begin{aligned} \Delta^*_{ijk} &= \sum_{x \in X} \beta^x_i (x_j - x_k) + e_i \\ \tilde{\Delta}^*_{ijk} &= \sum_{x \in \tilde{X}} (\beta^x_i + \beta^R_i \alpha^x_i)(x_j - x_k) + \tilde{e}_i \end{aligned} \tag{14}$$

With classical measurement errors, we have,

$$\begin{aligned}\mathbb{E}[\Delta^*_{ijk}|X_j, X_k] &= \sum_{x \in X} \beta_i^x (x_j - x_k) = (X_j - X_k)'\beta_i \\ \mathbb{E}[\tilde{\Delta}^*_{ijk}|\tilde{X}_j, \tilde{X}_k] &= \sum_{x \in \tilde{X}} (\beta_i^x + \beta_i^R \alpha_i^x)(x_j - x_k) = (\tilde{X}_j - \tilde{X}_k)'\tilde{\beta}_i\end{aligned} \tag{15}$$

We jointly estimate the above system of equations (14) using constrained least squares where we normalize the preference parameter on wages to one ($\beta_i^W = 1$). By construction, the estimating equations have no constant since they are derived by equating utility functions. We use block bootstrap at the respondent level to allow for arbitrary correlation among responses within each respondent. We describes the details of the joint estimation and the bootstrap algorithm in the Appendix A.3

5.1 Model Estimates

In this section we discuss the estimates of the preference and belief parameters of the job choice model described in the previous section. In the first subsection we discuss estimates for preference parameters focusing on the preference on the gender of the manager. Next we discuss preference and thus demand for good mentors. We discuss estimates of the belief parameters focusing on beliefs on mentorship by the gender of the manager in the final subsection. This order of discussion is natural because only after having an understanding of how much individuals care about the gender and the mentorship of the manager, it is useful to discuss beliefs on mentorship by gender of the manager. Since the utility parameter on wages is normalized to 1, and wages are in hundred thousand (lakh) INR, the estimates should be interpreted as valuations of each attribute in hundred thousand INR.[31] For better interpretability we also present these estimates by converting them to percent of average annual wages.

5.1.1 Preference to work for female managers

Table 8 shows estimates of equation system (14) representing indifferences in the complete and incomplete scenarios jointly exploiting variation in reported compensations across jobs with exogenously varying attributes. We first discuss the complete scenarios under which are the

[31] Alongside we present the purchasing power parity equivalent in USD.

pure preference parameter estimates on manager's gender, flexible hours and average rating of the manager's mentorship.

Table 8: Complete & Incomplete Scenarios: Jointly Estimated

Incomplete Scenarios			Complete Scenarios		
Parameters	in 10^5 INR	% of wages	*Parameters*	in 10^5 INR	% of wages
$\tilde{\beta}^G = \beta^G + \beta^R \alpha^G$	-0.007	-0.01%	β^G (Male Manager)	-0.118***	-1.7%
	(0.012)			(0.019)	
$\tilde{\beta}^H = \beta^H + \beta^R \alpha^H$	1.136***	16.1%	β^H (Flexible Hours)	0.776***	11.1%
	(0.068)			(0.027)	
$\tilde{\beta}^W = \beta^W + \beta^R \alpha^W$	1.132***	16%	β^W (Annual wages)	1	
	(0.059)				
			β^R (Mentorship)	0.793***	11.3%
				(0.029)	
Observations	11,600			11,600	

Notes: The table shows estimates from estimating the system of equations (14). Estimates are represented in two sets of units - the first is in hundred thousand INR and the second converts those units as a percent of average annual wages. Standard errors are computed using block bootstrapped at the individual level with 1000 replications. β^W is normalized to 1. ***, **, * denote statistical significance at 1, 5, and 10%, respectively.

The most striking result is the evidence of strong preference to work for female managers as shown in the panel of complete scenarios. The preference parameter for male managers (β^G) is negative and statistically significant. On average, individuals are willing to give up 12 thousand INR ($\approx$ \$570) to work for a female manager. This value corresponds to 1.7% of average annual wages. The 95% confidence interval of this value as a percent of average annual wages is in between 1.3% to 2.2%.[32]

For the incomplete scenarios, the 'biased' estimate of preference for male managers ($\tilde{\beta}^G$) is statistically indistinguishable from zero. The difference in the estimates across the complete (β^G) and the incomplete scenarios ($\tilde{\beta}^G$), provides evidence that in the absence of information on the mentorship quality of managers, individuals believe that male managers are better mentors. The results suggests that beliefs and preferences are operating in opposite directions to generate such estimates. This difference in the estimates between the complete and incomplete scenarios is evidence of statistical discrimination against female managers if and only if individuals prefer to work for managers who are better mentors. We now turn to discussing preferences on the mentorship attribute. As explained earlier in the identification section, if

[32]Bias corrected bootstrap confidence intervals.

individuals do not have preference on manager's mentorship then beliefs are fundamentally unidentified and as such this difference cannot be taken as evidence of statistical discrimination.

5.1.2 Demand for mentorship quality in Managers

Table 8 shows clear evidence of respondents caring about high quality mentorship in their managers. Individuals on average are willing to give up around 80 thousand INR ($\approx$ \$ 3,800) or 11% of their average annual wages to work for a mentor who ranks one point higher on a five point scale. Converting it to standard deviation units, given the variation in mentorship rating in the scenarios presented to the respondent, one standard deviation increase in mentorship is valued at 5.65% of average annual wages. This estimate could seem large at first glance. However it is important to keep in mind that sample of respondents in the survey are job-seekers. Under diminishing returns to mentorship, a marginal increase in mentorship is of much higher value to first time job seekers than experienced employees in the labor market. Additional descriptive evidence on individuals caring for mentorship is presented in Appendix A.4

A potential concern could revolve around how respondents interpret average mentorship quality if the respondent is very different from the current employees. In such a case, the mentorship variable could be uninformative to the respondent.[33] [34] Even though our scenarios in the survey does not go into the specifics of types of jobs or industry to reduce cognitive load, but if respondents were indeed thinking of jobs dominated by out-group employees, such that mentorship quality is uninformative, then we would have seen its evidence in the data. On the contrary, the lower bound of the 95% confidence interval on preference for high rated mentors is at 10.5% of average annual wages, far away from zero.[35]

[33] Recall mentorship rating was presented in the survey as the average mentorship rating of the manager provided by its current employees.

[34] An example could be a female getting a job offer from the construction sector which is heavily male dependent

[35] There is less concern about other interpretations of mentorship. For example without the specifics of the industry mentioned, it is unlikely that a English major would be answering questions by invoking upon themselves the extra cognitive load to think about jobs outside of their domain of specialties unless that is what their belief about their future job prospects are, which is unlikely.

5.1.3 Beliefs

The estimates of the belief parameter on male manager's mentorship (α^G) are obtained from estimates of the vectors $(\tilde{\beta}^G, \beta^G, \beta^R)$ using equation (13). The first row of Table 9 show that on average male managers are believed to have higher mentorship rating by 0.14 points (on a 5 point scale) or by 0.28 standard deviations than female managers. Classical measurement error in the reported compensations leading to inflation of standard errors in the estimation of preference parameters will trickle down to the standard errors of the belief parameters. In spite of that the belief parameter estimate is statistically significant. The estimates in conjunction with the model and evidence that mentorship is a highly sought after manager attribute, imply that in the absence of a manager's rating, both genders believe that male managers are better mentors than female managers.

