

Gendered Beliefs and the Job Application Decision: Evidence from a Large-Scale Field and Lab Experiment*

Kristine Koutout

November 2020

Abstract

This paper investigates how workers' job application decisions are affected by their beliefs about hiring managers' beliefs regarding the relative productivity of women and men. To this end, I combine a natural field experiment with a lab experiment. In the field experiment, I partner with a firm to solicit approximately 5,000 job applications using ads that randomize over the gender of the hiring manager and the gender associations of the product sector. I then recruit the same job-seekers to a structured online lab experiment to elicit their beliefs about hiring managers' beliefs, based on the manager's gender and product sector. Truth-telling is incentivized with the Binarized Scoring Rule, using a procedure I adapt from [Dustan, Koutout, and Leo \(2020\)](#). I find that men are more likely to apply for a job with a manager whom they believe has beliefs that favor men more. A one standard deviation increase in beliefs about the manager's beliefs increases the probability a man applies by 70%. On the other hand, women are unresponsive to their beliefs about managers' beliefs. These results have important implications for the sorting by gender behavior driving a large part of the gender wage gap.

*I am grateful to Andrea Moro, Andrew Dustan, Greg Leo, Myrna Wooders, Catherine Eckel, and Ebehi Iyoha for comments, as well as discussants and participants at the American Economic Association Annual Conference, Advances in Field Experiments Annual Conference, Economics Science Association North America Annual Conference, Southern Economics Association Annual Conference, and Winter Experimental Social Science Institute in Abu Dhabi. The pre-analysis plan for the full experiment has been registered with the AEA RCT Registry, #AEARCTR-0005472. Funding from the National Science Foundation Doctoral Dissertation Research Improvement Grant 1948723 and the Russell Sage Foundation Small Grants in Behavioral Economics is gratefully acknowledged.

1 Introduction

One of the most important factors contributing to the gender wage gap is that women and men work in different occupations and industries (Blau & Kahn, 2017). This sorting by gender behavior is at least partially driven by workers’ job application decisions, yet we know little about why women and men apply to different jobs.¹ I propose that workers’ beliefs about hiring managers’ beliefs—their *second-order beliefs*—may be an important factor contributing to workers’ job application decisions. If workers believe that managers believe one demographic group is more productive in a job than another group, then workers will likely sort based on that demographic characteristic. Through anticipation of discrimination, historically persistent beliefs about productivity differences across groups—*first-order beliefs*—could have reverberating consequences through workers’ second-order beliefs.

To determine how second-order beliefs affect workers’ job application decisions, I combine a large-scale natural field experiment with a structured online lab experiment. In the field experiment, I solicit applications for a *real* high-paying white-collar job using ads that vary the relevant second-order belief. Then, I invite participants in the field experiment to an online lab experiment that elicits their beliefs using a carefully designed incentive-compatible procedure. By combining the observation of natural behavior in the high-stakes labor market with the rigor of belief elicitation in structured lab experiments, I identify precisely the primitives affecting workers’ real-world job application decisions.

I vary the relevant second-order beliefs in the job ad by changing the hiring manager types that workers face. Manager type is characterized along two dimensions: manager gender and the gender associations of the product sector (which I refer to as *sector gender* throughout the paper). This variation captures two channels through which second-order beliefs could drive differential sorting behavior by gender.

First, workers may choose to apply for jobs based on the expected gender composition of management. Workers who believe that female managers’ beliefs favor women and male managers’ beliefs favor men may sort into a gender-matching job because they are more likely to face a same-gender manager, whom they believe is more likely to favor the worker’s own gender.² By explicitly varying the gender of the hiring manager in the job ads I use to solicit workers’ applications, I isolate how workers’ beliefs about the beliefs of different gender managers affects their job application decision separately from their expectations about the gender composition of management.

Second, workers may make their job application decisions based on job-specific

¹See Barbulescu and Bidwell (2013); Fernandez and Campero (2017); Fernandez and Friedrich (2011); Fernandez and Mors (2008); Flory, Leibbrandt, and List (2015) for studies that show women and men make different job application decisions.

²I refer to managers as female/male and workers as women/men throughout the paper. The lab experiment describes managers as female/male. The intent of the language is to indicate gender identity, not biological gender.

characteristics that affect managers' beliefs, and hence workers' beliefs about those beliefs. For example, if workers believe that managers of any gender believe women are better at teaching than men, and men are better at construction work, they may sort by gender into these occupations regardless of the gender composition of management. Alternatively, workers may believe that managers for different jobs hold differing beliefs about the relative productivity of women and men based on manager selection or learning. Workers could believe that men who become managers in men-dominated jobs have systematically different beliefs from men who become managers in women-dominated jobs, and likewise for women, or that beliefs are learned by managers through experience in their occupation or industry. In addition to manager gender, I vary sector gender to capture all of these potential differences in managers that could drive sorting behavior.

The experiment focuses on a single male-dominated occupation: an outside business to business (B2B) sales position at a regional wholesale distributor of disposable goods.³ The job is highly-compensated with commissions that typically lead to six-figure incomes and generous benefits. The firm I partner with is medium-sized with several hundred employees located in seven metropolitan statistical areas (MSAs) of the U.S. Importantly, the firm distributes several categories of disposable goods, including janitorial and office products. I exploit the firms' multiple product categories in the experiment design by asking workers to choose between applying for a job in the janitorial product sector, which is traditionally male, or in the office product sector, which is traditionally more female.

The field experiment proceeds in two stages. In the first stage, I solicit the contact information of workers interested in, or with experience in, the sales occupation through ads on job boards that do not reveal the manager's gender or product sector. This first stage allows me to observe both applicants and non-applicants, as well as elicit the beliefs of both groups in the later lab experiment. In the second stage, I treat workers with an advertisement for an outside B2B sales position. Workers may apply either to the office product sector or the janitorial product sector. I vary the gender of the hiring manager in each sector and, in doing so, vary the relevant second-order beliefs.

One week after workers receive the job advertisement, I send them an invitation to participate online in a paid economics study with Vanderbilt University. The online study is structured like a traditional lab experiment, in which I adapt the procedure I developed in [Dustan, Koutout, and Leo \(2020\)](#) to measure beliefs (and beliefs about beliefs) about the differences between two populations in some measurable characteristic. I design the procedure to target the beliefs relevant to the job application decisions the workers made in the field experiment.

The procedure elicits beliefs using the Binarized Scoring Rule ([Hossain & Okui, 2013](#)) to incentivize truthful revelation from all expected utility maximizers, regard-

³The 2018 American Community Survey (ACS) finds that 29% of wholesale and manufacturing sales representatives are women.

less of risk preferences, as well as some non-expected utility maximizers. First, I elicit participants' first-order beliefs about the difference between women's and men's sales ability.⁴ Then, I inform participants that managers responsible for hiring and supervising sales representatives reported their first-order belief in an earlier experiment. Participants then report what they believe those managers reported, which are the participants' second-order beliefs. I elicit participants' second-order beliefs by manager gender and sector gender, for a total of four second-order beliefs per participant, corresponding to each type of manager participants could have faced in the field experiment.

The main result is that men's, but not women's, second-order beliefs affect their job application decision. I obtain the result by combining the field experiment data on application behavior and the lab experiment data on beliefs for the subsample of job-seekers in the field experiment who participated in the lab experiment. Using the random assignment of the first manager type listed in the job ad, I show that men are more likely to apply for a job with a manager type whom they believe has beliefs that favor men more. The effect size is large, with a one standard deviation increase in the favorability of beliefs increasing the probability a man applies by 5.5 percentage points on a base of 7.8%. I find no such effect for women in the male-dominated occupation for which I recruit. Since women in the experiment are already selected on interest in a male-dominated occupation, I suggest that this selection could be partly based on unresponsiveness to anticipated discrimination.

In addition to the main result, I find independently policy-relevant results in the component experiments. I test three pre-specified hypotheses in the field experiment. First, men are about twice as likely to apply to the job as women. Second, workers do not systematically sort based on sector gender. Lastly, women are unresponsive to manager gender, but men are significantly more likely to apply to *female* managers.

In the lab component of the experiment, I find that participants' beliefs about managers' beliefs are consistent with sorting by gender behavior. Most participants believe that female managers in traditionally female sectors favor women and that male managers in traditionally male sectors favor men. Beliefs about managers whose gender does not match the sector gender—female managers in male sectors and male managers in female sectors—are less homogeneous, suggesting that policies like quotas and affirmative action may have mitigating effects on sorting by gender behavior. Importantly, I find little difference in women's and men's beliefs.

This paper makes three contributions to the existing literature. I present the first direct evidence that second-order beliefs affect labor market behavior. While beliefs about other people's beliefs may play a role in mechanisms studied in other papers, none to my knowledge explicitly study second-order beliefs. Second, the lab experiment provides the first measurement of second-order beliefs outside of an

⁴Specifically, I ask participants who they believe performed better on a sales ability assessment used by employers to screen workers—a randomly selected woman or a randomly selected man—and by how much.

undergraduate student population. Second-order beliefs have typically been studied in a game context using lab experiments on students and focus on beliefs about beliefs regarding strategic actions, not beliefs about beliefs regarding differences in populations. Lastly, the field experiment produces new evidence on differential job application decisions by gender and demonstrate that workers' beliefs about managers' beliefs are an important mechanism driving these job application decisions.

The second-order beliefs I study were first measured by [Dustan, Koutout, and Leo \(2020\)](#). In that paper, we develop a methodological framework for eliciting first- and second-order beliefs about the differences between two populations in some measurable characteristic in the lab. We find that a sample of undergraduate students *believe* that women and men hold different beliefs, despite observing no statistically distinguishable differences in the beliefs of women and men. Prior work measuring second-order beliefs elicited beliefs about the beliefs of another player regarding game play strategies. For example, [Manski and Neri \(2013\)](#) elicit second-order beliefs about actions in a simple game to study consistency between actions and beliefs and [Kneeland \(2015\)](#) uses a ring-network game to establish second- and higher-order levels of rationality. In this paper, I investigate whether second-order beliefs about the differences between populations are relevant to real-world decisions.

The closest paper to the field experiment component of this study is [Flory et al. \(2015\)](#), which uses a natural field experiment to study how women's and men's job application rates respond to competitive compensation schemes. They also vary the gender associations of the job and, in a follow-up to the main experiment, consider whether expectations of discrimination drive their results by varying manager gender. Unlike [Flory et al.](#), I find no effect of sector gender on application decisions. I also detect an effect of manager gender on men, while they find none. More importantly, I build on [Flory et al.](#) by combining the natural field experiment with a lab experiment eliciting beliefs to precisely identify the role of second-order beliefs in workers' job application decisions.

While second-order beliefs have not been explicitly studied in the labor market, theoretical and empirical work exists that is suggestive of their importance. Expectations of discrimination, as [Flory et al. \(2015\)](#) describe it, or anticipation of discrimination is a mechanism that potentially operates on workers' second-order beliefs. [Alston \(2019\)](#) shows that participants anticipate discrimination in a lab experiment in which they can pay to either reveal or hide their gender on a job application for a stereotypically male task. This anticipation of discrimination could be driven by participants who believe that managers in the synthetic labor market believe that women are less productive at the stereotypically male task; however, other explanations are possible. Workers could also anticipate discrimination based on the manager's tastes, which may be unrelated to the worker's beliefs about the manager's beliefs regarding the worker's productivity.

Similar mechanisms operate in equilibrium models of statistical discrimination, where beliefs become a self-fulfilling prophecy (see [Fang and Moro \(2010\)](#) for a

review). In a recent empirical paper, [Glover, Pallais, and Pariente \(2017\)](#) show that minority workers in a French grocery store chain exert less effort under more biased managers, as measured by the Implicit Association Test (IAT). The self-fulfilling prophecy [Glover et al.](#) uncover could be based on workers' beliefs about the beliefs of managers with respect to worker productivity; however, again, other explanations are possible. Only measurements of belief can elucidate their relevance, but second-order beliefs have never been directly measured using an incentive-compatible mechanism. The beliefs I elicit demonstrate that workers' primitives are in line with the self-fulfilling prophecy mechanism.

Numerous studies have shown that workers sort by gender at the job application decision. When applying for jobs, women sort into lower-paying jobs ([Barbulescu & Bidwell, 2013](#)), lower-level job queues ([Fernandez & Campero, 2017](#); [Fernandez & Mors, 2008](#)), more stereotypically female job titles ([Fernandez & Friedrich, 2011](#)), and less competitive compensation schemes ([Flory et al., 2015](#)). My results suggest that this sorting behavior could be due to workers' second-order beliefs.

Moreover, sorting by gender behavior is an important factor contributing to the gender wage gap. Large parts of the gender wage gap in numerous countries can be attributed to women and men working in different firms, occupations, and industries.⁵ This study suggests that workers' beliefs about managers' beliefs in these different jobs may contribute to this sorting behavior underlying a large part of gender wage gaps.

The paper proceeds as follows. I describe the experiment design and implementation in Section 2. Then, I present my main results on workers' beliefs and the job application decision in Section 3. I present additional results from the component field experiment in Section 4 and the component lab experiment in Section 5. I conclude with a discussion of the results in Section 6.

2 Experiment Design

In this section, I detail the experiment design and implementation for determining the causal effect of workers' second-order beliefs on their job application decisions. First, I describe the field experiment in subsection 2.1. I then turn to the lab experiment in subsection 2.2, including a careful specification of the beliefs of interest. Lastly, I discuss implementation in subsection 2.3 and present descriptive statistics for the full sample from the field experiment and the subsample who completed the lab experiment.

⁵See [Bayard, Hellerstein, Neumark, and Troske \(2003\)](#); [Blau and Kahn \(2017\)](#); [Carrington and Troske \(1998\)](#); [Groshen \(1991\)](#) for evidence in the U.S., [Card, Cardoso, and Kline \(2016\)](#); [Cardoso, Guimarães, and Portugal \(2016\)](#) for Portugal, [Amuedo-Dorantes and De la Rica \(2006\)](#) for Spain, [Fortin and Huberman \(2002\)](#) for Canada, [Gallen, Lesner, and Vejlin \(2019\)](#) for Denmark, and [Fafchamps, Söderbom, and Benhassine \(2009\)](#) for Africa.

2.1 Field Experiment

The field experiment proceeds in two stages. In the first stage, I solicit the contact information of workers with experience or interest in the sales occupation. In the second stage, I treat those workers with a job advertisement that varies the gender of the manager workers face in each sector and observe their job application decision.

The job for which I solicit applications is an outside business to business (B2B) sales representative position with a medium-sized regional wholesale distributor of multiple product sectors. The job is highly compensated, advertised as six-figure income potential, and includes generous benefits like health insurance, employer-paid life insurance, vacation days, and a 401(k) with matching contributions. The partner firm has an ongoing need for sales representatives and hires on a rolling basis. I solicit job applications from workers in all seven metropolitan statistical areas (MSAs) in which the firm is located.

This job has a number of characteristics that facilitate the study of workers' beliefs on the job application decision. First, an "outside" sales representative spends most of their time away from the office, as suggested by the title. This factor decreases the amount of interaction with the manager, minimizing the role of preferences associated with in-person exchanges (e.g. sexual harassment). Second, compensation is self-determined through commissions.⁶ This factor decreases the relevance of the manager's role in determining wages, increasing the relative weight of the manager's role in the hiring decision. Both of these factors contribute to isolating the probability of being hired as the determining factor in the worker's job application decision.