Table 9: Estimates of Belief Parameters

Belief Parameter	
α^G (Male managers)	0.140***
	(0.024)

Notes: Standard errors are computed using block bootstrap at the individual level with 1000 repetitions. Parameters are estimated by jointly estimating the complete and incomplete scenarios and using estimates of β^R $\tilde{\beta}^G$ and β^G presented in Table 8. The estimates come from the equation $\alpha^G = \frac{\tilde{\beta}^G - \beta^G}{\beta^R}$. ***, **, * denote statistical significance at 1, 5, and 10%, respectively.

Table 10: Estimates of Valuation of Beliefs

Belief Parameter	in 10^5 INR	% of wages
$\beta^R\alpha^G$ (Male managers)	0.112***	1.6%
	(0.024)	

Notes: Note that the estimates come from the equation $\beta^R\alpha^G = \tilde{\beta}^G - \beta^G$ for all $x \in \{G, W, H\}$ presented in Table 8. The standard errors are block bootstrapped at the individual level with 1000 repetitions. Parameters are estimated by jointly estimating the complete and incomplete scenarios. ***, **, * denote statistical significance at 1, 5, and 10%, respectively. Average annual wages was 7 lakh INR (≈ $38.8 thousand)

In in Table 10 we report estimates of the valuations of beliefs on male manager mentorship ($\beta^R\alpha^G$) in monetary units. The estimates of α^G are hard to interpret since they represent relative beliefs on a five point scale. To get an estimate of this we can multiply the belief on the difference in mentorship between male and female managers with how much wages individuals are willing to give up for a one point increase in mentorship rating (β^R) . Observe that the difference in the parameters between the complete and the incomplete scenarios provide us exactly that. $\tilde{\beta^x} - \beta^x = \beta^R \alpha^x$ for all $x \in \{G, W, H\}$. Thus estimates of $\beta^R \alpha^G$ provides us an estimate of the valuation of statistical discrimination.[36]

6 Heterogeneity

6.1 By Gender of respondent

In this section we present results by running the model separately for male and female respondents. that preferences on male managers (β^G) and beliefs on mentorship of male managers ($\beta^R \alpha^G$) on average are not statistically different between male and female respondents. But it is important to note that we do see economically lower statistical discrimination against female managers, among female respondents relative to male respondents.

Table 11: Jointly Estimated Complete & Incomplete Scenarios: By Gender of the Respondent

Incomplete Scenarios			**Complete Scenarios**		
Parameters	Female	Male	*Parameters*	Female	Male
$\tilde{\beta}^G = \beta^G + \beta^R \alpha^G$	-0.017	-0.000	β^G (Male Manager)	-0.096***	-0.135***
	(0.016)	(0.018)		(0.018)	(0.019)
$\tilde{\beta}^H = \beta^H + \beta^R \alpha^H$	1.047***	1.193***	β^H (Flexible Hours)	0.739***	0.803***
	(0.038)	(0.039)		(0.020)	(0.021)
$\tilde{\beta}^W = \beta^W + \beta^R \alpha^W$	1.350***	1.382***	β^W (Annual wages)	1	1
	(0.055)	(0.061)			
			β^R (Mentorship)	0.732***	0.835***
				(0.024)	(0.025)
Observations	4,840	6,760		4,840	6,760

Notes: Standard errors are computed using block bootstrapped allowing for arbitrary correlation at the individual level with 1000 replications. β^W is normalized to 1. Units are in hundred thousand INR. ***, **, * denote statistical significance at 1, 5, and 10%, respectively.

6.2 By Field of study

In this section we delve further and present results where we estimate the model separately for (a) respondents enrolled in different fields of study - Arts, Science and Engineering and (b) by gender of the respondents and field of study. Since sample size could go small here,

[36]Also note that $\frac{\mathbb{E}(\tilde{\beta^G}_i - \beta^G_i)}{\mathbb{E}(\beta^R_i)} \neq \mathbb{E}(\frac{\tilde{\beta^G}_i - \beta^G_i}{\beta^R_i}) = \mathbb{E}(\alpha^G_i)$. But $\mathbb{E}(\tilde{\beta^G}_i - \beta^G_i) = \mathbb{E}(\beta^R_i \alpha^G_i)$ provides us the average valuations of beliefs.

normal based confidence intervals may not be reliable (Hansen). Thus, for inference we use bias-corrected bootstrap confidence interval. In Table 12 we find that preferences and beliefs do not differ between Arts and Engineering. However respondents majoring in Science have lower statistical discrimination against female managers.

Table 12: Jointly Estimated Complete & Incomplete Scenarios: By Departments

	Arts		Science		Engineering	
Parameters	Preferences	Beliefs	Preferences	Beliefs	Preferences	Beliefs
$\tilde{\beta}^G$	-0.006		-0.015		-0.004	
	(0.021)		(0.020)		(0.019)	
β^G	-0.128***		-0.0987***		-0.119***	
	(0.0345)		(0.0268)		(0.0269)	
β^R	0.795***		0.785***		0.793***	
	(0.0430)		(0.0594)		(0.0566)	
$\beta^R\alpha^G$		0.122***		0.0841***		0.115***
		(0.0441)		(0.0309)		(0.0345)
Observations	10,360	10,360	5,120	5,120	7,720	7,720

Notes: The table report estimates of preference and belief parameters by running the model for each field of study (Arts, Science and Engineering). The table only shows coefficients of interest. β^G is the preference parameter for male managers. Note that $\beta^R\alpha^G = \tilde{\beta}^G - \beta^G$ is the valuation of beliefs of malae manager mentorship. It does not report the coefficients from the incomplete scenarios and that of flexible hours. As before, the preference parameter on wages is normalized to 1. The estimates are in units of hundred thousand of INR.

6.3 By Gender of Respondent and Field of study

To further investigate the result of lower statistical discrimination in Science majors, we run the model by the field of the study and the gender of the respondents in Table 13. We find that this result is driven by females in Science. In Table 13 we find no evidence of statistical discrimination towards female managers by female respondents who are enrolled in Science majors. Additionally among respondents majoring in Arts, we find that males are more than twice as statistically discriminatory as female respondents. Also among males, we find that statistical discrimination against female managers is the highest among males enrolled in Departments under Arts. Females in Engineering value mentorship rating at 7.4% of average annual wages which is the lowest among other gender-department pairs. However the standard error on this estimate is relatively much higher (it is in between three to five times as high as the standard errors on the estimates from other groups). This is potentially driven by very few females enrolled in Engineering departments (only about 10% of females in our sample are Engineering majors).

6.4 By Parental Education

In this section we estimate the model for two subsamples by relative parental education. In particular we estimate the model for respondents whose mothers are more educated than their fathers and vice versa. Our sample has 13% of individuals whose mothers are strictly more educated than their fathers. Table 14 reports estimates of the preference and belief parameters. For the sub sample, whose mothers are more educated than their fathers, we do not find evidence of statistical discrimination against female managers. The results do not change if we introduce another subsample whose mothers and fathers are equally educated.

6.5 Distributions

Each respondent in our survey answered questions in 20 scenarios with 3 jobs each. This gives us 40 unique data points of compensations across jobs of varying attributes for each respondent - 20 from the incomplete and 20 from the complete scenarios. We use this to estimate the model for each individual separately. Using the empirical distribution we can compute the sample mean preferences and given the standard deviation we can compute standard errors of the mean. But more interestingly, using the empirical distributions we carry out two informative exercises. First, we quantify the proportion of individuals who statistically discriminate against female managers. Second, we regress the estimated parameters on sample characteristics to understand correlation between discrimination and observable characteristics.[37] The second exercise in itself is also a simple test of heterogeneity.[38] We also address the concern that individual parameters may be estimated with noise and discuss bounds on the estimates under reasonable assumptions. For better interpretability of the figures we have converted the preference and belief parameters into percent of average annual wages.

The empirical CDF reveals that that 60% of respondents in absence of information on man-

[37]The parameter estimates which we regress on observable demographics could be noisy, being estimated from 40 observations. This will overstate the standard errors of the regression coefficients. Finding statistically significant estimates, in spite of overstated standard errors makes an even stronger case for heterogeneity.

[38]One could also take the route of testing for heterogeneity using wild bootstrap-t (Cameron, Gelbach & Miller (2008), Busso, Gregory & Kline (2013)). However in our case, the test is non-standard since it is testing at the boundary of the parameter space- testing the null hypothesis of zero variance against the alternative hypothesis of positive variance.

Table 13: Jointly Estimated Complete & Incomplete Scenarios: By Departments and Gender of respondent

Panel A: Female Respondents

	Arts		Science		Engineering	
Parameters	Preferences	Beliefs	Preferences	Beliefs	Preferences	Beliefs
$\tilde{\beta}^G$	-0.016		-0.031		0.002	
	(0.019)		(0.028)		(0.030)	
β^G	-0.0977***		-0.0817**		-0.107***	
	(0.0232)		(0.0368)		(0.0411)	
β^R	0.784***		0.653***		0.510**	
	(0.0469)		(0.0789)		(0.222)	
$\beta^R\alpha^G$		0.0816***		0.0507		0.109**
		(0.0300)		(0.0471)		(0.0547)
Observations	6,960	6,960	1,680	1,680	1,040	1,040

Panel A: Male Respondents

	Arts		Science		Engineering	
Parameters	Preferences	Beliefs	Preferences	Beliefs	Preferences	Beliefs
$\tilde{\beta}^G$	0.014		-0.007		-0.005	
	(0.050)		(0.027)		(0.022)	
β^G	-0.190**		-0.107***		-0.121***	
	(0.0952)		(0.0352)		(0.0296)	
β^R	0.819***		0.850***		0.836***	
	(0.0930)		(0.0709)		(0.0544)	
$\beta^R\alpha^G$		0.204*		0.100**		0.116***
		(0.119)		(0.0407)		(0.0370)
Observations	3,400	3,400	3,440	3,440	6,680	6,680

Notes: The table report estimates of preference and belief parameters by running the model for each possible pair of field of study (Arts, Science and Engineering) and gender of the respondent (male and female). β^G is the preference parameter for male managers. Note that $\beta^R\alpha^G = \tilde{\beta}^G - \beta^G$ is the valuation of beliefs of malae manager mentorship.. The table only shows coefficients of interest. It does not report the coefficients from the incomplete scenarios and that of flexible hours. As before, the preference parameter on wages is normalized to 1. The estimates are in units of hundred thousand of INR.

Table 14: Jointly Estimated Complete & Incomplete Scenarios: By Relative Parental Education

	Father more educated		Mother more educated	
Parameters	Preferences	Beliefs	Preferences	Beliefs
$\tilde{\beta}^G$	-0.005		-0.020	
	(0.013)		(0.033)	
β^G	-0.115***		-0.145***	
	(0.020)		(0.055)	
β^R	0.814***		0.644***	
	(0.032)		(0.096)	
$\beta^R\alpha^G$		0.109***		0.125
		(0.024)		(0.076)
Observations	20,200	20,200	3,000	3,000

Notes: The table report estimates of preference and belief parameters by running the model- first for the subsample whose mother's are more educated than their fathers and the second where the father is more educated than the mother. The table only shows coefficients of interest. β^G is the preference parameter for male managers. Note that $\beta^R\alpha^G = \tilde{\beta}^G - \beta^G$ is the valuation of beliefs of malae manager mentorship. It does not report the coefficients from the incomplete scenarios and that of flexible hours. As before, the preference parameter on wages is normalized to 1. The estimates are in units of hundred thousand of INR.

Figure 3: Distribution of Individual Beliefs

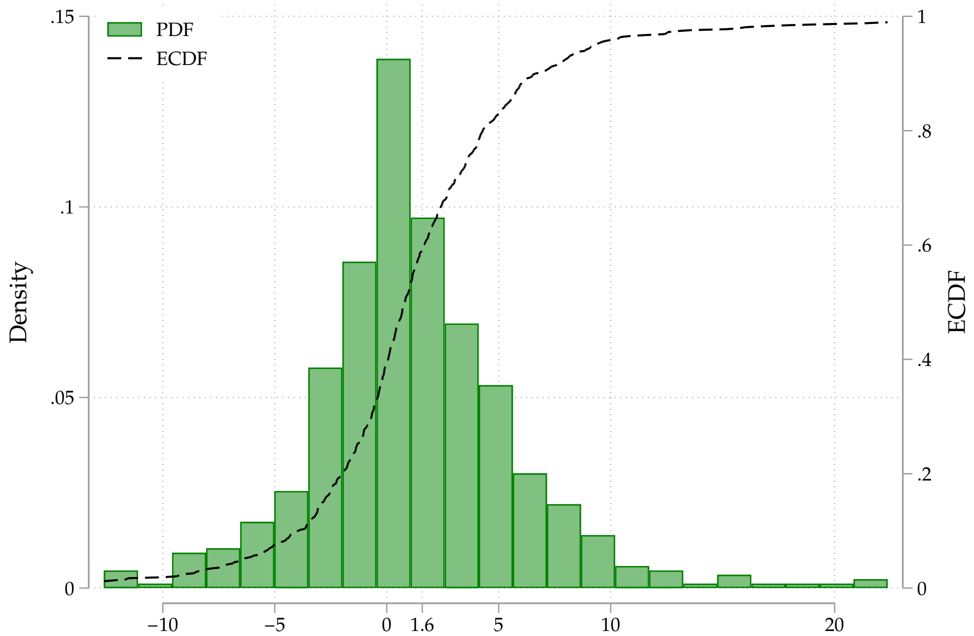

Notes: The figure plots the histogram and the smoothed empirical cumulative distribution function of individual beleifs obtained by estimating the model for each individual. The units on the x-axis are in percent of average of annual wages.