In addition, the flexibility in hours and place of work of the outside sales position increases the appeal of this job to women workers (Goldin, 2014; Wiswall & Zafar, 2018). Since wholesale sales is a male-dominated occupation, women make up a smaller proportion of the population of workers interested in this job. The flexibility factor increases my ability to recruit women workers and is noted in the job advertisement to augment its salience.

Lastly, the multiple product sectors in which the firm operates enable the study of gendered product sectors. I can vary the product sector in which the worker would potentially be employed without varying any firm-specific characteristics including: firm size, location, culture, and gender composition. The two product sectors I advertise are the office products sector, which I characterize to workers as traditionally female, and the janitorial products sector, which I characterize as traditionally male, though workers familiar with the sectors will already be aware of their gendered nature.

⁶Initially, workers are paid a salary in a type of "trial period"; however, successful workers are transitioned to commission-based pay.

2.1.1 Stage 1: Create Pool of Workers

In the first stage of the field experiment, I solicit the contact information of workers with experience or interest in the sales occupation in order to develop a pool of workers. By establishing a pool of workers, I can observe the behavior of both applicants and non-applicants, as well as elicit the beliefs of both groups in the lab experiment. I solicit the contact information of workers in two ways. First, I place a generic advertisement on the job boards ZipRecruiter and Monster, as well as send the same generic advertisement through “Sponsored InMail” on LinkedIn. The generic advertisement does not include the firm’s name or information about the relevant product sectors. Rather, it provides a list of job responsibilities, general compensation information,⁷ and basic qualifications.⁸ Workers are encouraged to respond to the ad to learn specifics about the company and a full job description. Appendix A contains the generic job advertisement.

Workers respond to the generic advertisement by clicking “Apply” on ZipRecruiter or Monster, or clicking “Learn More” on LinkedIn. The apply button on ZipRecruiter and Monster pre-fills workers’ contact information and allows workers the option of attaching a resume, though a resume is not required. The button on LinkedIn takes workers to the website of the firm set up to manage the recruiting process, where they fill out a form with their contact information.

The second way in which I solicit contact information is through the resume databases on ZipRecruiter and Monster. Workers may choose to include their resume in the database to make it searchable by employers. From these resumes, I obtain the contact information of workers to include them in the experiment. I download the resumes of workers within 50 miles of one of the seven MSAs in which the firm is located who have been active on the website in the last 10 days, and who either 1) have a history of working in a sales position, or 2) express interest in a sales position through their profile on the job board or their objective statement on the resume. These workers, along with those who provide their contact information directly as described above, constitute the “full sample” of workers.

2.1.2 Stage 2: Invite workers to apply for a job

In the second stage of the field experiment, I email the full sample of workers the complete job advertisement. Both the office product sector and the janitorial product sector are listed in the ad, with a different hiring manager associated with each sector. Workers who choose to apply must also choose a product sector.

Treatment varies the gender of the manager in each product sector. The manager may be either male or female and may be assigned to either the male or female sector. The combination of manager gender and sector gender characterize the four manager

⁷No exact figures are provided for compensation. The wording describing the compensation was prescribed by the firm.

⁸A bachelor’s degree, with experience in outside sales preferred.

types, listed in Table 1. Workers never choose between a male and a female manager in the same sector because it would deviate from job advertisement norms to ask workers to choose between two hiring managers for the same job.

Table 1: Manager Types

Type	Manager Gender	Product Sector
$\{m, j\}$	Male	Janitorial (M)
$\{f, j\}$	Female	Janitorial (M)
$\{m, o\}$	Male	Office (F)
$\{f, o\}$	Female	Office (F)

Since workers always face one manager in the office sector and one in the janitorial sector, there are four treatments, listed in Table 2. Treatment 1 has a male manager in both sectors. Treatment 2 has a female manager in both treatments. Treatment 3 has a “sector-matching manager” in each sector, meaning that the gender of the manager matches the gender of the sector in both sectors. Workers in this treatment face a female manager in the office sector and a male manager in the janitorial sector. Lastly, treatment 4 has a “non-sector-matching manager” in each sector, meaning the gender of the manager does *not* match the gender of the sector in both sectors. Workers in this treatment face a female manager in the janitorial sector and a male manager in the office sector.

Table 2: Treatments in the Field Experiment

Treatment	Office Sector (F) Manager Gender	Janitorial Sector (M) Manager Gender
T1: All Male Managers	Male $\{m, o\}$	Male $\{m, j\}$
T2: All Female Managers	Female $\{m, o\}$	Female $\{f, j\}$
T3: Sector-Matching Managers	Female $\{f, o\}$	Male $\{m, j\}$
T4: Non-Sector-Matching Managers	Male $\{f, o\}$	Female $\{f, j\}$

The job advertisement is delivered from a “do not reply” e-mail address at the recruiting firm. The ad first either 1) thanks them for their interest if they provided their contact information or 2) tells them where their resume was found and states that they may be a good match for the position. Next, the ad provides detailed job and firm information. Last, the ad lists the two product sectors in bold with information specific to each sector and the name and e-mail address of the hiring manager listed under each sector heading. The order the sectors appear in the e-mail is randomized within treatment groups. Workers are informed that they may submit application materials to only one hiring manager. Appendix B contains the treatment e-mail.

All treatment information is included in the product sector section of the job

advertisement. First, I signal the gender associations of each product sector by describing the job titles for the types of customers with whom the prospective worker should be comfortable engaging. The ad describes the typical customers in the janitorial sector as “facility managers, operations, and maintenance workers.”⁹ Similarly, the ad describes the typical customers in the office sector as “office managers, receptionists, and administrative workers.”¹⁰ The job titles of typical customers are based on my previous experience in these sectors and were also validated by the sales managers at the firm.

I signal the gender of the manager in each product sector through the manager’s first name. The top two female and male names from 1975, along with the top four last names, in the U.S. are used as pseudonyms for real hiring managers at the firm. In addition to naming the manager to whom workers should send their application, the e-mail address of the manager is in the format `firstname.lastname@[recruiting firm name].com` to ensure salience of manager gender as signaled by the manager’s first name.

Workers submit a job application in two parts. First, they e-mail their resume and (optional) cover letter to the manager of their choice. Second, they complete a job application form on the recruiting firm’s website with education information and a brief job history. As a part of the application form, workers must provide the e-mail address of the manager to whom the application should be routed. In this way, participants must confirm in two places the manager type to which they wish to apply. Again, this process ensures the salience of manager type by forcing participants to look for and either write or copy and paste the manager’s name from the job advertisement. Workers who submit application materials to more than one manager are informed that their application materials cannot be processed and advised to resubmit their application materials to only one manager.

One week after the worker receives the treatment e-mail, I send an e-mail recruiting them to an online lab experiment with Vanderbilt University. I carefully ensure there are no similarities between the treatment e-mail and the recruitment e-mail so that workers do not know that there is any connection between the field and lab experiment. Such a connection might bias the lab results since monetary compensation in a experiment cannot compete with the high-stakes of a job application. In other words, if the worker believes that there is a non-zero (even if trivial) chance that their potential employer sees their responses on the lab experiment, the worker would report what they believe the employer wants rather than their true beliefs. Rather than assuring workers (truthfully) that the employer does not see responses to the lab experiment, or even if a worker completes the lab experiment, I avoid any connection between the field and lab portions of the experiment. Appendix C

⁹The 2018 Current Population Survey (CPS) finds that 96% of maintenance and repair workers, 68% of general and operations managers, and 66% of janitors and building workers are men.

¹⁰The 2018 CPS finds that 94% of secretaries and administrative assistants, 72% of first-line supervisors of office and administrative support, and 71% of office and administrative support workers are women.

contains the recruitment e-mail.

2.2 Lab Experiment

The goal of the lab experiment is to elicit workers’ second-order beliefs about the first-order beliefs of managers they face in the labor market regarding the relative productivity of women and men workers. Before discussing *how* I elicit those beliefs, I must carefully define *what* these beliefs of interest are. Incorrect or imprecise specifications of beliefs can lead to the elicitation of beliefs that do not accurately answer the research question of interest. For example, [Dustan et al. \(2020\)](#) shows that eliciting the difference in the beliefs about medians of two populations (as opposed to what I elicit: the belief about the median of the differences) could lead to the conclusion that people believe there are no differences in those populations, when in fact people believe there are large differences.

Let X^w be the distribution of sales productivity in the population of women workers and X^m be the distribution of sales productivity in the population of men workers. I want to elicit whether workers believe that a manager believes that $Pr(X^w > X^m) \geq 0.5$ or that $Pr(X^w < X^m) \geq 0.5$ (or both). Now, let $X^d = X^m - X^w$. Then, I want to learn whether workers believe that a manager believes that $Pr(X^d > 0) \geq 0.5$ or that $Pr(X^d < 0) \geq 0.5$. This property can be inferred from the median of X^d , which is the median of the distribution of differences in women’s and men’s sales productivity.

The idea that this property captures is intuitive. If a manager must choose between an equally qualified (on paper) woman and man, does the manager believe it is more likely that the woman or the man will be more productive in the sales job? The worker’s belief about the manager’s belief about this property then reveals whether the worker believes they are more or less likely to be hired, given the same resume as a worker of the opposite gender.

To elicit this property of workers’ subjective belief distributions, I adapt the procedure developed in [Dustan et al. \(2020\)](#) to measure first- and second-order beliefs about the difference in some measurable characteristic between two populations. To begin, the procedure requires two auxiliary samples: one to incentivize the first-order belief elicitation and one to incentivize the second-order belief elicitation. A small sample measuring productivity in sales is needed to elicit managers’ beliefs about the relative productivity of women and men. A small sample measuring managers’ first-order beliefs about relative productivity is needed to elicit workers’ beliefs about managers’ beliefs. These samples are used only to incentivize participants to truthfully reveal their beliefs and are not part of the results of this paper.

2.2.1 Auxiliary Sample Measuring Sales Productivity

The measurable characteristic of interest in this experiment is the productivity of workers in a sales job. I proxy this productivity with performance on a sales ability

assessment used by employers to screen workers. Sales assessments like the one I use are generally comprised of psychological questions in which the worker is asked to respond to a series of hypothetical situations. I obtained a small sample of workers who completed the assessment as a part of a job application from a firm that has tested more than 2 million workers.¹¹

Workers who choose to participate in the online lab experiment, henceforth “participants”, are told that the “sales assessment is used by employers to evaluate job candidates for positions in outside sales” and given information about the types of characteristics that are evaluated (e.g. “controls emotions and handles rejection”).¹² The metric I use from the sales assessment is the worker’s percentile rank in the population of test-takers.¹³ This measure has the benefit of being unit-free, so that the results can be interpreted without the potentially confounding issue of participants’ beliefs about the support of the measure.

2.2.2 Auxiliary Sample Measuring Managers’ First-Order Beliefs

The second sample measures the first-order beliefs of managers. Since I want to elicit the second-order beliefs directly relevant to a worker’s job application decision, I further refine this population to “managers responsible for hiring and supervising outside sales representatives.” I collected a small sample of first-order beliefs from managers at the partner firm using the elicitation procedure I describe next.¹⁴ The only difference is that I do not elicit managers’ second-order beliefs, just their first-order beliefs.¹⁵

2.2.3 Eliciting Participants’ First- and Second-Order Beliefs

The belief elicitation proceeds as follows. I elicit first-order beliefs by asking participants to report who they believe did better on the sales ability assessment—a randomly selected man or a randomly selected woman—and by how much. After the first-order belief elicitation, participants learn that managers responsible for hiring and supervising outside sales representatives answered the same question they just did, as described in Section 2.2.2. Participants are then asked to report what they believe a randomly drawn manager from each manager-type population chose,

¹¹The firm provided the results of the sales assessment for two anonymous women and two anonymous men. For a discussion of why only one draw from a population is necessary to incentivize truth-telling without deception, see [Dustan et al. \(2020\)](#).

¹²I do not provide participants with the assessment itself or even sample questions. My goal is to learn about beliefs regarding productivity in sales, not beliefs about how women and men perform on a particular assessment.

¹³Note that these are not any workers from the field experiment.

¹⁴I sampled two women and two men.

¹⁵I cannot elicit managers’ second-order beliefs at this point because I do not have a sample of first-order beliefs to incentivize truth-telling.

Figure 1: First-Order Belief Slider





Figure 2: Second-Order Belief Slider



which are their second-order beliefs. Manager types are described as “(gender) manager in a traditionally (gender) product sector.” The first-order belief elicitation is incentivized using the sample measuring worker productivity, described in Section 2.2.1, and the second-order belief elicitation is incentivized using the sample measuring managers’ first-order beliefs, described in Section 2.2.2.

Participants report their first-order belief using the slider in Figure 1 and report their second-order beliefs using the slider in Figure 2. The slider always begins in the center of the support, at gender neutrality. At this point on the slider, the text on the title bar with the grey background reads “I believe that they did the same” for the first-order belief slider and “I believe that the manager guessed they did the same” for the second-order beliefs slider. Participants move the slider to the left to report that (the manager guessed) the woman outperformed the man and move the slider to the right to report that the man outperformed the woman on the sales assessment. The text on the title bar of the slider updates at each point on the slider so that participants can see their guess in words.

As the participant moves the slider, the last row reporting the probabilities for winning the \$15 prize also update according to the Binarized Scoring Rule (BSR).¹⁶

¹⁶It remains the state-of-the-art at present despite recent concern that it may be subject to a “pull-to-center” effect (Danz, Vesterlund, & Wilson, 2020), meaning that participants report a belief somewhere between their true belief and the center of the support.

The defining feature of the BSR is that participants maximize their *probability* of winning a prize, rather than maximizing the *size* of the prize. This change in objective function means that the BSR is incentive-compatible for all expected utility maximizers as well as some non-expected utility maximizers. The critical assumption is *stochastic monotonicity*, meaning that the participant must prefer higher probabilities of winning the prize to lower probabilities, which covers a wide range of preferences.

The BSR works as follows. Let x be a draw of the random variable of interest X (in my case, the difference in percentile rank on the sales ability assessment between a randomly drawn man and a randomly drawn woman), θ the participant’s guess about x , and \bar{K} the number of points in the support of the distribution of X . Then, the participant’s probability of winning some prize is

$$\Pr(\text{Win the Prize}) = 1 - \frac{l(x, \theta)}{\bar{K}} \quad (1)$$

where $l(\cdot)$ is a loss function that maps the participant’s error to a value. The form of the loss function determines what *function* of the distribution of X is elicited by the payment rule. The prize in this experiment is \$15.