Figure 4: Distribution of Individual Preferences

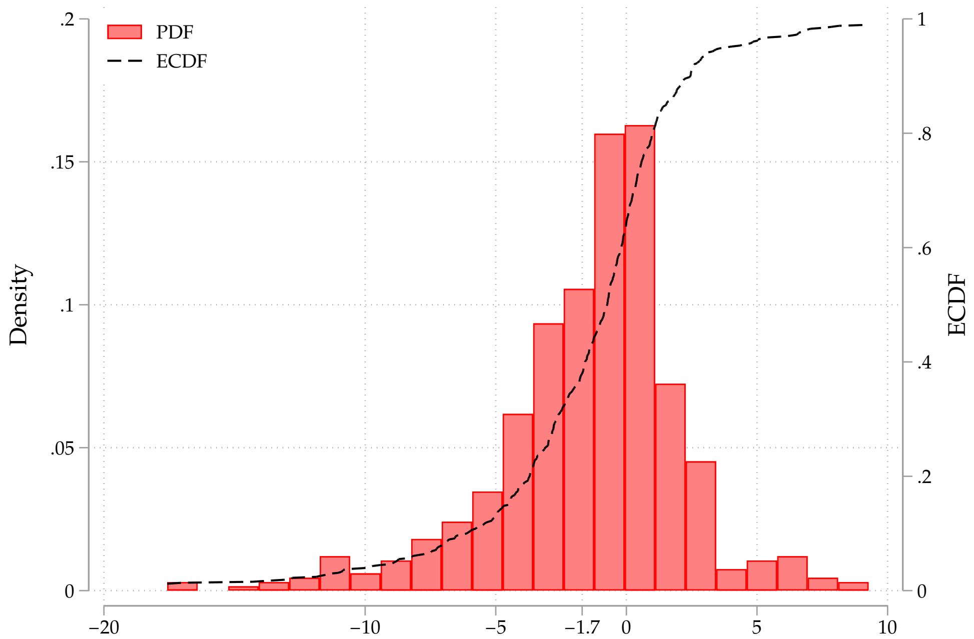

Notes: The figure plots the histogram and the smoothed empirical cumulative distribution function of individual preferences obtained by estimating the model for each individual. The units on the x-axis are in percent of average of annual wages.

ager quality believe female managers to be of worse quality shown in Figure ??. If the true underlying distribution of preferences, beliefs and noise are symmetric then this number is a lower bound on the proportion of individuals who statistically discriminate against female managers. This is because under symmetry the median is unaffected and hence the true cumulative distribution will intersect zero at a lower point than what we see in Figure ??. Upon comparison of the estimate of the mean and its standard error obtained by estimating the model individual by individual with the model model, we find that they are very close. Hence even though individual estimates could be noisy, this would not have been observed had that the noise on average not been zero. Similarly in Figure ?? we observe that individuals preferences upon receiving information on manager quality reveals that atleast 62% of individuals prefer working for female managers. Also it is important to note that it is not everyone who prefers to work for female. The reason we find an average to work for female managers is because of two reasons. First, there are more individuals who prefer working for female managers. Second, the amount of money that would make an individual who prefers to work for female managers to switch to work for male managers is higher than the amount needed to make an individual preferring to work for a male manager to switch and work for a female manager.

7 Validity of estimates of the belief parameter

In the penultimate section of the survey we directly elicit beliefs on manager's mentorship. We do this to corroborate the results that we obtain from the model. This part of the survey presented respondents with with 10 jobs. Each job had manager's name, wages and flexible hours. With each job we provided a slider scale ranging from zero to five and asked respondents to report their expected rating for each manager in each of the jobs. On this data, we project a linear model of reported expected ratings on the gender of the manager, annual wages and flexible hours, in alignment with the parametrization of beliefs in the model of job choice. We use this data as a check on the estimates of the belief parameters obtained from our model which exploits the variation in compensations between the complete and the incomplete scenarios. The objective is to check whether the result on statistical discrimination against female managers that we found from the estimation of the model, holds in the direct belief elicitation as well.

Table 15: Estimates of Expected Beliefs from Directly Elicited Belief data

	Model			Data on beliefs		
Belief Parameters	All	Females	Males	All	Females	Males
α^G (Male managers)	0.140***	0.108***	0.161***	0.091***	0.099***	0.084***
	(0.024)	(0.033)	(0.033)	(0.022)	(0.034)	(0.033)
α^W (Annual wages)	0.465***	0.478***	0.457***	0.428***	0.425***	0.429***
	(0.053)	(0.077)	(0.074)	(0.003)	(0.004)	(0.004)
α^H (Flexible hours)	0.449***	0.422***	0.467***	0.358***	0.352***	0.362***
	(0.040)	(0.060)	(0.056)	(0.021)	(0.032)	(0.030)

Notes: Standard errors bootstrapped at the individual level with 1000 repetitions. These estimates come from projecting a linear model on the data of reported expected manager's rating in each of 10 different jobs with annual wages, flexibility of hours and manager's name. The linear projection takes the form of $R_{ij} = \alpha^G G_j + \alpha^W W_j + \alpha^H H_j + \eta_{ij}$.

Table 15 reports the estimates by all individuals and also by the gender of the respondent. We see that this data not only matches the qualitative conclusion from the belief parameter estimates of statistical discrimination on female managers, but also the average estimates are quite close. Note that this data by itself can only help us estimate α^x and not $\beta^r \alpha^x$ for all

$x \in \{G, H, W\}$. However the primary objective of this data is served by corroborating that the sign of α^G is positive in both the estimates obtained from the job choice model using choice and compensation data as well as from the directly elicited belief data. This is an important result because it gives us the confidence in the estimates which tell us that individuals do believe females have lower mentorship quality. The concern of individuals manipulating their responses and misreporting their preferences for female managers is not an issue given this corroborating evidence.[39]

The objective is not to see how close are the estimates of α^G are, because of two additional reasons. First, they are obtained form two unrelated sources of the survey. Second, the estimates of α from the model as explained before could be noisy for individuals whose β^R is close to zero. This is added reason for why we emphasize on using $\beta^R\alpha^G$ as our measure of valuation of statistical discrimination instead of only α^G.

8 Robustness

A concern is whether individuals followed the instructions provided to them in the beginning of the survey. Identification is a result of individuals assuming that the jobs are identical in any other characteristic which is not observable to them. To ensure this and discipline our estimates, we designed specific questions at the end of the survey both direct and indirect, to infer on whether individuals actually followed the instructions. We dropped 2.2% of the sample as a result of this restriction and re-estimated the model. The estimates were robust to this restriction.

We also dropped 1% of the sample who finished the survey in less than 15 minutes. The choice of 15 minutes is somewhat ad-hoc, but it is driven by the distribution of time completion of the survey - the 1st percentile of completion duration was at 13.1 minutes. Estimates were robust to this restriction as well.

[39]It is necessary to discuss interpretations had the result from the direct belief elicitation appeared opposite or showed no effect on the manager's gender. Such a situation cannot completely be considered as evidence that respondents had manipulated their responses. This is because one cannot tell apart this explanation with an equally plausible explanation of learning from the complete scenarios.

9 Additional Discussions and Implications

An important consequence of our design of the information experiment is that under certain assumptions, it allows us to separate out beliefs and preferences.[40] Observe that given our results in the incomplete scenarios, it is important to note that observing indifference in the data is not necessarily an evidence of no discrimination. This can happen when preference and beliefs operate in opposite directions as it does in the incomplete scenarios.[41] We cannot comment on whether such beliefs held by individuals are biased or not. In order to do so one would need data on the population distributions of mentorship by gender of manager. In general, there is consistent evidence that individuals do not have a good sense of population distributions (Wiswall & Zafar (2015), Bordalo et al. (2016), Alesina et al. (2018), Bursztyn et al. (2018), Alesina & Stantcheva (2020)). Even though our paper can not speak to bias of beliefs, it is an important avenue to pursue, because under biased beliefs, search and job matches in equilibrium will be sub-optimal. There are many other instances in which our research applies. Examples include job-seekers who have alumni networks in firms and can get some information on manager and manager quality thus affecting their search and consequently the final match. Another example is of individuals who are already employed but seeking to switch teams within firms.

There are two primary implications of our research. First, without considering supply side selection, we are missing out on the complete picture of group level inequalities. Additionally studying beliefs and preferences of the supply side is essential because it affects search and equilibrium matches. Second is how firms might respond differently to such beliefs and preferences of workers, which could generate different rates of promotion of females to managerial position. This provides multiple potential avenues for future research extending from this paper. If firms have strong priors that matching employees with their preferred managers increases match productivity, this could potentially lead to females being promoted at higher rates, conditional on productivity. However females could still be promoted at lower rates if firms have high enough discriminatory preferences against females and are willing to forgo

[40] Analogous to separating taste-based and statistical discrimination.

[41] As explained previously workers may care about other aspects of the manager apart from mentorship and may use it as a measure of overall manager quality. In that case, we can still point-identify beliefs. But this added flexibility comes at the expense of giving up point identification of preferences (But still the results on preference remain suggestive).

increased profits resulting from more efficient matches. Thus in a way worker preferences could be used in testing for discriminatory practices by the firm at the executive levels those who decide whom to promote to managerial positions - one where workers prefer to work for female managers and conditional on productivity females are promoted at lower rates than males. It will be also interesting to explore whether individual preferences and perceptions about average preferences differ and more importantly whether perceptions are incorrect. In the spirit of Bursztyn, González & Yanagizawa-Drott (2018) this could serve as an additional reason on why female managers are promoted at lower rates.[42] Another interesting avenue for future research is to explore whether having more female managers allows firms to compete profitably for employees with otherwise similar firms, if there is overall preference to work for female managers. Additionally if there exists asymmetric information between incumbent and competing firms (Pinkston (2009), Kahn (2013)) then there are even higher information rents to be taken advantage of.

10 Conclusion

In this paper we provide novel documentation of discrimination by employees on manager's gender in choosing jobs and investigate whether it could contribute in generating glass ceilings for females in managerial positions. To answer this question, we design and conduct a hypothetical job choice survey involving a novel information experiment among job seeking students at a highly selective university in India. We present respondents with a series of hypothetical job scenarios consisting of jobs with exogenously varying attributes (annual wages, flexible hours, manager's name and mentorship rating of manager). Respondents are asked to choose their most preferred job and report wage compensations that make them indifferent between jobs. In the first ten scenarios (incomplete scenarios) individuals were shown jobs with all attributes except the mentorship rating of the manager. In the last ten scenarios (complete scenarios) individuals observed all the above attributes in each job. We identify preferences using the variation in compensations within the complete scenarios. The varia-

[42] Bursztyn, González & Yanagizawa-Drott (2018) show that in Saudi Arabia, husbands individually prefer having their wives to participate in the labor force but misperceive the social norms and believe that such preferences are uncommon in average. Upon revealing information about preferences held on average, thus correcting their misperceptions, the authors find that husbands enroll their wives in a costly training program designed for labor force participation.

tion in compensations between the complete and incomplete scenarios provide us necessary variation to identify beliefs on mentorship. We find that in the absence of information on manager mentorship—chocies driven by both preferences and beliefs— employees are indifferent between male and female managers. However, in presence of information on manager mentorship, we find strong preference to work for female managers. using a structural model of job choice, we estimate that individuals are willing to give up on average 1.7% of average annual wages to work for female managers. Hence in absence of additional information on manager mentorship, female managers are believed to be worse mentors than male managers. We quantify this statistical discrimination against female managers at around 1.6% of average annual wages. We find both genders prefer to work for female managers. Our documentation of heterogeneity in the preferences and beliefs of individuals correlate with demographics in the directions we would expect them (e.g. We do not find evidence of negative beliefs on female manager mentorship among female respondents majoring in Science). We corroborate these results—on negative beliefs on female manager mentorship—obtained from the job choice model using additional data on directly elicited beliefs. We collect this data at the end of our survey and information experiment, where we directly ask about expected mentorship rating of managers in different jobs. As discussed in details in the previous section, our results suggest that discrimination by firm executives in generating glass ceilings for females at managerial levels could be under-biased, if conditional on quality females are still promoted at lower rates. To that extent, our paper sheds light on how can preferences of employees on their managers could bee used as tests of discrimination by firm executives who decide on promotion, with additional data.

References

Abel, M. (2019), 'Do workers discriminate against female bosses?'.

Abel, M. & Buchman, D. (2020), The effect of manager gender and performance feedback: Experimental evidence from india, Technical report, Institute of Labor Economics (IZA).

Adams-Prassl, A. & Andrew, A. (2019), 'Preferences and beliefs in the marriage market for young brides'.

Agan, A. & Starr, S. (2018), 'Ban the box, criminal records, and racial discrimination: A field experiment', *The Quarterly Journal of Economics* **133**(1), 191–235.

Alam, M. M. U. (2019), Identification in models of discrimination. Working Paper.

Alesina, A., Miano, A. & Stantcheva, S. (2018), Immigration and redistribution, Technical report, National Bureau of Economic Research.

Alesina, A. & Stantcheva, S. (2020), Diversity, immigration, and redistribution, *in* 'AEA Papers and Proceedings', Vol. 110, pp. 329–34.

Altonji, J. G. & Pierret, C. R. (2001), 'Employer learning and statistical discrimination', *The Quarterly Journal of Economics* **116**(1), 313–350.

Arrow, K. J. (1972), 'Some mathematical models of race discrimination in the labor market', *Racial discrimination in economic life* pp. 187–204.

Arrow, K. J. (1998), 'What has economics to say about racial discrimination?', *Journal of economic perspectives* **12**(2), 91–100.

Athey, S., Avery, C. & Zemsky, P. (2000), 'Mentoring and diversity', *American Economic Review* **90**(4), 765–786.

Becker, G. S. (1971), 'The economics of discrimination', *Second Edition, Chicago: University of Chicago Press* .

Bertrand, M. & Duflo, E. (2017), Field experiments on discrimination, *in* 'Handbook of economic field experiments', Vol. 1, Elsevier, pp. 309–393.

Bertrand, M. & Mullainathan, S. (2004), 'Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination', *American economic review* **94**(4), 991–1013.

Blass, A. A., Lach, S. & Manski, C. F. (2010), 'Using elicited choice probabilities to estimate random utility models: Preferences for electricity reliability', *International Economic Review* **51**(2), 421–440.

Blau, F. D., Currie, J. M., Croson, R. T. & Ginther, D. K. (2010), 'Can mentoring help female assistant professors? interim results from a randomized trial', *American Economic Review* **100**(2), 348–52.

Bohannon, R. W. (1985), 'Mentorship: A relationship important to professional development: A special communication', *Physical therapy* **65**(6), 920–923.

Bohren, J. A., Imas, A. & Rosenberg, M. (2019), 'The dynamics of discrimination: Theory and evidence', *American economic review* **109**(10), 3395–3436.

Boneva, T. & Rauh, C. (2018), 'Parental beliefs about returns to educational investmentsthe later the better?', *Journal of the European Economic Association* **16**(6), 1669–1711.

Bordalo, P., Coffman, K., Gennaioli, N. & Shleifer, A. (2016), 'Stereotypes', *The Quarterly Journal of Economics* **131**(4), 1753–1794.