I choose the loss function $l(\cdot)$ to elicit the median. The median of the distribution of differences between women and men is the lowest cost, in terms of time and cognitive load, function to elicit the property $\Pr(X^d > 0) \geq 0.5$ or $\Pr(X^d < 0) \geq 0.5$. For example, eliciting the mode of the ternary distribution with outcomes “woman outperforms man”, “man outperforms woman”, and “woman and man perform the same” would require twice the number of elicitations and provide less relevant information. Eliciting the probabilities $\Pr(X^d > 0) \geq 0.5$ and $\Pr(X^d < 0) \geq 0.5$ directly would similarly require twice the number of elicitations to obtain the same information about the relative productivity of women and men.¹⁷ See [Dustan et al. \(2020\)](#) for a detailed discussion of the appropriate function to elicit to learn about beliefs about differences between two populations. The loss function for the median is the absolute error term, so the participant’s probability of winning the \$15 prize in this experiment is determined by:

$$\Pr(\text{Win the Prize}) = 1 - \frac{|x - \theta|}{\bar{K}} \quad (2)$$

The last factor to determine the payment structure is the number of points in the support over which participants can report their belief, \bar{K} . The choice of \bar{K} affects the incentives for the participant to be precise—the larger and denser the support, the lower the marginal cost of being one unit away from the outcome of the random

¹⁷Though eliciting the probabilities directly requires twice the number of elicitations, it also yields additional information that elicitation of the median does not. This additional information is not sufficiently useful, however, to warrant the higher costs of elicitation.

draw. At the same time, small or sparse supports could limit participants’ ability to express extreme beliefs or their ability to be precise if outcomes are binned. I choose $\bar{K} = 21$ to balance these factors, which corresponds to ten points of support on either side of gender neutrality.

The interactive slider communicates all information about the payment rule, without requiring the participant to understand the BSR or the equation that determines the payment probabilities.¹⁸ When the participant chooses a point on the slider, the last row of the chart updates to reflect the probability of winning the \$15 prize based on their selection, for *each possible outcome of the random draw*.¹⁹ The chart can be read as “If the actual outcome of the random draw is: (the woman did better by/the man did better by) x percentile points, you win \$15 with probability y ” where x is the number in the first row and y is the corresponding number in the last row. For example, in Figure 1, the participant has a 100% chance of winning the \$15 prize if the outcome of the random draw is exactly their guess that the man did better by 20 percentile points, a 95% chance of winning if the outcome is that the man did better by 15 percentile points, etc.

2.3 Implementation

The field experiment was implemented in two waves. The first wave ran from February to March 2020 and was prematurely halted due to the firm’s temporary hiring freeze in response to COVID-19. Recruiting for the online lab experiment continued through May 2020. The second wave ran from August to September 2020. Recruiting for the online lab experiment continued through October 2020.

Workers were recruited to the lab experiment through an e-mail invitation from Vanderbilt University to participate in a paid study. The study was advertised as a 10 to 15 minute online survey that paid \$5 for completion. Participants were told they could earn another \$15 based on their responses, as well as earn tickets into a drawing for a \$500 Amazon gift card. To begin the experiment, participants clicked on a link in the e-mail.

Instructions were both written and recorded as audio. Participants had to press “Play” and listen to the audio track on each page before they could proceed to the next page, or stay on the page for the length of the audio file plus 25 seconds. After the initial instructions, participants completed an example with feedback to ensure they understood how to report their beliefs and how they would be paid.

Participants began by reporting their first-order belief, followed by the four second-order belief elicitations with respect to each manager type—female in female

¹⁸Part of the innovation of the procedure developed by [Dustan et al. \(2020\)](#) was this method of implementing the BSR. The typical implementation taught participants the BSR or the underlying equation (see for example [Hossain and Okui \(2013\)](#), [Babcock, Recalde, Vesterlund, and Weingart \(2017\)](#), or [Dianat, Echenique, and Yariv \(2018\)](#)).

¹⁹Recall that the random draw is of a woman and a man who completed the sales assessment for the first-order belief, and of a manager of the relevant type for the second-order belief.

sector, female in male sector, male in female sector, and male in male sector.²⁰ Then, participants completed a number of other short tasks. These included, in order: unincentivized recall questions on the gender and sector of the last manager type for which they reported their second-order belief; incentivized beliefs about managers’ beliefs about customer preferences; incentivized beliefs about the traditional gender associations of a number of product sectors (including office and janitorial products); incentivized risk preference elicitation; unincentivized manager type preference elicitation; unincentivized confidence in own sales ability; and a number of survey questions covering employment status, demographics, and self-reported confidence in beliefs and risk preferences. All incentivized tasks were incentivized using tickets into the drawing for the \$500 Amazon gift card, with the exception of risk preferences. I elicited risk preferences using the method developed in [Eckel and Grossman \(2008\)](#), paying participants up to \$2.80.

The measure I use from the additional short tasks in the analysis apart from the demographics are the unincentivized manager type preference elicitation. I collect preferences for manager type by asking participants 1) which gender manager they prefer, and 2) which type of product sector they prefer.²¹ Participants could choose female, male, or no preference for each.

Table 3: Sample Sizes in the Field Experiment, by Treatment

	T1	T2	T3	T4	Total
Wave 1	538	537	535	533	2,143
<i>Women</i>	183	183	182	181	729
<i>Men</i>	336	334	332	331	1,333
<i>Unknown</i>	19	20	21	21	81
Wave 2	680	693	687	679	2,739
<i>Women</i>	268	273	274	262	1,077
<i>Men</i>	389	397	393	394	1,573
<i>Unknown</i>	23	23	20	23	89
Total	1,218	1,230	1,222	1,212	4,882

Notes: Columns reference the four treatment groups as listed in [Table 2](#). Rows reference subsamples for women, men, and people of unknown gender.

One reported belief out of the five elicited in total was randomly selected for payment at the end of the survey. Participants were informed of their guess and

²⁰The four second-order beliefs are elicited in random order.

²¹On the same page, I also ask participants about their preferences over workplace gender composition and provide an open comment box for them to make any notes on their preferences. The intention of the open comment box is to allow participants to reduce the cost of revealing gendered preferences by explaining them.

the random draw from the relevant population it was compared to in order to determine their probability of winning the \$15 prize. Participants then generated a random number between 0 and 100 using an embedded random number generator to determine if they won the prize. If their number was equal to or lower than their probability of winning, determined by equation 2, they won. After informing them of their total earnings, including tickets into the drawing for the \$500 Amazon gift card for their responses on the additional tasks, participants were directed to another site through a link to fill out a receipt. Payment was sent to participants through PayPal within 24 hours. Screenshots of the full experiment are included in Appendix D.

A total of 4,882 workers were treated in the field experiment, with 598 of those workers choosing to participate in the lab experiment. Table 3 reports the number of participants in each treatment of the field experiment by wave and by gender. Workers do not report their gender in the field experiment, so gender was predicted using the first names of all children born in the U.S. from 1964 to 1998. If more than 80% of children born with a given name were female, the worker was classified as a woman. If fewer than 20% of children born with a given name were female, the worker was classified as a man. 93% of workers could be classified using these rules. Another 3% of workers were classified using their photo,²² middle name,²³ or other clear marker of gender on their resume.²⁴ Workers who participated in the lab experiment self-reported their gender. Approximately 1% of workers in the field experiment were classified using this self-reported gender. The remaining 3% of workers were unable to be classified and dropped from the analysis.²⁵

Since I do not collect demographic information from participants in the field experiment, I cannot compare their characteristics to the subsample who participated in the lab experiment. Instead, I compare the demographic characteristics of participants in the lab experiment to those of wholesale and manufacturing sales representatives in the 2018 American Community Survey in the seven metropolitan statistical areas (MSAs) I recruit from in Table 4. My sample has more women, racial diversity, people who have never been married, and younger people with slightly more children.

3 Worker Beliefs and the Job Application Decision

This section presents the main result of the paper, which I obtain from combining the application choice data from the field experiment with the data on beliefs from

²²Photos could be found on resumes, ZipRecruiter profiles, and LinkedIn profiles, where the link to the worker’s LinkedIn profile was included on their resume.

²³Using the same classification strategy described for first names.

²⁴Examples include references to a men’s basketball scholarship and an award for “Top 10 Female in Sales.”

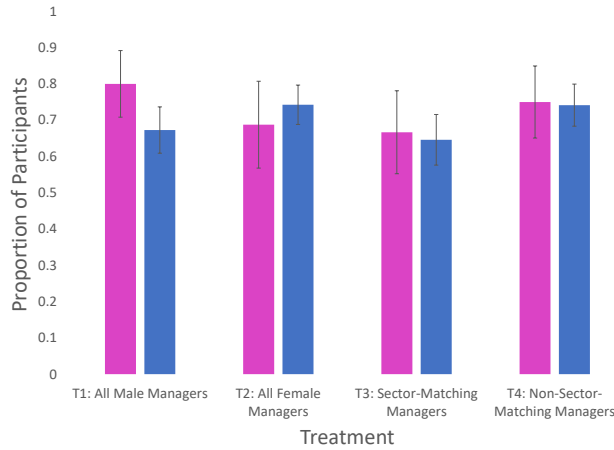
²⁵Self-reported gender in the lab replaced predicted gender in the field for 11 participants whose self-reported gender in the lab experiment contradicted predicted gender from the birth certificate algorithm.

Table 4: Comparison of Demographics in Lab Sample and in 2018 American Community Survey

	Experiment	2018 ACS	Difference
Woman	0.392 (0.020)	0.285 (0.017)	0.107*** (0.026)
White Non-Hispanic	0.677 (0.020)	0.892 (0.012)	-0.215*** (0.022)
Black/African-American	0.201 (0.017)	0.067 (0.010)	0.133*** (0.019)
Other Race/Ethnicity	0.105 (0.013)	0.041 (0.008)	0.064*** (0.014)
Married	0.446 (0.021)	0.693 (0.018)	-0.247*** (0.027)
Never Married	0.353 (0.020)	0.190 (0.015)	0.163*** (0.025)
Other Marital Status	0.201 (0.017)	0.117 (0.012)	0.084*** (0.020)
Age	40.966 (0.570)	47.373 (0.561)	-6.407*** (0.806)
Number of Children	1.135 (0.059)	0.800 (0.040)	0.335*** (0.069)
Number of Children <5yo	0.142 (0.020)	0.142 (0.016)	0.000 (0.026)
Wage Income			
\$0-\$25,000	0.171 (0.016)	0.244 (0.016)	-0.074*** (0.023)
\$25,000-\$50,000	0.231 (0.018)	0.175 (0.015)	0.055** (0.023)
\$50,000-\$75,000	0.258 (0.018)	0.175 (0.015)	0.082*** (0.023)
\$75,000-\$100,000	0.167 (0.016)	0.120 (0.012)	0.047** (0.020)
\$100,000-\$150,000	0.091 (0.012)	0.167 (0.014)	-0.076*** (0.019)
\$150,000+	0.083 (0.012)	0.118 (0.012)	-0.035** (0.017)
Observations	598	684	1282

Notes: Columns reference the lab experiment sample, the 2018 American Community Survey sample of wholesale or manufacturing sales representatives in one of the seven MSAs I recruited from in the experiment, and the differences in those two samples. Standard errors are in parentheses.

Figure 3: Proportion of Workers who Choose the First Manager Listed, Conditional on Applying



Notes: From left to right on each graph, the bars show the proportion of workers who applied to the first manager type listed, conditional on applying for the job, in treatment 1 with a male manager in both sectors, treatment 2 with a female manager in both sectors, treatment 3 with a sector-matching manager in both sectors, and treatment 4 with a non-sector-matching manager in both sectors. Standard error bars are robust.

the lab experiment. The question of interest is: how do workers' beliefs about managers' beliefs affect their job application decision? To answer it, I test the following hypothesis:

Hypothesis 1. *More favorable second-order beliefs about a manager type increase the likelihood that a worker applies for a job with that manager type.*

The specification to test this hypothesis is informed by a behavioral model of job choice I outline below, and by an empirical pattern that emerges from the field experiment: 71.10% ($SE = 2.62\%$) of workers who apply choose the first manager listed in the job. This proportion is significantly different from 50% ($p = 0.000$). Figure 3 shows that this proportion is consistent across treatment groups and gender. There are no significant differences between treatment groups and the proportions are significantly different from 50% in all treatments ($p = 0.000$ for treatments 1, 2, and 4, and $p = 0.007$ for treatment 3). There is also no significant difference ($p = 0.682$) in the probability a woman worker chooses the first manager (73%, $SE = 5.2\%$) and the probability a man worker chooses the first manager (70%, $SE = 3.0\%$).

Workers' application behavior underscores the importance of the order of options that has been recognized in a variety of fields, including psychology, marketing, and

political science.²⁶ For example, a “primacy effect” favoring the first option has been identified on election ballots (Van Erkel & Thijssen, 2016), lists of website links (Murphy, Hofacker, & Mizerski, 2006), and in wine tastings (Mantonakis, Rodero, Lesschaeve, & Hastie, 2009). This behavioral phenomenon has been modeled as a framing effect in economics (Salant & Rubinstein, 2008) and studied in the context of revealed preference as “limited attention” (Masatlioglu, Nakajima, & Ozbay, 2012).

Based on this result and the literature, I develop the following theoretical framework (its formalization is presented in Appendix E) to inform the empirical specification I use to test my main result. I model workers’ job application options in a dynamic setting under the assumption that workers’ (second-order) beliefs reflect their expectations about the probability of being hired. For simplicity, I assume workers’ beliefs about the hiring manager enter their utility directly; therefore, *ceteris paribus*, workers prefer to apply for a job with a manager type for whom they hold higher beliefs.

I assume that workers observe the first manager type, as defined by the manager’s gender and sector, listed in the job ad at no cost. After observing the first manager type, workers choose among the following three options: 1) apply to the first (observed) manager type, 2) pay a (psychic) cost to observe the second manager type (this cost could be motivated by a limited attention model), or 3) proceed to the next time period without applying. If the worker chooses option (2) by paying the cost to observe the second manager type, the worker chooses between applying to one of the two observed manager types or proceeding to the next time period without applying. Workers who proceed to the next time period choose between applying to another randomly drawn manager type or their reservation wage. This model informs the empirical specification detailed next.

3.1 Empirical Specification

The dependent variable of interest is the worker’s decision to apply for the job $Pr(\text{Apply})$. The model suggests that the worker’s second-order belief about the first manager type listed in the job ad b_i^{first} enters into the worker’s decision to apply to the first manager type, to observe the second manager type, to apply to the second manager type conditional on observing it, or to proceed to the next time period without applying; therefore, it is the main independent variable of interest.²⁷ According to Hypothesis 1, the more favorable beliefs the worker has about the beliefs of the first manager, the more likely the worker should be to apply to the job.

In addition to the primary belief of interest, the model suggests that two other measures of beliefs also affect workers’ decision. I refer to these beliefs as the “alternative” beliefs since they represent the worker’s dynamic decision. The first measure of alternative beliefs captures the possibility that the worker chooses to observe the

²⁶The order of options was randomized in the experiment for this reason.

²⁷Since beliefs are reported as man minus woman, I change the sign on the beliefs of women workers.

second manager type in the ad. As discussed above, workers who pay the cost of observing the second manager type choose between applying to one of the two manager types they have observed or not applying. *If* they choose to apply, the worker chooses the manager with the most favorable beliefs among the two manager types they faced; therefore, the first measure of alternative beliefs is the most favorable belief out of the two facing manager types, which I define as b_i^{\max} . That is, $b_i^{\max} = b_i^{\text{first}}$ if the first manager type is also the most favorable manager of the two types the worker faced, but $b_i^{\max} = b_i^{\text{second}}$ otherwise, where b_i^{second} is the worker’s belief about the belief of the second manager type.