Bursztyn, L., González, A. L. & Yanagizawa-Drott, D. (2018), Misperceived social norms: Female labor force participation in saudi arabia, Technical report, National Bureau of Economic Research.

Busso, M., Gregory, J. & Kline, P. (2013), 'Assessing the incidence and efficiency of a prominent place based policy', *American Economic Review* **103**(2), 897–947.

Cameron, A. C., Gelbach, J. B. & Miller, D. L. (2008), 'Bootstrap-based improvements for inference with clustered errors', *The Review of Economics and Statistics* **90**(3), 414–427.

Charité, J., Fisman, R. & Kuziemko, I. (2015), Reference points and redistributive preferences: Experimental evidence, Technical report, National Bureau of Economic Research.

Delavande, A. & Manski, C. F. (2015), 'Using elicited choice probabilities in hypothetical elections to study decisions to vote', *Electoral studies* **38**, 28–37.

Dey, M. S. & Flinn, C. J. (2005), 'An equilibrium model of health insurance provision and wage determination', *Econometrica* **73**(2), 571–627.

Falk, A., Kosse, F. & Pinger, P. (2020), 'Mentoring and schooling decisions: Causal evidence'.

Fisman, R. & ONeill, M. (2009), 'Gender differences in beliefs on the returns to effort evidence from the world values survey', *Journal of Human Resources* **44**(4), 858–870.

Flabbi, L., Macis, M., Moro, A. & Schivardi, F. (2019), 'Do female executives make a difference? the impact of female leadership on gender gaps and firm performance', *The Economic Journal* **129**(622), 2390–2423.

Guryan, J. & Charles, K. K. (2013), 'Taste-based or statistical discrimination: The economics of discrimination returns to its roots', *The Economic Journal* **123**(572), F417–F432.

Heckman, J. J. (1998), 'Detecting discrimination', *Journal of economic perspectives* **12**(2), 101–116.

Kahn, L. B. (2013), 'Asymmetric information between employers', *American Economic Journal: Applied Economics* **5**(4), 165–205.

Kuziemko, I., Norton, M. I., Saez, E. & Stantcheva, S. (2015), 'How elastic are preferences for redistribution? evidence from randomized survey experiments', *American Economic Review* **105**(4), 1478–1508.

Lang, K. & Lehmann, J.-Y. K. (2012), 'Racial discrimination in the labor market: Theory and empirics', *Journal of Economic Literature* **50**(4), 959–1006.

Lange, F. (2007), 'The speed of employer learning', *Journal of Labor Economics* **25**(1), 1–35.

Lyle, D. S. & Smith, J. Z. (2014), 'The effect of high-performing mentors on junior officer promotion in the us army', *Journal of Labor Economics* **32**(2), 229–258.

Mas, A. & Pallais, A. (2017), 'Valuing alternative work arrangements', *American Economic Review* **107**(12), 3722–59.

Phelps, E. S. (1972), 'The statistical theory of racism and sexism', *The american economic review* **62**(4), 659–661.

Pinkston, J. C. (2009), 'A model of asymmetric employer learning with testable implications', *The Review of Economic Studies* **76**(1), 367–394.

Richard, O. C., Ismail, K. M., Bhuian, S. N. & Taylor, E. C. (2009), 'Mentoring in supervisor–subordinate dyads: Antecedents, consequences, and test of a mediation model of mentorship', *Journal of Business Research* **62**(11), 1110–1118.

Rosen, S. (1986), 'The theory of equalizing differences', *Handbook of labor economics* **1**, 641–692.

Sorkin, I. (2018), 'Ranking firms using revealed preference', *The quarterly journal of economics* **133**(3), 1331–1393.

Taber, C. & Vejlin, R. (2020), 'Estimation of a roy/search/compensating differential model of the labor market', *Econometrica* **88**(3), 1031–1069.

Wiswall, M. & Zafar, B. (2015), 'Determinants of college major choice: Identification using an information experiment', *The Review of Economic Studies* **82**(2), 791–824.

Wiswall, M. & Zafar, B. (2018), 'Preference for the workplace, investment in human capital, and gender', *The Quarterly Journal of Economics* **133**(1), 457–507.

A Appendix

A.1 Appendix Figures

Figure 1: Definitions of attributes

INSTRUCTIONS

In each question you will see a choice scenario.

A scenario will have 3 jobs (X, Y and Z), which you have to assume have been offered to you.

Each job will have 4 characteristics:

Manager: First name of the team's manager.

Annual wages: Gross annual salary (in lakhs).

Flexible hours: Whether the job allows for flexible hours or not.

Manager rating: This is the average rating of the mentorship of the manger, provided by this manager's current employees in an anonymous survey.

This is a measure of how good of a mentor is this manager to its subordinates.

This rating is on a scale of 1-5. The scale points mean as follows: 1: Poor; 2: Fair 3: Good, 4: Very good and 5: Excellent.

There will be 20 such choice scenarios.

Figure 2: Instructions

INSTRUCTIONS

In each scenario we will first ask you:

A. To choose one job among 3 job options.

Your instructions are:

To assume that all other characteristics, which are NOT MENTIONED here, are THE SAME in all the jobs. For example work from home under each manager is not shown here. You are to assume that its either available or not available in all three jobs. No job is different in anything that you do not observe,

In each scenario, we will then ask you for:

B. If you were to negotiate wages, how much ***minimum increase in wages*** *you would need in each of the other two jobs, for you choose them instead.*

Your instructions are:

State your wage increase, assuming that it does not change anything else about the job.

Let us give you an example from a survey of food preference.

A. Here you will see me choose one dish among 3 different dishes.

B. Then you will see me state how much MINIMUM drop in prices in the other two dishes would make me choose each of them instead.

Figure 3: Expected Rating

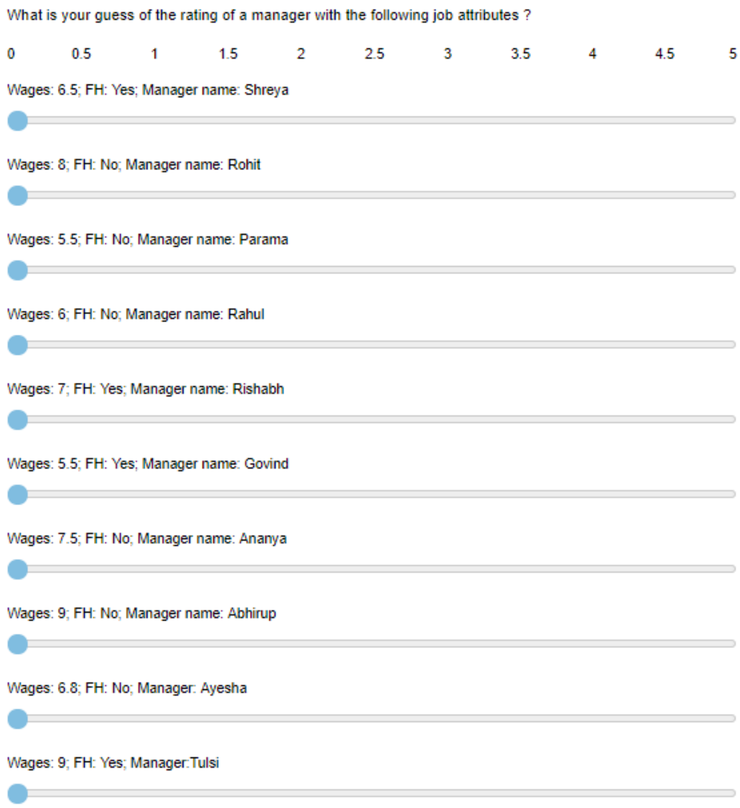

A.2 Details of Administration and Collection of Data

The online survey was designed and implemented on Qualtrics. Recruitment of students was done by the research assistants, who had previous experience in recruiting students for surveys and RCTs. In the midst of a nationwide lock-down, we administered the online survey in three key steps.

Step 1: The RAs, based on their previous experience, sent a sign-up link to each department's class' student representative, who distributed the link in the class lists. The sign-up sheet apart from containing the consent form, for them to sign, had asked about their email addresses to which the link of the survey would be sent, basic demographics, department, faculty of study (Arts/Science/Engineering) and level (Bachelors or Masters) and year of study. Students were also allowed to choose the date and the time at which they would like to take the survey. They had a choice among 4 dates starting April 8th to April 11th. On their selected date they had a choice of 1 among 6 time slots : 10am, 12 noon, 5pm, 7pm, 9pm and 11pm. We observed that pilots which had specified time-slots along with dates had higher completion rates than those with just dates. The sign-up form ended with a summary. The sign-up form was designed to automatically send them their signed-up form to their email address. This enabled us to automatically have the consent form's signed copy sent to the participant.

Step 2: Upon receiving sign-ups we scheduled emails to be sent out with unique links to the survey for each participant, an hour before each they were scheduled to participate in the survey. Hence that link could not be used on two different devices, to fill up the survey. The survey was also designed to prevent ballot-boxing i.e. once a survey is completed using a link, that link when clicked on again would show a confirmation of the survey being already completed. The links were designed to expire within 24 hours. Thanks to the extensive pilot done before, we did not face any technical difficulties while implementing the survey. Debriefs with pilot participants were extremely helpful in rewording the questions in order to maximize correct communication and understanding.

Step 3: The mode and details of online payment were selected at the last section of the survey. The options included direct bank transfers, PayPal and UPI (Unified Payment Interface). The payment was processed to the list of students who were verified within the pre-stated timeline for each mode of payment.

A.3 Estimation details

We implement the joint estimation of individual indifference in the complete and incomplete scenarios in a fully interacted model by stacking up the matrices of observables across the complete and the incomplete scenarios. In particular we estimate the following constrained least squares regression with no constant,

$$\begin{bmatrix} \mathbf{\Delta_{jk}} \\ \vdots \\ \tilde{\mathbf{\Delta}}_{\mathbf{jk}} \\ \vdots \end{bmatrix} = \begin{bmatrix} \mathbf{X_j} - \mathbf{X_k} & \mathbf{0} \\ \vdots & \vdots \\ \mathbf{0} & \tilde{\mathbf{X}}_{\mathbf{j}} - \tilde{\mathbf{X}}_{\mathbf{k}} \\ \vdots & \vdots \end{bmatrix}' (\boldsymbol{\beta} \quad \tilde{\boldsymbol{\beta}}) + \mathbf{e} \tag{16}$$

where the constraint is the normalization for the preference parameter on wages to be equal to one. The standard errors are computed using block bootstrap at the student level. This accounts for any arbitrary correlation between responses at the student level.

Using estimates of β and $\tilde{\beta}$ we estimate $\beta^R\alpha$ which gives us an interpretable measure of the statistical discrimination in terms of willingness to pay.

The block bootstrap algorithm is as follows: Sample contains N individuals.

1. Generate B block bootstrap samples of N individuals each.

   For each $b = 1, .., B$

2. Estimate the model for each member in the bootstrap sample, by bootstrapping each member's responses.

   Obtain $\hat{\beta}_i^{(b)}$ and compute its sample mean $\hat{\beta}^{(b)} = \sum_{i=1}^{N} \hat{\beta}_i^{(b)}$

Compute mean and standard deviation of the B estimates in hand to generate estimates of the bootstrap mean and bootstrap standard error.

$$\hat{\mathbb{E}}(\beta_i) = \frac{1}{B} \sum_{b=1}^{B} \hat{\beta}^{(b)}$$

$$\widehat{\text{std error}}(\beta_i) = SD(\hat{\beta}^{(b)})$$

Figure 4: Bootstrap Distribution of Beliefs and Preferences on Male Mangers

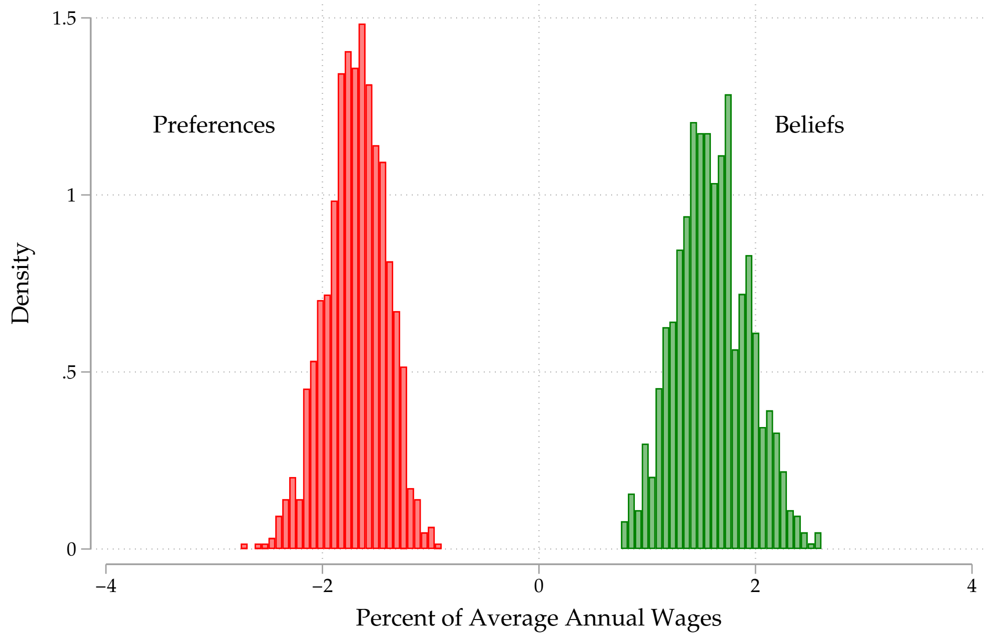

Notes: The figure shows bootstrap discributions of beliefs and preferences in units of percentage of average annual wages. The bootstrap distirbutions are obtained from 1000 block bootsrapped samples, using the alogrithm described in Appendix A.3. These bootstrap distributions are used for estimation of means and standard erros of preferences and beleifs.

A.4 Descriptive evidence on mentorship

After respondents answered choice and compensations questions for twenty scenarios, we asked them descriptive questions on different qualities of a manger along with demographic questions. In particular, we asked to rank three different qualities of managers that they would care about the most- competence, mentorship and peasant to work with. These rankings are plotted in Figure 5 for each of the above qualities. 40.8% of individuals rank mentorship above the other two qualities and more than 75% rank it at either 1 or 2. This shows a stark difference on how on average job-seekers care about mentorship in their managers above other qualities.