The second measure of alternative beliefs I include is the average of all second-order beliefs except for the first manager listed b_i^{avg} . This measure controls for the alternatives the worker would face if they do not apply in the experimental period, unconditional on whether they observe the second manager type. A worker who faces a manager with a relatively high belief might still choose not to apply if the alternative manager types have even more favorable beliefs. Similarly, a worker who faces a manager with relatively low beliefs might still apply if the alternative manager types have even less favorable beliefs. In essence, controlling for the average means that the coefficient on the worker’s belief about the belief of the first manager captures how more favorable beliefs *relative to* the beliefs of the other manager types affect the worker’s application decision.

To causally interpret the relationship between second-order beliefs and the job application decision, I must control for any factor that affects job application behavior *and also varies with the randomly assigned manager types*. Manager gender-specific preferences or sector-specific preferences may be such a factor; therefore, I control for the vector of preferences ρ_i elicited in the lab experiment in the estimating equation. There are four indicator variables corresponding to participants reporting 1) preferences for the gender of the first manager, 2) preferences for the opposite-gender of the first manager, 3) preferences for the sector gender of the first manager, and 4) preferences for the opposite-gender sector of the first manager.

In summary, I estimate the relationship between the second-order belief of interest b_i^{first} and choosing to apply for the job $Pr(\text{Apply})$ separately for women and men using the equation:

$$Pr(\text{Apply}) = \alpha \mathbf{T}_i + \beta_1 b_i^{\text{first}} + \beta_2 b_i^{\max} + \beta_3 b_i^{\text{avg}} + \theta \rho_i + \gamma \mathbf{X}_i + \epsilon_i \quad (3)$$

\mathbf{T}_i is a vector of all treatment indicator variables and the vector of controls \mathbf{X}_i includes the worker’s recruitment wave, city, and job board, as well as whether or not the worker reports actively searching for a job in the lab experiment. Note there is no constant, so the vector of coefficients on \mathbf{T}_i reports average application rates in each treatment, conditional on the other variables. I estimate a linear probability model in my preferred specifications. Appendix G contains estimates of a Probit model, with the same qualitative results.

The descriptive statistics for each of the main variables are listed in Table 5. The probability of applying and the preference parameters are binary variables. As discussed in Section 2.2, beliefs are reported about the median of the distribution of differences between women and men in percentile rank on the sales ability assessment. Positive values correspond to beliefs that favor the worker’s own gender and negative values correspond to beliefs that favor the opposite-gender of the worker. So, on average, women in the sample believe that the median first manager type listed in their job ad reports that the median difference between women and men is three percentile rank points in favor of women. Men, on average, believe that the median first manager type listed in their job ad reports that the median difference is 1 percentile rank point in favor of women (because the sign is negative).

Table 5: Descriptive Statistics


	Women	Men
Pr(Apply)	0.076 (0.266)	0.179 (0.384)
Belief about First Manager	3.296 (29.054)	-1.254 (25.554)
Most Favorable Belief of Facing Managers	18.879 (22.165)	15.299 (20.605)
Average Belief about Managers, except First	-0.045 (14.946)	0.204 (11.956)
Prefer First Manager’s Gender	0.251 (0.435)	0.217 (0.412)
Prefer Opposite-Gender of First Manager	0.283 (0.451)	0.191 (0.394)
Prefer First Manager’s Sector	0.184 (0.388)	0.197 (0.398)
Prefer Opposite Sector of First Manager	0.161 (0.369)	0.225 (0.418)
Observations	223	351

Notes: Columns reference the sample of women and the sample of men who participated in the lab experiment. Standard deviations are in parentheses.

3.2 Results

I find that second-order beliefs about the first manager positively and significantly affect the worker’s job application decision through the random assignment of the first manager type listed in the job ad for men, but not women. Figure 10 reports the coefficients and their 95% confidence intervals in four specifications of equation 3, progressively including more explanatory variables. The *Base Model* includes only the belief about the first manager b_i^{first} and the vector of treatment indicator variables. The *Add Alternatives* specification adds b_i^{max} and b_i^{avg} . The *Add Alternatives & Preferences* specification adds the vector of preference variables ρ_i . Lastly, the *Add Alternatives, Preferences, & Controls* specification adds the vector of controls \mathbf{X}_i .

Figure 4: Regression Coefficient on Belief about First Manager by Gender, Main Specification



Notes: 95% confidence intervals are shown. Standard errors are robust. Coefficient is scaled to a one unit increase in the favorability of beliefs. Panel (a) reports the coefficient on the regression with men only and panel (b) reports the same for women only. In each graph, the first estimate includes only the vector of treatment variables and the belief of interest. The second estimate adds measures of alternative beliefs, the third estimate adds preferences, and the fourth estimate includes controls for the wave, city, and job board in which the worker was recruited and whether the worker reports actively searching for a job.

The coefficients on the worker’s belief is positive and statistically significant at the 5% level for men in all specifications. In the full specification, a one unit increase in the worker’s belief about the belief of the first manager listed in the job ad increases the probability a man applies by 1.28 percentage points ($SE = 0.53$). This effect is

approximately 16% of the base application rate of 7.8% for men workers, indicating that the effect is also economically significant.

On the other hand, I find no evidence of a relationship between women’s second-order belief and their job application decision. A one unit increase in beliefs changes the probability a woman worker applies between $[-0.67, 0.94]$ percentage points in the full specification. While the sample size is smaller for women (220 compared to 341 men in the full specification), the coefficient is also an order of magnitude smaller for women compared to men (0.14 versus 1.28 percentage points). As I discuss in more detail in Section 6, this null result could be due to the selection of women into the male-dominated occupation for which I recruit.

Result 1. *More favorable second-order beliefs about a manager type increase the likelihood that a man worker applies for a job with that manager type. I find no evidence that second-order beliefs affect the likelihood a woman worker applies for a job.*

Regression charts are contained in Appendix F. In the main specification, the vector of treatment variables is statistically significant for men in all but the full specification including controls. No other variables are statistically significant for men.

Similarly, the vector of treatment variables is statistically significant in all but the full specification with controls for women.²⁸ In addition, for women, preferences for the gender of the first manager are negative and statistically significant with a relatively large coefficient. Further investigation shows this effect is driven by women with male managers first, suggesting that even though women report a preference for working with a male manager, they may be less likely to apply to a male hiring manager.²⁹ Lastly, average beliefs is positive and statistically significant for women in the full specification with controls. This coefficient suggests that women with overall more favorable beliefs are more likely to apply.


3.3 Robustness to Alternative Specifications

This result is robust to alternative specifications. First, I control for the first manager listed interacted with treatment. That means I include eight treatment variables, instead of four, corresponding to both the treatment group and the first manager listed, in the specification in equation 3. Figure 11 shows that the size of the coefficient is the same for men. The estimate is noisier since I lose degrees of freedom with little gain in explanatory value, but the coefficient remains significant at the 10% level. There is still no detectable effect of beliefs on women’s job application decision.

²⁸With one minor exception, treatment 3 is also not statistically significant once preferences are added to the specification.

²⁹Recall that the preferences I elicit in the lab experiment are coarse and unincentivized.

Figure 5: Regression Coefficient on Belief about First Manager by Gender, Other Specifications



Notes: 95% confidence intervals are shown. Standard errors are robust. Coefficient is scaled to a one unit increase in the favorability of beliefs. Panel (a) reports the coefficient on the regression with men only and panel (b) reports the same for women only. In each graph, specification (1) interacts treatment with first manager type. Specification (2) adds first-order beliefs. Specification (3) replaces the most favorable belief of the two facing managers with the belief about the second manager type. Specification (4) replaces the most favorable belief of the two facing managers and the average belief of all other managers, except the first, with the most favorable belief of all managers, excepting the first. Specification (5) replaces the average belief of all other managers, except the first, with the most favorable belief of the non-facing managers. Specification (6) replaces the most favorable belief of the two facing managers and the average belief of all other managers, except the first, with the belief about the second manager type and the average belief of the non-facing managers.

Second, I add first-order beliefs. If workers' beliefs about the relative productivity of women and men affect their job application decision directly, and second-order beliefs are informed by workers' first-order beliefs, the effect of second-order beliefs on the job application decision could be confounded. I find no evidence this is the case. The coefficient remains the same and statistically significant at the 5% level.

I also consider other specifications of the alternative beliefs. The main specification is determined by a theoretical model because there are many reasonable ways to control for the alternatives in the worker's dynamic application decision. I show that the results are robust to many other specifications of the alternative beliefs.

First, I replace the most favorable belief of the two facing managers with the belief about the second manager listed in the job ad. The former alternative assumes that

the worker’s belief about the second manager type only matters if it is more favorable than the worker’s belief about the first manager type, while the latter allows the worker’s belief about the second manager type to enter the worker’s decision even if it is less favorable. Second, I replace both alternative beliefs with the most favorable belief out of all of the manager types that are not the first manager listed. The idea of this alternative belief is that the worker decides only between the first manager type and taking another draw, implicitly assuming the worker does not consider the second manager type. Third, I replace the worker’s average belief about all manager types except the first with the most favorable belief about the non-facing manager types. This alternative treats the worker’s decision as a choice between the most favorable of the two manager types the worker faces in the experiment and the most favorable of the two manager types the worker does *not* face in the experiment. Lastly, I consider a combination of the alternatives by replacing the most favorable belief about the two facing managers with the belief about the second manager *and* replacing the average of all other managers, excepting the first, with the average of just the two non-facing managers. This specification assumes the worker chooses between the first manager type, the second manager type, or taking another draw in expectation of the two manager types the worker does not face in the experiment.

Figure 11 shows that different ways of controlling for the alternative beliefs has little effect on the estimated coefficient on worker’s beliefs about the first manager type listed for men. Moreover, the coefficient remains significant at the 5% level in all specifications. On the other hand, the estimated coefficient for women remains small and not statistically distinguishable from zero, though it does change sign to be negative in these other specifications.

4 Field Experiment Results

This section describes results from the component field experiment that are policy-relevant independent of the main result. In the field experiment, I observe workers’ job application decisions after randomly assigning workers to a treatment that determines the gender of the manager in each of two sectors: the office products sector (female) and the janitorial products sector (male). I specify three hypotheses in the pre-analysis plan for the field experiment data, which I test in this section.

Hypothesis 2. *Men workers are more likely to apply than women workers.*

Hypothesis 3. *Workers are more likely to apply to their same-gender product sector.*

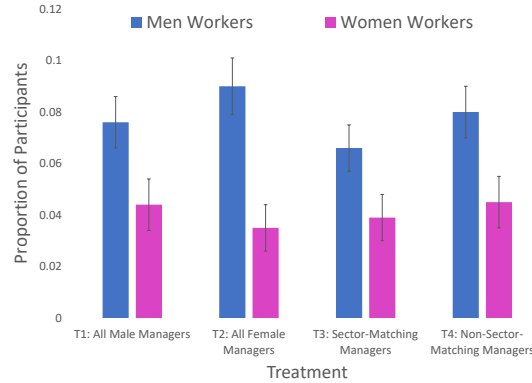
Hypothesis 4. *Workers are more likely to apply to their same-gender manager.*

I find strong evidence in favor of Hypothesis 2 that men workers are more likely to apply than women workers. Men are almost twice as likely to apply for a job (7.8%, $SE = 0.50\%$) compared to women workers (4.1%, $SE = 0.47\%$) and this

difference is statistically significant ($p = 0.000$).³⁰ Moreover, Figure 6 shows that women are less likely to apply than men in every treatment group ($p = 0.0158$, $p = 0.000$, $p = 0.025$, $p = 0.010$ for treatments 1, 2, 3, and 4 respectively).

Result 2. *Men workers are more likely to apply to the job than women workers, overall and in each treatment.*

Figure 6: Proportion of Participants in each Treatment who Applied for the Job



Notes: From left to right on each graph, the bars show the proportion of participants who applied for the job in treatment 1 with a male manager in both sectors, treatment 2 with a female manager in both sectors, treatment 3 with a sector-matching manager in both sectors, and treatment 4 with a non-sector-matching manager in both sectors. Standard error bars are robust.


I find no evidence in favor of Hypothesis 3 that workers are more likely to apply to their same-gender sector. Figure 7 shows that both women (57%, $SE = 5.8\%$) and men (53%, $SE = 3.3\%$) are slightly more likely to choose the office sector than the janitorial sector, conditional on applying, but these proportions are not statistically different from 50% ($p = 0.123$ for women and $p = 0.840$ for men).³¹ Unconditional on applying, I find no statistically significant differences between the probability a woman applies to the office sector (2.3%, $SE = 0.35\%$) compared to the janitorial sector (1.8%, $SE = 0.31\%$), nor the probability a man applies to the janitorial sector (3.7%, $SE = 0.35\%$) compared to the office sector (4.2%, $SE = 0.37\%$) ($p = 0.120$ for women and $p = 0.840$ for men).

Result 3. *I find no evidence that workers are more likely to choose their same-gender sector.*

³⁰p-values are reported for one-sided tests of proportions unless otherwise noted in this subsection since the pre-specified hypotheses are one-sided.

³¹For men, I use the one-sided test on the probability they apply to the janitorial sector since that was the pre-specified hypothesis. The probability men apply to the office sector is also not statistically different from 50% using a two-sided test ($p = 0.160$).

Figure 7: Proportion of Participants in each Treatment who Chose the Janitorial Sector, Conditional on Applying



Notes: From left to right on each graph, the bars show the proportion of participants who applied to the janitorial sector, conditional on applying for the job, in treatment 1 with a male manager in both sectors, treatment 2 with a female manager in both sectors, treatment 3 with a sector-matching manager in both sectors, and treatment 4 with a non-sector-matching manager in both sectors. Standard error bars are robust.

I cannot do a direct comparison of application rates for a female versus a male manager to test Hypothesis 4 that workers are more likely to apply to a same-gender manager because there are no treatments that hold the sector constant and vary the gender of the manager. Rather, I specify the following tests in the pre-analysis plan.

1. The probability a woman worker applies to the office sector is higher when there is a female compared to a male manager in the office sector, and a male manager in the janitorial sector.
2. The probability a man worker applies to the janitorial sector is higher when there is a male compared to a female manager in the janitorial sector, and a female manager in the office sector.
3. The probability a woman worker applies to a female manager, conditional on a female manager being available, is higher than the probability she applies to a male manager, conditional on a male manager being available.
4. The probability a man worker applies to a male manager, conditional on a male manager being available, is higher than the probability he applies to a female manager, conditional on a female manager being available.

I find no evidence in support of Hypothesis 4 for women using either test 1 or test 3. For test 1, women are no more likely to apply for a job with a female manager

than a male manager in the office sector when the alternative is a male manager in the janitorial sector ($p = 0.594$).³² In fact, the point estimate is slightly in favor of the male manager (2.7%, $SE = 0.76\%$ apply to the male manager compared to 2.4%, $SE = 0.72\%$ who apply to the female manager). Furthermore, test 3 finds that women are no more likely to apply for a job with a female manager when a female is available compared to a male manager when a male is available ($p = 0.599$) and the point estimate is again slightly in favor of male managers (2.8%, $SE = 0.44\%$ apply to a male manager, when available, compared to 2.6%, $SE = 0.45\%$ who apply to a female manager, when available).

Evidence is strongly in favor of the *opposite* conclusion from Hypothesis 4 for men workers. For test 2, men are statistically more likely to apply for a job with a female manager (4.2%, $SE = 0.75\%$) than a male manager (2.2%, $SE = 0.55\%$) in the janitorial sector when the manager in the office sector is female ($p = 0.029$ in a two-sided test). Moreover, test 4 finds that men are more likely to apply to a female manager when a female is available (76%, $SE = 3.3\%$) compared to a male manager when a male is available (59%, $SE = 3.9\%$) and this difference is statistically significant ($p = 0.018$ in a two-sided test).

Result 4. *I find no evidence that workers are more likely to apply to their same-gender manager. Men workers are more likely to apply to a female manager than a male manager.*

5 Lab Experiment Results

The component lab experiment provides independent evidence on workers' second-order beliefs that have policy implications beyond the main results described in Section 3. In the lab experiment, I elicit workers' beliefs about the beliefs of managers who are relevant to their job application decisions in the field experiment. I first present results on the primary beliefs of interest in Section 5.1: second-order beliefs. Then, I show the results for workers' first-order beliefs in Section 5.2 and their relationship to second-order beliefs, or intra-participant beliefs, in Section 5.3.

5.1 Evidence on Second-Order Beliefs

Average second-order beliefs, reported by manager type and by gender in Table 6, show that participants believe different manager types hold different beliefs.³³ I

³²Test 1 (and test 2 for men) was designed with the idea that workers would be most responsive to the availability of a same-gender manager in their same-gender sector. Alternative tests, such as whether women are responsive to the availability of a same-gender manager in the opposite-gender sector, are also null.

³³As discussed in Section 2, the form of the payment function dictates that participants optimally report the median of their subjective belief distributions, so the means in Table 6 are the averages of the medians reported by participants.

reject equality of the four averages using the sample of all participants in column (1) ($p = 0.000$ for all tests).³⁴ The differences in participants’ beliefs about the beliefs of different manager types is further evidenced by Figure 8, which aggregates beliefs into a ternary outcome. I reject equality of the four ternary belief distributions ($p = 0.000$ for all tests using the Wilcoxon signed-rank test).^{35,36} These differences are not reflected in differences between women’s and men’s second-order beliefs. Table H shows that women report slightly more extreme beliefs than men about the sector-matching managers (less than one unit on the 21-point scale); however, there are no differences in women’s and men’s beliefs about non-sector-matching manager types.

Table 6: Average Second-Order Beliefs

Sample	(1) All	(2) Women	(3) Men	(4) Difference
Female Manager in Female Sector	-17.352 (0.897)	-20.426 (1.438)	-15.440 (1.139)	4.986*** (1.833)
Female Manager in Male Sector	0.634 (0.963)	-0.157 (1.668)	1.122 (1.170)	1.279 (1.981)
Male Manager in Female Sector	-5.781 (1.020)	-5.404 (1.722)	-6.037 (1.264)	-0.633 (2.098)
Male Manager in Male Sector	20.712 (0.890)	22.825 (1.573)	19.432 (1.055)	-3.393* (1.824)
Observations	576	223	352	575

Notes: Columns (1) to (3) reference subsamples. Column (4) reports the differences between the women and men subsamples. Standard errors are reported in parentheses underneath the means.

I next consider whether workers with the second-order beliefs of participants in this lab experiment would hypothetically sort by gender in the labor market if they acted exclusively on their second-order beliefs. That is, do participants believe that female managers and managers in the female sector favor women? And do they believe that male managers and managers in the male sector favor men?

Participants’ beliefs are consistent with sorting by gender behavior. Most participants believe that sector-matching manager types favor the manager’s own gender.

³⁴All tests in this subsection are two-sided t-tests or tests of proportions, as appropriate.

³⁵The Wilcoxon signed-rank test for matched pairs accounts for intra-participant dependence, meaning that the samples for each distribution tested are not independent of each other because they are composed of the same participants.

³⁶I also reject the equality of the four cardinal distributions, which are reported in Appendix H (again, $p = 0.000$ for all tests).

Figure 8: Ternary Distributions of Second-Order Beliefs



Notes: From left to right on the graphs, the bars show the proportion of participants who report a belief that *the most likely* outcome of a random draw of the relevant manager type is that the manager reports the woman outperforms the man, the woman performs the same as the man, the man outperforms the woman on the sales assessment.

76% ($SE = 1.8\%$) of participants believe that *most* female managers in female sectors believe that the most likely outcome of a random draw is that the woman outperforms the man, while 79% ($SE = 1.7\%$) believe that *most* male managers in male sectors believe the opposite, that the most likely outcome of a random draw is

that the man outperforms the woman.³⁷

Non-sector-matching managers mitigate sorting by gender behavior. Only 56% ($SE = 2.1\%$) of participants believe that most male managers in female sectors believe that the most likely outcome of a random draw is that the woman outperforms the man. Similarly, just 46% ($SE = 2.1\%$) of participants believe that most female managers in male sectors believe that the most likely outcome is that the man outperforms the woman.

5.2 Evidence on First-Order Beliefs

I find that participants’ first-order belief distribution is centered around zero. Table 7 reports that, on average, participants reveal a belief that the median percentile rank difference between a man and a woman is approximately one percentile point in favor of the woman (-0.77 , $SE = 0.91$). I cannot reject the null hypothesis that the average first-order belief is zero ($p = 0.398$).

Table 7: First-Order Beliefs

	All	Women	Men	Difference
Average Belief	-0.773 (0.914)	-2.152 (1.584)	0.099 (1.108)	2.252 (1.878)
Ternary Beliefs				
W>M	0.465 (0.021)	0.475 (0.034)	0.460 (0.027)	-0.015 (0.043)
W=M	0.148 (0.015)	0.135 (0.023)	0.153 (0.019)	0.019 (0.030)
W<M	0.387 (0.020)	0.390 (0.033)	0.386 (0.026)	-0.004 (0.042)
Observations	576	223	352	575


Notes: Columns (1) to (3) reference subsamples. Standard errors are reported in parentheses underneath the means. Column (4) reports the differences between the women and men subsamples. The rows labeled “W>M”, “W=M”, and “W<M” report the proportion of participants who believe that the most likely outcome of a random draw is the woman outperforms the man, the woman performs the same as the man, and the man outperforms the woman respectively.

The ternary first-order belief distribution shown in Figure 9 and reported in

³⁷This interpretation is derived from the use of the median. A participant who reports a negative median believes there is at least a 50% probability that the manager reports a belief that the woman outperforms the man. Similarly, a participant who reports a positive median believes there is at least a 50% probability that the manager reports a belief that the man outperforms the woman.

Table 7 shows that the gender neutrality suggested by the central tendency of the distribution conceals considerable heterogeneity in beliefs.³⁸ 85% ($SE = 1.5\%$) of participants believe there is some difference in women’s and men’s sales ability; however, they do not agree on which gender outperforms the other. Participants are more likely to believe that the most likely outcome of a random draw is the woman outperforms the man (46%, $SE = 2.1\%$) than to believe the opposite, that the most likely outcome is the man outperforms the woman (39%, $SE = 2.0\%$), and this difference is statistically significant ($p = 0.007$).

Figure 9: Ternary Distribution of First-Order Beliefs



Notes: From left to right on the graph, the bars show the proportion of participants who believe that the most likely outcome of a random draw is the woman outperforms the man, the woman performs the same as the man, and the man outperforms the woman on the sales assessment.

I find no statistically distinguishable differences in the first-order beliefs reported by women and men participants. Table 7 shows that there are no differences in the average cardinal beliefs nor ternary belief distribution outcomes. I also cannot reject at conventional significance levels that women’s and men’s first-order cardinal ($p = 0.297$) or ternary ($p = 0.873$) belief distributions are identical using the Wilcoxon rank-sum test.

5.3 Intra-Participant Beliefs

To evaluate how participants’ first-order beliefs compare to their second-order beliefs about managers’ beliefs, I calculate intra-participant beliefs by taking the difference between the participant’s first-order belief and their second-order belief about the

³⁸Cardinal belief distributions in Appendix H also demonstrate this heterogeneity

relevant manager type. For women, their first-order beliefs are compared to their second-order beliefs about female managers in each sector while, for men, their first-order beliefs are compared to their second-order beliefs about male managers in each sector. The intuition for focusing on a comparison with same-gender manager types is to compare participants’ first-order beliefs with their second-order beliefs about managers “like them.” Since differences are taken as first-order belief minus second-order belief (and beliefs are reported as man minus woman), positive figures indicate that the participant believes the manager’s belief favors men more than the participant’s own belief.

Table 8 shows that participants believe that they share similar first-order beliefs with managers of their same-gender in the *opposite-gender* sector, but that their same-gender managers’ beliefs in the *same-gender* sector favor the opposite gender much more than participants do themselves. Women believe that the beliefs of female managers in the female sector favor men 18.27 percentile points ($SE = 2.06$) more than they themselves do ($p = 0.000$). On the other hand, I find no statistically distinguishable difference between women’s first-order beliefs and their beliefs about female managers in the *male* sector ($p = 0.342$). A similar pattern emerges for men. Men believe that the beliefs of male managers in the male sector favor women 19.33 percentile points ($SE = 1.42$) more than they do ($p = 0.000$), but believe that the beliefs of male managers in the female sector favor women only marginally more than themselves (2.00, $SE = 2.09$; $p = 0.000$).

Table 8: Intra-Participant Cardinal Beliefs

	All	Women	Men	Differences
Female Sector	10.843 (1.318)	18.274 (2.063)	6.136 (1.665)	-12.137 (2.660)
Male Sector	-12.609 (1.242)	-1.996 (2.094)	-19.332 (1.424)	-17.337 (2.445)
Observations	575	223	352	575

Notes: Columns (1) to (3) reference subsamples. Standard errors are reported in parentheses underneath the means. Column (4) reports the differences between the women and men subsamples.

6 Discussion

In this paper, I combine a large-scale natural field experiment with a structured on-line lab experiment to show that workers’ beliefs about managers’ beliefs affect job application behavior, but I find this is true only for men workers, not women work-

ers. This result raises the question of why women in the experiment are unresponsive to their second-order beliefs, particularly since they *believe* that most managers in the male sector, and male managers more generally, favor men in their beliefs. One explanation is that women who choose male-dominated occupations, like the sales job for which I recruit workers, may be selected based on the fact that they do not respond to their second-order beliefs. These women may have preferences such that choosing a non-conforming job or succeeding in adverse conditions overrides expectations about being hired. This experiment is not designed to study the selection of women into the occupation I recruit for, but future work should consider beliefs in a female-dominated occupation to determine if the opposite pattern arises there.

Despite the null results for women workers, the sorting patterns of men workers based on their second-order beliefs have important implications for the gender wage gap. Men believe that most male managers, and managers in male sectors, favor men. Sorting based on these beliefs increases the supply of workers to male occupations and industries relative to female occupations and industries, potentially crowding out women. Moreover, selection out of female occupations and industries is itself important if this sorting behavior leads to skill mismatch, where men who would otherwise choose female-dominated jobs choose not to because of the higher likelihood of female managers.

In addition to the main result showing that men’s beliefs about managers’ beliefs affects their job application decision, I observe interesting behavior in the field experiment. First, men are almost twice as likely to apply as women, even though they are selected into the field experiment in the same way. This result could mean that women have a higher cost of applying or lower expectations of being hired and/or wages to overcome that cost. Policy aimed at recruiting women to the applicant pool may be an effective tool to increase representation in men-dominated jobs.

Second, I find no evidence that women or men sort based on the sector gender. The beliefs I elicit in the lab experiment are explicitly characterized with respect to “traditionally female” or “traditionally male” sectors; however, the signal of sector gender in the field experiment may not be completely effective. It is also possible that, once sorted into a men-dominated (or women-dominated) occupation, workers are less likely to further sort within that occupation.

Third, I find that women are unresponsive to manager gender, but men are more likely to apply to their *opposite-gender* manager. The null results for women with respect to manager gender is in line with the results of [Flory et al. \(2015\)](#) finding that workers selected into a men-dominated occupation are unresponsive to manager gender. This result is also consistent with the unresponsiveness of women to their second-order beliefs. On the other hand, men’s higher likelihood for applying to female managers is unexplained by the experiment. The fact that men are more likely to apply to female managers does not conflict with the main result that men act on their second-order beliefs, but it does indicate that there is another factor contributing to the job application decision.

The lab experiment also provides new insights. First, I find that workers' beliefs are consistent with sorting by gender behavior, a result that provides support for other literatures studying mechanisms that require workers to believe that different types of managers hold different beliefs. Moreover, the fact that workers overwhelmingly believe that sector-matching managers favor their own gender, but have less homogeneous beliefs about non-sector-matching managers, suggests a role for policies like quotas and affirmative action. Managers whose gender is not traditionally associated with a sector may attract more workers of their gender.

Second, I find few differences in women's and men's beliefs. Women report slightly more extreme second-order beliefs about sector-matching managers, but the size of the difference is quite small. Moreover, I find no differences in either second-order beliefs about non-sector matching managers or in first-order beliefs. This latter result may seem surprising, but it is consistent with the results in [Dustan et al. \(2020\)](#) and other studies. Women and men generally have the same beliefs about the differences between women and men, suggesting that factors influencing first-order beliefs are not gender-specific.

Third, women and men believe that their beliefs are reflected by managers of their same gender in the *opposite-gender* product sectors. For both women and men, participants believe that managers of their same-gender in gender-matching sectors have more extreme beliefs about the participant's gender. This result underscores the extreme beliefs that participants believe that sector-matching managers hold.

In summary, I provide evidence on an important new mechanism affecting workers' job application decision. I study this mechanism within the microcosm of one occupation; however, my results suggest that second-order beliefs could be driving broader sorting patterns in the labor market. Workers with the beliefs I elicit would be more likely to choose occupations and industries that match their own gender, contributing to the sorting by gender behavior underlying a large part of the gender wage gap.

References

- Alston, M. (2019). The (perceived) cost of being female: An experimental investigation of strategic responses to discrimination. *Working paper*.
- Amuedo-Dorantes, C., & De la Rica, S. (2006). The role of segregation and pay structure on the gender wage gap: evidence from matched employer-employee data for Spain. *The BE Journal of Economic Analysis & Policy*, 5(1).
- Babcock, L., Recalde, M. P., Vesterlund, L., & Weingart, L. (2017). Gender differences in accepting and receiving requests for tasks with low promotability. *American Economic Review*, 107(3), 714–47.
- Barbulescu, R., & Bidwell, M. (2013). Do women choose different jobs from men? mechanisms of application segregation in the market for managerial workers. *Organization Science*, 24(3), 737–756.