Figure 5: Ranking of manager qualities

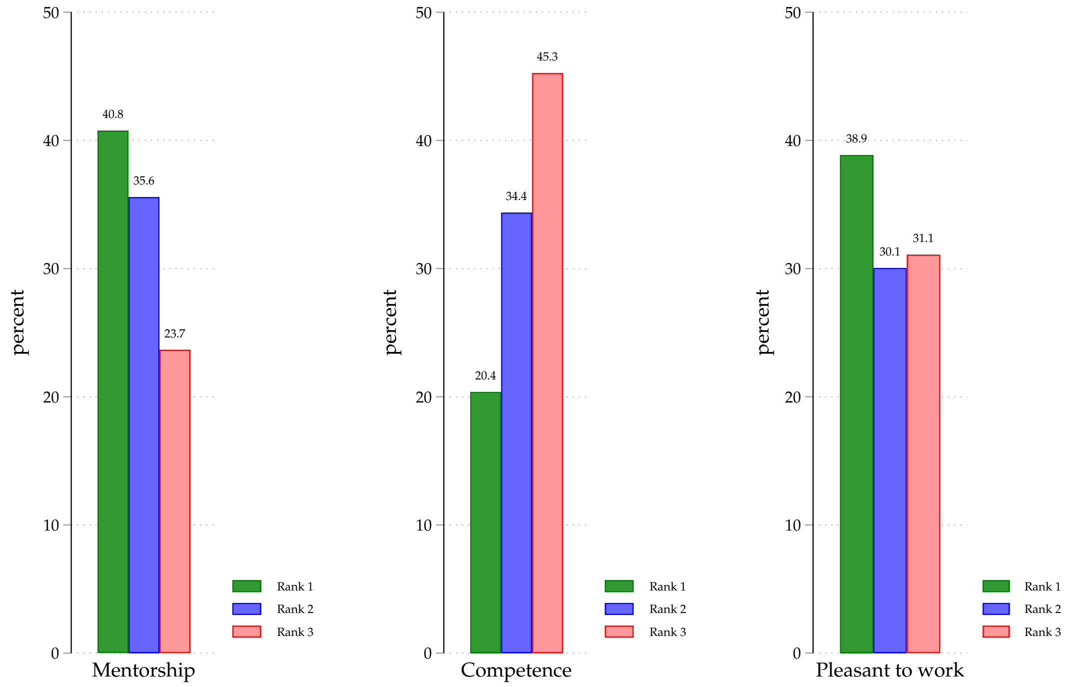

Notes: This figure plots data on what respondents care about in their managers among the following qualities - mentorship, competence, pleasant to work.

A.5 Generic model

In this section, we delineate a more generic model than the one presented in the main paper. The identifying assumptions remain the same. However, we relax the interpretation of the attribute of mentorship of the manager. In this specification, individuals care about overall manager quality (Q) besides caring about wages, flexible hours and gender of the manager. The mentorship rating acts as a signal of overall manager quality. Individuals care about this overall manager quality. The purpose of this generic model is to show that the results of statistical discrimination still holds.

Redefining the set of attributes to $A \equiv \{G, W, H, Q\}$, the utility of individual i takes the same linear form,

$$U_{ij} = \sum_{x \in A} \beta_i^x x_j + \epsilon_{ij} \tag{17}$$

Observe that now in both the complete and the incomplete scenarios, individuals will need to form expectations on the manager quality. In the incomplete scenario they will not have information on the manager's mentorship rating whereas in the complete scenarios they will have information about it. The expected utilities in the complete and the incomplete scenarios take the following forms-

$$\begin{aligned}\text{Incomplete Scenarios: } \mathbb{E}_i[U_{ij}|\tilde{X}_j] &= \sum_{x\in A\setminus Q} \beta_i^x x_j + \beta_i^Q \mathbb{E}_i(Q_j|\tilde{X}_j) + \mathbb{E}_i(\epsilon_{ij}|\tilde{X}_j)\\ \text{Complete Scnearios: } \mathbb{E}_i[U_{ij}|X_j] &= \sum_{x\in A\setminus Q} \beta_i^x x_j + \beta_i^Q \mathbb{E}_i(Q_j|X_j) + \mathbb{E}_i(\epsilon_{ij}|X_j)\end{aligned} \tag{18}$$

As explained above in both the expected utilities the individuals forms expectations on manager quality. But in the complete scenario the individual has the additional information of manager's mentorship rating. We parameterize the expectation on manager's quality in the following way-

$$\begin{aligned}\text{Complete Scenarios: } \mathbb{E}_i(Q_j|X_j) &= \sum_{x\in A\setminus Q} \gamma_i^x x_j + \gamma_i^R R_j + \mathbb{E}_i(\zeta_i|X_j)\\ \text{Incomplete Scenarios: } \mathbb{E}_i(Q_j|\tilde{X}_j) &= \sum_{x\in A\setminus Q} \gamma_i^x x_j + \gamma_i^R \mathbb{E}[R_j|\tilde{X}_j] + \mathbb{E}_i(\zeta_i|\tilde{X}_j)\\ &= \sum_{x\in A\setminus Q} \tilde{\gamma_i}^x x_j + \mathbb{E}_i(\zeta_i|\tilde{X}_j)\end{aligned} \tag{19}$$

Incorporating the above in the expected utility functions in both scenarios, we have

$$\begin{aligned}\text{Complete Scenarios: } \mathbb{E}_i[U_{ij}|X_j] &= \sum_{x\in A\setminus Q} (\beta_i^x + \beta_i^Q \gamma_i^x) x_j + \mathbb{E}_i(\epsilon_{ij}|X_j)\\ \text{Incomplete Scenarios: } \mathbb{E}_i[U_{ij}|\tilde{X}_j] &= \sum_{x\in A\setminus Q} (\beta_i^x + \beta_i^Q \tilde{\gamma}_i^x) x_j + \mathbb{E}_i(\epsilon_{ij}|\tilde{X}_j)\end{aligned} \tag{20}$$

Then with the same set of identifying assumptions, given the reported compensations and normalizing $\beta^W = 1$, we have the following indifference conditions

$$\begin{aligned}\Delta_{ijk} &= \sum_{x\in A\setminus Q} \underbrace{(\beta_i^x + \beta_i^Q \gamma_i^x)}_{\beta_i^{x(CS)}} (x_j - x_k)\\ \tilde{\Delta}_{ijk} &= \sum_{x\in A\setminus Q} \underbrace{(\beta_i^x + \beta_i^Q \tilde{\gamma}_i^x)}_{\beta_i^{x(IS)}} (x_j - x_k)\end{aligned} \tag{21}$$

The differences in the coefficients infront of the gender differences across the complete and the incomplete scenarios give us.

$$\beta_i^{G(CS)} - \beta_i^{G(IS)} = \beta_i^{Q}(\gamma_i^{G} - \tilde{\gamma}_i^{G}) \tag{22}$$

γ encapsulates the information on manager quality given X_j whereas $\tilde{\gamma}$ encapsulates the information on manager quality given $\tilde{X}_j$ i.e. in the absence of the information on manager quality. In presence of statistical discrimination against female managers this should be negative. This is what our estimates show, given that individual care positively about the quality of the manager. Thus under this model specification statistical discrimination is identified. Observe the analogy with the model presented in the main paper. Here too, if the individual does not care about the quality of the manager (i.e. $\beta^{Q} = 0$), the parameters identified from the complete and the incomplete scenarios must be identical, because the variation in information revelation will have no affect.