- Bayard, K., Hellerstein, J., Neumark, D., & Troske, K. (2003). New evidence on sex segregation and sex differences in wages from matched employee-employer data. *Journal of Labor Economics*, 21(4), 887–922.
- Blau, F. D., & Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, 55(3), 789–865.
- Card, D., Cardoso, A. R., & Kline, P. (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly Journal of Economics*, 131(2), 633–686.
- Cardoso, A. R., Guimarães, P., & Portugal, P. (2016). What drives the gender wage gap? a look at the role of firm and job-title heterogeneity. *Oxford Economic Papers*, 68(2), 506–524.
- Carrington, W. J., & Troske, K. R. (1998). Sex segregation in us manufacturing. *ILR Review*, 51(3), 445–464.
- Danz, D., Vesterlund, L., & Wilson, A. J. (2020). *Belief elicitation: Limiting truth telling with information on incentives* (Tech. Rep.). National Bureau of Economic Research.
- Dianat, A., Echenique, F., & Yariv, L. (2018). Statistical discrimination and affirmative action in the lab.
- Dustan, A., Koutout, K., & Leo, G. (2020). Beliefs about beliefs about gender. *Available at SSRN*.
- Eckel, C. C., & Grossman, P. J. (2008). Forecasting risk attitudes: An experimental study using actual and forecast gamble choices. *Journal of Economic Behavior & Organization*, 68(1), 1–17.
- Fafchamps, M., Söderbom, M., & Benhassine, N. (2009). Wage gaps and job sorting in african manufacturing. *Journal of African Economies*, 18(5), 824–868.
- Fang, H., & Moro, A. (2010). *Theories of statistical discrimination and affirmative action: A survey*. Handbook of Social Economics.
- Fernandez, R. M., & Campero, S. (2017). Gender sorting and the glass ceiling in high-tech firms. *Ilr Review*, 70(1), 73–104.
- Fernandez, R. M., & Friedrich, C. (2011). Gender sorting at the application interface. *Industrial Relations: A Journal of Economy and Society*, 50(4), 591–609.
- Fernandez, R. M., & Mors, M. L. (2008). Competing for jobs: Labor queues and gender sorting in the hiring process. *Social Science Research*, 37(4), 1061–1080.
- Flory, J. A., Leibbrandt, A., & List, J. A. (2015). Do competitive workplaces deter female workers? a large-scale natural field experiment on job entry decisions. *The Review of Economic Studies*, 82(1), 122–155.
- Fortin, N. M., & Huberman, M. (2002). Occupational gender segregation and women’s wages in canada: An historical perspective. *Canadian Public Policy/Analyse de Politiques*, S11–S39.
- Gallen, Y., Lesner, R. V., & Vejlin, R. (2019). The labor market gender gap in denmark: Sorting out the past 30 years. *Labour Economics*, 56, 58–67.

- Glover, D., Pallais, A., & Pariente, W. (2017). Discrimination as a self-fulfilling prophecy: Evidence from french grocery stores. *The Quarterly Journal of Economics*, *132*(3), 1219–1260.
- Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, *104*(4), 1091–1119.
- Groshen, E. L. (1991). The structure of the female/male wage differential: Is it who you are, what you do, or where you work? *Journal of Human Resources*, 457–472.
- Hossain, T., & Okui, R. (2013). The binarized scoring rule. *Review of Economic Studies*, *80*(3), 984–1001.
- Kneeland, T. (2015). Identifying higher-order rationality. *Econometrica*, *83*(5), 2065–2079.
- Manski, C. F., & Neri, C. (2013). First-and second-order subjective expectations in strategic decision-making: Experimental evidence. *Games and Economic Behavior*, *81*, 232–254.
- Mantonakis, A., Rodero, P., Lesschaeve, I., & Hastie, R. (2009). Order in choice: Effects of serial position on preferences. *Psychological Science*, *20*(11), 1309–1312.
- Masatlioglu, Y., Nakajima, D., & Ozbay, E. Y. (2012). Revealed attention. *American Economic Review*, *102*(5), 2183–2205.
- Murphy, J., Hofacker, C., & Mizerski, R. (2006). Primacy and recency effects on clicking behavior. *Journal of Computer-Mediated Communication*, *11*(2), 522–535.
- Salant, Y., & Rubinstein, A. (2008). (a, f): choice with frames. *The Review of Economic Studies*, *75*(4), 1287–1296.
- Van Erkel, P. F., & Thijssen, P. (2016). The first one wins: Distilling the primacy effect. *Electoral Studies*, *44*, 245–254.
- Wiswall, M., & Zafar, B. (2018). Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics*, *133*(1), 457–507.

Appendices

A Job Advertisement

We are recruiting top talent for multiple outside business to business sales positions in and around (City). The positions involve:

- interfacing with new and existing customers
- developing leads
- cold calling prospects
- working on quotations
- assisting with order fulfillment
- entertaining clients
- partnering with vendor representatives

Starting package includes salary commensurate with experience and comprehensive benefits, with the potential to move to an incentive-based system with unlimited six-figure earning potential. Positions require a minimum of a bachelor's degree, with some sales experience preferred.

Press 'Apply' to express your interest in receiving company information and detailed job descriptions. Thank you for your interest in these competitive positions!

B Treatment E-mail

Dear «First Name» «Last Name»,

Thank you for your interest in hearing more about our open positions!

We are currently hiring new outside business to business sales representatives in «City» for «Firm Name». «Firm Name» is a wholesale distributor of disposable goods including janitorial and office products with a strong regional presence.³⁹ The position we are hiring for involves interfacing with new and existing customers, developing leads, cold calling prospects, working on quotations, assisting with order fulfillment, entertaining clients, and partnering with vendor representatives.

«Firm Name» has a solid infrastructure in both their «City & Corporate Offices» to support your sales efforts, so that maximum time can be devoted to interfacing with customers. Most meetings take place in-person, but both hours and place of work are flexible. Compensation for the position consists of a competitive base salary plus expenses, potentially moving to an incentive-based system with unlimited six-figure earning potential. The comprehensive benefits package includes:

- Health insurance
- Generously funded Health Savings Account
- 401(k) plan with matching contributions offered
- Employer paid life insurance provided
- Dental, vision, long term disability and AFLAC plans offered
- Employee Assistance Program offered and paid 100%
- Paid Time Off

Minimum Requirements:

- 4 year degree from an accredited university or college.
- Evidence of achievement and progression of results in school, career, or outside interests.
- Basic understanding of business transactions- bids and proposals, requisitions, purchase orders, logistics, payments, etc.
- Firm understanding of gross profit margins.

³⁹The order of the product sectors here and later in the treatment e-mail was randomized.

- Excellent written and verbal communication skills- able to recognize and recall nuance from conversations as well as present written data/information.
- Able to prospect, persuade, and close business accounts- ideal candidate will be able to get commitments from potential customers even when the initial answer is no.
- Excellent time management and prioritization.
- Previous experience in outside B2B sales is preferred.

Janitorial Products

You should be comfortable working with janitorial product buyers such as facility managers, operations, and maintenance workers. The most successful team members assess problems and offer solutions that improve the productivity and cost efficiency of customers' facilities, while being profitable to the business.

To apply for this position, please send your resume and cover letter (optional) to the hiring manager «**Manager Name & E-mail**», then complete the application at the link below.

<http://www.«Recruiting Firm Name».com/job-application-«Firm Abbreviation»>

Office Products

You should be comfortable working with office product buyers such as office managers, receptionists, and administrative workers. The most successful team members assess problems and offer solutions that improve the productivity and cost efficiency of customers' offices, while being profitable to the business.

To apply for this position, please send your resume and cover letter (optional) to the hiring manager «**Manager Name & E-mail**», then complete the application at the link below.

<http://www.«Recruiting Firm Name».com/job-application-«Firm Abbreviation»>

Please **submit your application to only one manager**. Application materials sent to more than one manager will not be processed.⁴⁰

We look forward to reviewing your application. Thank you again for your interest!

«Recruiting Firm Name»

⁴⁰In the first wave, these two sentences were instead “To avoid duplication, **please submit your application to only one manager**.” . The change as made in the second round to reduce administrative burden from contacting workers who submitted application materials to more than one manager and increase worker compliance with protocol.

C Recruitment E-mail

Dear (First Name) (Last Name),

Vanderbilt University is conducting a paid economics study, in which you will be asked to make educated guesses using your career experience in sales. You are guaranteed a \$5 payment just for completing the study, but you can earn \$15 more based on how accurate your guesses are. In addition, you will have the opportunity to earn tickets into a drawing for a \$500 Amazon gift card!

The survey takes just 10-15 minutes to complete, including listening to instructions about making your educated guesses. Your participation is crucial to obtaining a representative sample, as only a limited number of individuals have been invited.

You will be paid via PayPal within 24 hours and can deposit the money directly into your bank account or spend it online. You do not need a Paypal account to receive payment. Thank you in advance for assisting Vanderbilt University in this important research!

Note that your participation in the survey is completely voluntary and your responses are kept anonymous. Your responses to the survey are used for research purposes only- there will be no commercial use of the survey and your data will never be shared. This study has been approved by the Vanderbilt University Institutional Review Board. If you have any questions, contact Kristine Koutout at kristine.f.koutout@vanderbilt.edu.

To participate, click on your personal link below. Please be sure to take the survey on a large screen like a laptop or tablet rather than a phone. If you need to leave the survey at any time, you can return to your place using the link below.

(Link)

Thank you again for your participation! If you want to opt out of receiving reminders about this survey, please reply to this e-mail.


Kristine Koutout
Vanderbilt University


D Experiment Screenshots

Thank you for participating!

Please click on the audio file to listen to the instructions. You can follow along in the text below.

Press Play.



Note that the "Next" button on pages with instructions (like this one) will only become available after the audio file has played through. If you prefer to read the instructions, please press the  button to mute the audio while it plays or silence your device.

Please also note that the link you used to reach this survey is unique to you. If you need to leave the survey, you can always return to exactly where you left off by clicking on the link again.

The "Progress Bar" at the bottom of each page tracks how far along you are in the survey. Once you reach the end, you will learn how much money you earned total. You are guaranteed to earn at least \$5.00. See **Payment** information and **Instructions** below for more details on how your total earnings will be determined.

Anonymity

Your responses on this survey are completely anonymous and will be used for research purposes only. Information like your name and e-mail address is not stored with your responses to the survey.

This survey has been reviewed by the Institutional Review Board (IRB) at Vanderbilt University to ensure your privacy is protected. You may direct any questions or concerns about this survey to the Vanderbilt IRB at (866) 224-8273. For more information, click [here](#).

Payment

For simply completing this 10 to 15 minute survey, you will be paid \$5.00. You may earn an additional \$15.00 prize based on your educated guesses in the next section of this survey. Your payment will be sent via PayPal to the e-mail address of your choice within 24 hours.

In addition, you will have the opportunity to be entered into a drawing for a \$500.00 Amazon gift card. The drawing for the \$500.00 Amazon gift card will occur at the conclusion of this online survey. You will receive more information about winning the \$500.00 Amazon gift card after you make your guesses.

Instructions


In this survey, you will be asked to make five educated guesses about how people performed on an assessment of sales ability. One guess out of those five will be randomly selected by a computer at the end of the survey to determine your total payment.

Your selected guess will be compared to the *actual performance* of randomly selected people who took the assessment of sales ability as a part of their job application for a sales position. The closer your educated guess is compared to the actual outcome of those randomly selected people, the more likely you are to earn the \$15.00 prize. A lottery system determines how likely you are to earn the \$15.00 prize based on the accuracy of your guess. (If you are interested, the lottery system is carefully designed so that it is mathematically optimal to submit your best guess about the median outcome.) **You are most likely to earn the \$15.00 prize if you submit your true best guess.**

To demonstrate how the lottery system works, I will go through an example with you.

Next

Progress Bar



Example


Press Play.



In this example, you are asked to make an educated guess about which city is closer to St. Louis, Missouri: Washington, D.C. or Phoenix, Arizona. You will not be paid for this example; it is only to ensure that you understand how to make your guess.

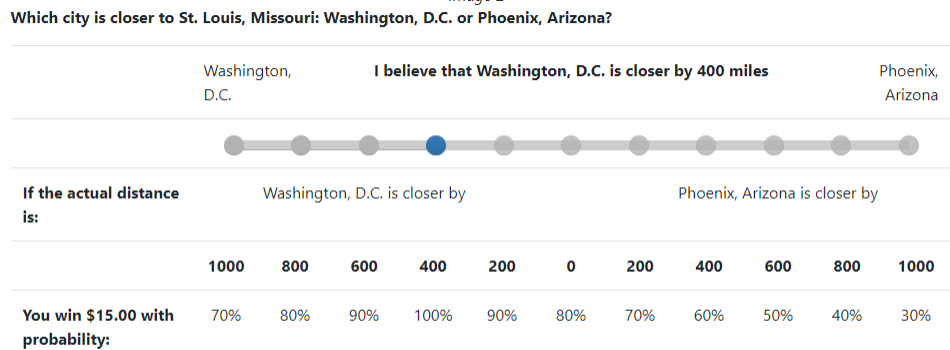
You will see a screen like the following:

Image 1



Suppose you believe Washington, D.C. is 400 miles closer to St. Louis, Missouri than Phoenix, Arizona. You would move the slider in the section that says "Washington, D.C." until it says "400." Note that the slider now says "I believe that Washington, D.C. is closer by 400 miles."


Image 2



The chart below the slider shows your probability of winning the \$15.00 prize based on what the actual distance is and your guess. For example, if your guess is accurate and Washington, D.C. is 400 miles closer than Phoenix, Arizona, you win the \$15.00 prize for sure (100%).

Image 3


Which city is closer to St. Louis, Missouri: Washington, D.C. or Phoenix, Arizona?



If Washington, D.C. is actually 200 miles closer, you have a 90% chance of winning the \$15.00 prize.

Image 4


Which city is closer to St. Louis, Missouri: Washington, D.C. or Phoenix, Arizona?



On the other hand, if Phoenix, Arizona is actually closer than Washington, D.C. by 200 miles, your chance of winning the \$15.00 prize falls to 70%.

Image 5


Which city is closer to St. Louis, Missouri: Washington, D.C. or Phoenix, Arizona?



As you move the slider, the chart will update to show the probabilities of winning the \$15.00 prize at each possible value of the actual distance. So, if you decided Washington, D.C. was actually 800 miles closer than Phoenix, Arizona, the chart would change when you moved the slider.


Image 6

Which city is closer to St. Louis, Missouri: Washington, D.C. or Phoenix, Arizona?



You will now have an opportunity to test the slider and make your guess. Remember, this example is just for practice and you will not be paid for the results.


Next



Example Task

In this example, you are asked to make an educated guess about which city is closer to St. Louis, Missouri. You will not be paid for this example; it is only to ensure that you understand how to make your guess.

Which city is closer to St. Louis, Missouri: Washington, D.C. or Phoenix, Arizona?




Next

Progress Bar

Example Results

The results show you the actual outcome and your probability of winning the \$15.00 prize based on your guess.

Washington, D.C. is actually 600 miles closer to St. Louis, Missouri than Phoenix, Arizona is. You would have won \$15.00 with 70% probability.



Now, press the random number generator (RNG). The RNG selects a random number between 1 and 100. If that number is equal to or lower than your probability of winning, you would earn the \$15.00 prize. In other words, if you have an 80% chance of winning, then you would win the \$15.00 prize if the RNG selects a number that is 80 or lower.

RNG 93

I'm sorry, you would not have won the \$15.00 prize.

Next

Progress Bar

Task Instructions

Press Play.



You will now make five educated guesses that determine whether you earn the \$15.00 prize. Remember that each guess is related to the performance of individuals on an assessment of sales ability. This sales assessment is used by employers to evaluate job candidates for positions in outside sales.


The assessment evaluates job candidates in three broad categories: "will to sell" (e.g. desire and commitment to success), "sales DNA" (e.g. controls emotions and handles rejection), and "selling competencies" (e.g. hunting and closing sales).

Individual performance on the assessment is measured as a percentile in the total population of job candidates who have taken this assessment (more than 2 million people). So, an individual who scored 90% on the sales assessment performed better than 90% of the people who completed this assessment. An individual who scored 25% performed better than 25% of people who completed the assessment.

Consider your choices carefully. Remember that one of the guesses you make will be randomly selected by a computer to determine your payment. Each guess is equally likely to be selected but you will not know which guess is selected for payment until the end of the survey. It is in your best interest to treat each guess as if it is the one that determines your payment.

Next

Progress Bar



Task 1

A computer at the end of this survey will randomly select a man and a woman who completed the sales ability assessment when applying for a sales position. **Who do you believe did better on the sales ability assessment, the randomly selected man or the randomly selected woman, and by how much?**

	Woman	I believe that they did the same																		Man	
If the actual outcome is:	The woman did better by									The man did better by											
	50+	45	40	35	30	25	20	15	10	5	0	5	10	15	20	25	30	35	40	45	50+
You win \$15.00 with probability:	50%	55%	60%	65%	70%	75%	80%	85%	90%	95%	100%	95%	90%	85%	80%	75%	70%	65%	60%	55%	50%
Next																					
	Progress Bar																				

Additional Task Instructions


Press Play.



In an earlier version of this survey, managers responsible for hiring and supervising outside sales representatives answered the same question you just did. They were given the same instructions you received at the beginning of this survey, including the information about the anonymity of their responses, and asked to make their true best guess. These managers were also paid based on how close their educated guess was to the actual outcome when a computer selected a random man and a random woman. **You must now make educated guesses about what those managers chose as their guesses.** You will make your last four guesses, each corresponding to a different type of manager. Consider your choices carefully. Again, any one of your guesses could be randomly selected to determine your payment at the end of this survey and each is equally likely.


Next

Progress Bar



Task 2, Female Manager in a Traditionally Female Product Sector

A computer at the end of this survey will randomly select a female manager responsible for hiring and managing sales representatives in a traditionally female product sector. **What did she choose when asked "Who do you believe did better on the sales ability assessment, the randomly selected man or the randomly selected woman, and by how much?"**




Next




Task 3, Male Manager in a Traditionally Female Product Sector

A computer at the end of this survey will randomly select a male manager responsible for hiring and managing sales representatives in a traditionally female product sector. **What did he choose when asked "Who do you believe did better on the sales ability assessment, the randomly selected man or the randomly selected woman, and by how much?"**




Next




Task 4, Female Manager in a Traditionally Male Product Sector

A computer at the end of this survey will randomly select a female manager responsible for hiring and managing sales representatives in a traditionally male product sector. **What did she choose when asked "Who do you believe did better on the sales ability assessment, the randomly selected man or the randomly selected woman, and by how much?"**




Next




Task 5, Male Manager in a Traditionally Male Product Sector

A computer at the end of this survey will randomly select a male manager responsible for hiring and managing sales representatives in a traditionally male product sector. **What did he choose when asked "Who do you believe did better on the sales ability assessment, the randomly selected man or the randomly selected woman, and by how much?"**




Next



Task 6

Recall the last educated guess you made and answer the questions below.

In the last question, did we ask you about a male or female manager?


 

In the last question, did we ask you about a traditionally male or traditionally female product sector?

Next

Progress Bar



Task 7

Press Play.

▶ 0:00 / 0:42 ———— 🔊 ⋮

You will now have the opportunity to earn tickets that will be entered into a drawing for a \$500.00 Amazon gift card. At the conclusion of this study, one ticket will be drawn to determine the winner, who will receive the gift card via e-mail. **The more tickets you earn, the greater chance you have of winning the \$500.00 Amazon gift card.**

In an earlier version of this survey, managers responsible for hiring and supervising outside sales representatives were asked the following question:

"In general, do you think customers in your product sector prefer working with a male sales representative, a female sales representative, or have no preferences either way?"

You must now guess what each type of manager chose. For each correct guess you make below, you earn 5 tickets for the drawing to win the \$500.00 Amazon gift card.

Consider a **male manager in a traditionally male sector**. What do you believe he chose when asked:

"In general, do you think customers in your product sector prefer working with a male sales representative, a female sales representative, or have no preferences either way?"

 ▾

Consider a **female manager in a traditionally male sector**. What do you believe she chose when asked:

"In general, do you think customers in your product sector prefer working with a male sales representative, a female sales representative, or have no preferences either way?"

 ▾

Consider a **male manager in a traditionally female sector**. What do you believe he chose when asked:

"In general, do you think customers in your product sector prefer working with a male sales representative, a female sales representative, or have no preferences either way?"

 ▾


Consider a **female manager in a traditionally female sector**. What do you believe she chose when asked:

"In general, do you think customers in your product sector prefer working with a male sales representative, a female sales representative, or have no preferences either way?"

 ▾

Next

Progress Bar



Task 8

Press Play.



You will now have another opportunity to earn tickets that will be entered into the drawing for a \$500.00 Amazon gift card.

You must guess whether the product sectors listed below are traditionally male or traditionally female. If your guess matches the most commonly guessed answer at the end of this study, you earn 5 tickets into the drawing for the \$500.00 Amazon gift card. Each guess is worth 5 tickets, so you can earn up to a total of 45 tickets into the drawing.

For example, if most people who take this survey guess that "pharmaceuticals" is a traditionally male product sector, then you earn 5 tickets if you also guessed that the product sector is traditionally male.

Banking Services:

Beauty Products:

Financial Services:

Janitorial Products:

Manufacturing Machinery:

Office Products:

Packaging Products:

Pharmaceuticals:

Travel Services:


Next

Progress Bar



Task 9

Press Play.



You will now have **one more opportunity to win cash**, paid via PayPal within 24 hours.


Choose one of the seven gambles listed below. Each gamble has two possible outcomes: Low and High. The two outcomes are equally likely. At the end of this survey, a random number generator (RNG) will select a number between 1 and 100. If the random number is 50 or lower, you win the "Low" dollar amount. If the number is 51 or higher, you win the "High" dollar amount.

For example, suppose you choose gamble 4. Then, if the random number generated is 50 or lower, you win \$0.45. If the random number generated is 51 or higher, you win \$2.15.

Gamble	Low	High
<input type="radio"/> 1	\$1.00	\$1.00
<input type="radio"/> 2	\$0.75	\$1.45
<input type="radio"/> 3	\$0.60	\$1.80
<input type="radio"/> 4	\$0.45	\$2.15
<input type="radio"/> 5	\$0.30	\$2.50
<input type="radio"/> 6	\$0.15	\$2.65
<input type="radio"/> 7	\$0	\$2.80

Next

Progress Bar



Task 10

In general, do you prefer to work with mostly men, mostly women, an equal number of each, or have no preferences either way?

In general, do you prefer a male manager, a female manager, or have no preferences either way?

In general, do you prefer to work in a traditionally male product sector, a traditionally female product sector, or have no preferences either way?

We welcome any comments you have about your preferences. Remember that all responses in this survey are anonymous.

Next

Progress Bar



Task 11

Consider a hypothetical situation in which you and 9 other sales representatives are hired at the same time to promote a new product line. At the end of one year, how do you think you will rank in total sales relative to the other sales representatives hired at the same time as you?

Rank 1 for highest through 10 for lowest.

Hypothetical rank:

1 2 3 4 5 6 7 8 9 10

Next

Progress Bar



Survey Questions

How accurate do you think the educated guesses you made in this survey are?

How often do you take risks in general?

How often do you take risks in your career?

Are you currently employed full-time, employed part-time, or unemployed?

If currently employed, are you paid by salary, hourly, or commission?

What was your total earned income last year?

Are you actively seeking employment or looking to change jobs?

If you are not currently employed, answer the questions based on your last job.

Is your direct supervisor male, female, or other?

Do you work with mostly women, mostly men, or an equal number of each?

What is your age?

What is your gender?

What is your race/ethnicity?


What is your marital status?

How many children do you have?

How many children under 5 do you have?

See Results


Progress Bar



Results

Your payment will be based on the following choice:

You guessed that they did the same and the actual outcome was the man did better by 15% , so you have a 85% chance of winning \$15.00.



Now, press the random number generator (RNG). The RNG selects a random number between 1 and 100. If that number is equal to or lower than your probability of winning, you would earn the additional \$15.00. In other words, you have a 85% chance of winning, so you win the \$15.00 if the RNG selects a number that is 85 or lower.

RNG 89

Task 9

You chose gamble 5 from the table below. Now, press the random number generator (RNG). The RNG selects a random number between 1 and 100. If that number is equal to or lower than 50, you win \$0.30. If the RNG selects a number greater than 50, you win \$2.50.

Gamble	Low	High
1	\$1.00	\$1.00
2	\$0.75	\$1.45
3	\$0.60	\$1.80
4	\$0.45	\$2.15
5	\$0.30	\$2.50
6	\$0.15	\$2.65
7	\$0	\$2.80

RNG 5

Based on the gamble you chose, **you earn \$0.30.**

Results

Task 7

You said you thought a male manager in a traditionally male sector chose "Male" when asked:

"In general, do you think customers in your product sector prefer working with a male sales representative, a female sales representative, or have no preferences either way?"

In fact, the manager chose "Male", so based on your answer **you earn 5 tickets** for the \$500.00 Amazon gift card drawing.

You said you thought a female manager in a traditionally male sector chose "Male" when asked:

"In general, do you think customers in your product sector prefer working with a male sales representative, a female sales representative, or have no preferences either way?"

In fact, the manager chose "No Preference", so based on your answer **you earn 0 tickets** for the \$500.00 Amazon gift card drawing.

You said you thought a male manager in a traditionally female sector chose "Female" when asked:

"In general, do you think customers in your product sector prefer working with a male sales representative, a female sales representative, or have no preferences either way?"

In fact, the manager chose "Male", so based on your answer **you earn 0 tickets** for the \$500.00 Amazon gift card drawing.

You said you thought a female manager in a traditionally female sector chose "Male" when asked:

"In general, do you think customers in your product sector prefer working with a male sales representative, a female sales representative, or have no preferences either way?"

In fact, the manager chose "No Preference", so based on your answer **you earn 0 tickets** for the \$500.00 Amazon gift card drawing.

You have earned a total of 5 tickets for the \$500.00 Amazon gift card drawing, in addition to those you may earn based on your answers about whether certain product sectors are traditionally male or female. If your ticket is drawn at the conclusion of this study, you will be contacted via e-mail to receive the gift card.

Thank you for completing the survey! You have earned a total of \$5.30. **Please fill out the below receipt to receive payment.** This information is collected separately to ensure that your responses are anonymous. Your payment will be sent within 24 hours of receiving this receipt.

Receipt

E Appendix E: Theoretical Framework

I outline a simple theoretical framework to show how workers' second-order beliefs could affect the job application decision. The model incorporates design elements from the experiment to better represent the context of the worker's decision. Workers may be male or female, but the model is symmetric by gender, so I analyze the case of one gender.

I study the behavior of risk-neutral workers in a market in which there are four jobs. The jobs are characterized by manager type, as listed in Table 9. Manager type is determined by the manager's own gender, indexed by $g \in \{m, f\}$, and the product sector, indexed by $s \in \{j, o\}$ where j and o correspond to the male (janitorial) and female (office) product sectors from the field experiment respectively.

Table 9: Manager Types

Type	Manager Gender	Product Sector
$\{m, j\}$	Male	Janitorial (M)
$\{f, j\}$	Female	Janitorial (M)
$\{m, o\}$	Male	Office (F)
$\{f, o\}$	Female	Office (F)

There are two time periods $t \in \{1, 2\}$. In the first time period, there are two job advertisement slots $a \in \{1, 2\}$. For every worker $i \in I$, Nature randomly draws one manager type from the office sector $\{g, o\} = \{\{m, o\}, \{f, o\}\}$ and one manager type from the janitorial sector $\{g, j\} = \{\{m, j\}, \{f, j\}\}$. Nature then randomly assigns each manager type drawn to a job slot. In total, there are eight possible combinations of draws in time period 1 listed in Table 10. In the second time period, Nature takes another draw of one of the four manager types in Table 1. The manager type randomly drawn in time period t advertised in slot a (if in time period 1) is $\{g, s\}_i^{t,a}$. There are an equal proportion of each manager type in the population.

Table 10: Eight Possible Job Advertisement Draws in Time Period 1

Draw	First Job Ad	Second Job Ad	Draw	First Job Ad	Second Job Ad
	Janitorial Sector (M)	Office Sector (F)		Office Sector (M)	Janitorial Sector (F)
	Manager Gender	Manager Gender		Manager Gender	Manager Gender
1	Male $\{m, j\}$	Male $\{m, o\}$	5	Male $\{m, j\}$	Male $\{m, o\}$
2	Female $\{f, j\}$	Female $\{m, o\}$	6	Female $\{f, j\}$	Female $\{m, o\}$
3	Male $\{m, j\}$	Female $\{f, o\}$	7	Male $\{m, j\}$	Female $\{f, o\}$
4	Female $\{f, j\}$	Male $\{f, o\}$	8	Female $\{f, j\}$	Male $\{f, o\}$

In time period 1, workers face the manager type in job slot 1 $\{g, s\}_i^{1,1}$. They may choose to either 1) apply to the first manager type $y(\{g, s\}_i^{1,1}) = 1$, 2) pay a cost c

to view the manager type in slot 2 $\{g, s\}_i^{1,2}$, or 3) proceed to time period 2 without applying. If the workers chooses to pay c to view the manager type in slot 2, they may then choose to either 1) apply to the first manager type $y(\{g, s\}_i^{1,1}) = 1$, 2) apply to the second manager type $y(\{g, s\}_i^{1,2}) = 1$, or 3) proceed to time period 2 without applying.

In time period 2, workers who did not apply in time period 1 face the manager type $\{g, s\}_i^2$. They may choose to either 1) apply to the manager $y(\{g, s\}_i^2) = 1$, or 2) accept their reservation wage r .

This set-up mimics the experiment design, in which treatment is the assignment of managers to sectors and the order in which those sectors appear in the job advertisement. The first period corresponds to the period in which workers participate in the experiment. The cost of viewing the job advertised second is due to workers' limited attention.

Workers have second-order beliefs $b_i(\{g, s\}^{t,a}) \in \mathbb{R}$, which are randomly drawn from normally distributed manager type-specific distributions $b_i(\{g, s\}^{t,a}) \sim N(\mu_{gs}, \sigma_{gs}^2)$. The random variable $b_i(\{g, s\}^{t,a})$ is defined to be a function of the worker's subjective distribution of managers' $j \in J$ first-order beliefs b_j about the gender-specific distribution of worker productivity that serves as a sufficient statistic for expected probability of being hired for worker i .⁴¹ Specifically, $b_i(\{g, s\}^{t,a}) = f[\mathbb{E}_i^{gs}(b_j)]$, where $f(\cdot)$ is a one-to-one mapping with $f(\cdot) > 0$ and the expected value is taken over worker i 's manager type-specific subjective distribution of b_j .

The worker's utility function is defined to be a function of second-order beliefs $U(b_i(\{g, s\}^{t,a}))$, where $U'(\cdot) > 0$. So, worker i 's utility from applying for a job with manager type $\{g, s\}_i^{t,a}$ is completely determined by the worker's expectation about the probability of being hired based on their second-order beliefs about that manager type $b_i(\{g, s\}^{t,a})$.

The worker's application decision can be determined through backwards induction. To simplify the expressions of the workers' decision rules, I use $b^{t,a}$ to represent $U(b_i(\{g, s\}^{t,a}))$. In time period 2, worker i applies for a job with manager type $\{g, s\}_i^2$ when the utility from applying is greater than the utility from taking the reservation wage $b^2 > U(r)$.

In time period 1, workers who chose to pay c must choose whether to apply for a job with one of the manager types drawn in period 1, $\{g, s\}_i^{1,1}$ or $\{g, s\}_i^{1,2}$, or to proceed to time period 2, where utility is discounted by a factor of δ . Worker i chooses to apply when

$$\max(b^{1,1}, b^{1,2}) > \delta \max(\mathbb{E}(b_i^2), U(r)) \quad (4)$$

⁴¹I abstract away from worker gender since the cases are symmetric. The distribution of worker productivity is the only point that differs for the genders. Women workers consider managers' first-order beliefs about the distribution of women workers, while men workers consider managers' first-order beliefs about the distribution of men workers.

Lastly, I consider workers at the beginning of time period 1. The worker has three choices. One, the worker can apply to the first manager type $y\left(\{g, s\}_i^{1,1}\right) = 1$ and receive utility $b^{1,1}$. Two, the worker can proceed directly to time period 2 and receive, in expectation, $\delta \max(\mathbb{E}(b_i^2), U(r))$. Three, the worker can pay the cost c of observing the second manager type. The expected utility from paying the cost c to observe the second manager type is

$$\frac{1}{2}b^{1,2}(m)\mathbb{1}[b^{1,2}(m) > b^{1,1}] + \frac{1}{2}b^{1,2}(f)\mathbb{1}[b^{1,2}(f) > b^{1,1}] + b^{1,2}[1 - Pr(b^{1,2} > b^{1,1})] - c \quad (5)$$

There are three cases. One, the expected utility from either of the two manager types in job slot 2 is higher than the utility from the known manager type in slot 1 $b^{1,2}(m) > b^{1,1}$ and $b^{1,2}(f) > b^{1,1}$. In this case, the worker would choose the manager type in slot 2 over the manager type in slot 1 regardless of the gender of the manager in slot 2; therefore, $1 - Pr(b^{1,2} > b^{1,1})$ is zero and the worker's decision is based on whether the difference between the expected utility of the manager type in slot 2 and the known utility of the manager type in slot 1 is greater than the cost of observing the second manager type $\frac{1}{2}b^{2,2}(m) + \frac{1}{2}b^{2,2}(f) - b^{1,2} > c$.

In the second case, utility from one manager type in job slot 2 is greater than the utility from the manager type in job slot 1 $b^{1,2}(g) > b^{1,1}$, but the utility from the other possible manager type in slot 2 is lower $b^{1,2}(g') < b^{1,1}$.⁴² Then, there is an equal probability that the worker chooses the manager type in slot 1 and the manager type in slot 2, based on their draw of the manager type in slot 2. For the worker to pay the cost c to observe that manager type, half the difference between the utility from the manager type in slot 2 that the worker would choose if drawn and the known manager type in slot 1 must be greater than cost of observing the manager type in slot 2. The half captures the probability that the worker draws this more favorable manager type, since the worker would pay the cost of observing the manager type in slot 2, but still choose the manager type in slot 1, if the worker drew the less favorable manager type.

Lastly, it could be the case that $b^{1,2}(m) < b^{1,1}$ and $b^{1,2}(f) < b^{1,1}$. Then, the worker would never pay the cost of observing the manager type in slot 2 based on that draw of manager type in slot 1. Note that this last case must necessarily be true for at least one of the four manager types that the worker can draw since there must exist a maximum, but there is not necessarily the converse case. Even though there must be a minimum, it is not necessarily true that there is a sufficiently large difference between the minimum and the expected utility from the manager types in the opposite sector to induce the worker to pay the cost to observe the manager type in slot 2.

⁴²Where g is one gender and g' the other.

In summary, the worker applies in period 1 when either $b^{1,1}$ or $\max(b^{1,1}, b^{1,2})$, conditional on paying the cost c to observe $b^{1,2}$, are larger than the expected utility from taking another draw of manager types in the next time period $\delta \max(\mathbb{E}(b_i^2), U(r))$. There are, therefore, three functions of beliefs that I include in the empirical specification: $b^{1,1}$, $\max(b^{1,1}, b^{1,2})$, and $\delta \max(\mathbb{E}(b_i^2), U(r))$. To make the last function tractable, I summarize $\delta \max(\mathbb{E}(b_i^2), U(r))$ with the average of all beliefs except the manager in job slot 1 in time period 1, since a worker who proceeds to time period 2 must be searching for a manager type not observed in time period 1. I do not similarly exclude the manager type in job slot 2 in the average since the worker decides endogenously whether to observe that manager type.

F Regression Tables

Table 11: Main Specification, Men

	(1)	(2)	(3)	(4)
Belief about First Manager	0.002** (0.001)	0.002** (0.001)	0.002** (0.001)	0.003** (0.001)
T1: Male Managers	0.142*** (0.038)	0.175*** (0.052)	0.181*** (0.055)	0.205 (0.176)
T2: Female Managers	0.235*** (0.050)	0.235*** (0.050)	0.250*** (0.061)	0.263 (0.183)
T3: Sector-Matching Managers	0.162*** (0.037)	0.184*** (0.048)	0.196*** (0.052)	0.200 (0.178)
T4: Non-Sector-Matching Managers	0.197*** (0.041)	0.211*** (0.042)	0.223*** (0.049)	0.251 (0.175)
Most Favorable Belief of Facing Managers		-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)
Average Belief about Non-First Manager		0.003 (0.002)	0.003 (0.002)	0.003 (0.002)
Prefer Gender of First Manager			-0.010 (0.052)	0.003 (0.051)
Prefer Opposite-Gender of First Manager			-0.034 (0.053)	-0.012 (0.055)
Prefer Sector of First Manager			-0.044 (0.055)	-0.058 (0.058)
Prefer Opposite Sector of First Manager			0.031 (0.053)	0.009 (0.051)
Controls	N	N	N	Y
Observations	351	351	351	341
R-Squared	0.193	0.198	0.202	0.257

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Robust standard errors are reported in parentheses. Controls include the wave, city, and job board in which the worker was recruited, as well as whether the worker reports actively searching for a job.

Table 12: Main Specification, Women

	(1)	(2)	(3)	(4)
Belief about First Manager	-0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)
T1: Male Managers	0.055* (0.032)	0.063* (0.036)	0.079* (0.046)	0.026 (0.054)
T2: Female Managers	0.097** (0.039)	0.125*** (0.047)	0.130** (0.051)	0.075 (0.066)
T3: Sector-Matching Managers	0.050* (0.028)	0.070* (0.041)	0.082 (0.053)	0.041 (0.064)
T4: Non-Sector-Matching Managers	0.108** (0.045)	0.123** (0.051)	0.135** (0.061)	0.086 (0.073)
Most Favorable Belief of Facing Managers		-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)
Average Belief about Non-First Manager		0.002 (0.001)	0.002 (0.001)	0.003* (0.001)
Prefer Gender of First Manager			-0.084** (0.034)	-0.082** (0.036)
Prefer Opposite-Gender of First Manager			-0.010 (0.051)	-0.040 (0.051)
Prefer Sector of First Manager			0.039 (0.051)	0.054 (0.050)
Prefer Opposite Sector of First Manager			0.015 (0.050)	0.016 (0.046)
Controls	N	N	N	Y
Observations	223	223	223	220
R-Squared	0.084	0.095	0.111	0.309

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Robust standard errors are reported in parentheses. Controls include the wave, city, and job board in which the worker was recruited, as well as whether the worker reports actively searching for a job.

Table 13: Other Specification of Treatment

	(1)	(2)
Belief about First Manager	0.002* (0.001)	-0.000 (0.001)
Most Favorable Belief of Facing Managers	-0.001 (0.002)	-0.001 (0.001)
Average Belief about Non-First Manager	0.003 (0.002)	0.003* (0.002)
T1: Male Managers, Jan First	0.176 (0.185)	0.035 (0.081)
T2: Female Managers, Jan First	0.200 (0.192)	0.077 (0.082)
T3: Sector-Matching Managers, Jan First	0.207 (0.182)	-0.045 (0.085)
T4: Non-Sector-Matching Managers, Jan First	0.311* (0.184)	0.042 (0.086)
T1: Male Managers, Off First	0.178 (0.184)	-0.016 (0.058)
T2: Female Managers, Off First	0.271 (0.196)	0.054 (0.076)
T3: Sector-Matching Managers, Off First	0.135 (0.193)	0.082 (0.070)
T4: Non-Sector-Matching Managers, Off First	0.134 (0.183)	0.119 (0.097)
Prefer Gender of First Manager	0.013 (0.051)	-0.078** (0.037)
Prefer Opposite-Gender of First Manager	-0.017 (0.056)	-0.041 (0.053)
Prefer Sector of First Manager	-0.071 (0.061)	0.044 (0.050)
Prefer Opposite Sector of First Manager	0.031 (0.054)	0.017 (0.045)
Controls	Y	Y
Observations	341	220
R-Squared	0.270	0.326

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Robust standard errors are reported in parentheses. Controls include the wave, city, and job board in which the worker was recruited, as well as whether the worker reports actively searching for a job.

Table 14: Other Specifications of Alternative Beliefs, Men

	(1)	(2)	(3)	(4)	(5)
Belief about First Manager	0.003** (0.001)	0.002** (0.001)	0.002** (0.001)	0.002** (0.001)	0.002** (0.001)
First-Order Belief	-0.001 (0.001)				
Most Favorable Belief of Facing Managers	-0.001 (0.001)			-0.001 (0.001)	
Average Belief about Non-First Manager	0.003 (0.002)	0.004* (0.002)			
Most Favorable Belief of All Other Managers			0.000 (0.001)		
Belief about Second Manager		-0.002 (0.001)			-0.000 (0.001)
Most Favorable Belief of Non-Facing Managers				0.002* (0.001)	
Average Belief of Non-Facing Managers					0.003* (0.002)
T1: Male Managers	0.191 (0.175)	0.210 (0.171)	0.168 (0.173)	0.174 (0.178)	0.210 (0.171)
T2: Female Managers	0.255 (0.182)	0.247 (0.178)	0.252 (0.185)	0.218 (0.189)	0.247 (0.178)
T3: Sector-Matching Managers	0.187 (0.177)	0.196 (0.172)	0.171 (0.175)	0.165 (0.180)	0.196 (0.172)
T4: Non-Sector-Matching Managers	0.237 (0.174)	0.237 (0.171)	0.228 (0.178)	0.191 (0.181)	0.237 (0.171)
Prefer Gender of First Manager	0.004 (0.051)	0.005 (0.051)	0.006 (0.052)	0.011 (0.051)	0.005 (0.051)
Prefer Opposite-Gender of First Manager	-0.007 (0.055)	-0.008 (0.055)	-0.013 (0.055)	-0.005 (0.055)	-0.008 (0.055)
Prefer Sector of First Manager	-0.056 (0.058)	-0.064 (0.058)	-0.064 (0.056)	-0.068 (0.058)	-0.064 (0.058)
Prefer Opposite Sector of First Manager	0.011 (0.051)	0.010 (0.051)	0.008 (0.052)	0.015 (0.052)	0.010 (0.051)
Controls	Y	Y	Y	Y	Y
Observations	341	341	341	341	341
R-Squared	0.259	0.260	0.252	0.259	0.260

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Robust standard errors are reported in parentheses. Controls include the wave, city, and job board in which the worker was recruited, as well as whether the worker reports actively searching for a job. Column (1) adds first-order beliefs. Column (2) replaces the most favorable belief of the two facing managers with the belief about the second manager type. Column (3) replaces the most favorable belief of the two facing managers and the average belief of all other managers, except the first, with the most favorable belief of all managers, excepting the first. Column (4) replaces the average belief of all other managers, except the first, with the most favorable belief of the non-facing managers. Column (5) replaces the most favorable belief of the two facing managers and the average belief of all other managers, except the first, with the belief about the second manager type and the average belief of the non-facing managers.

Table 15: Other Specifications of Alternative Beliefs, Women


	(1)	(2)	(3)	(4)	(5)
Belief about First Manager	0.000 (0.001)	-0.001 (0.001)	-0.000 (0.001)	-0.000 (0.001)	-0.001 (0.001)
First-Order Belief	0.000 (0.001)				
Most Favorable Belief of Facing Managers	-0.001 (0.001)			-0.001 (0.001)	
Average Belief about Non-First Manager	0.002 (0.002)	0.004** (0.002)			
Most Favorable Belief of All Other Managers			0.002** (0.001)		
Belief about Second Manager		-0.002** (0.001)			-0.001 (0.001)
Most Favorable Belief of Non-Facing Managers				0.002** (0.001)	
Average Belief of Non-Facing Managers					0.003** (0.001)
T1: Male Managers	0.027 (0.054)	-0.036 (0.058)	-0.049 (0.052)	-0.050 (0.059)	-0.036 (0.058)
T2: Female Managers	0.075 (0.067)	0.052 (0.063)	-0.019 (0.054)	0.019 (0.060)	0.052 (0.063)
T3: Sector-Matching Managers	0.040 (0.065)	-0.018 (0.059)	-0.045 (0.051)	-0.019 (0.059)	-0.018 (0.059)
T4: Non-Sector-Matching Managers	0.086 (0.074)	0.051 (0.072)	-0.004 (0.068)	0.009 (0.073)	0.051 (0.072)
Prefer Gender of First Manager	-0.079** (0.036)	-0.083** (0.037)	-0.091** (0.038)	-0.083** (0.036)	-0.083** (0.037)
Prefer Opposite-Gender of First Manager	-0.040 (0.051)	-0.036 (0.050)	-0.035 (0.052)	-0.034 (0.050)	-0.036 (0.050)
Prefer Sector of First Manager	0.053 (0.050)	0.051 (0.051)	0.050 (0.050)	0.049 (0.051)	0.051 (0.051)
Prefer Opposite Sector of First Manager	0.018 (0.047)	0.015 (0.045)	0.011 (0.046)	0.012 (0.044)	0.015 (0.045)
Controls	Y	Y	Y	Y	Y
Observations	220	220	220	220	220
R-Squared	0.310	0.321	0.305	0.314	0.321

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Robust standard errors are reported in parentheses. Controls include the wave, city, and job board in which the worker was recruited, as well as whether the worker reports actively searching for a job. Column (1) adds first-order beliefs. Column (2) replaces the most favorable belief of the two facing managers with the belief about the second manager type. Column (3) replaces the most favorable belief of the two facing managers and the average belief of all other managers, except the first, with the most favorable belief of all managers, excepting the first. Column (4) replaces the average belief of all other managers, except the first, with the most favorable belief of the non-facing managers. Column (5) replaces the most favorable belief of the two facing managers and the average belief of all other managers, except the first, with the belief about the second manager type and the average belief of the non-facing managers.


G Probit Regressions

Figure 10: Regression Coefficient on Belief about First Manager by Gender, Main Specification



Notes: 95% confidence intervals are shown. Standard errors are robust. Coefficient is scaled to a one unit increase in the favorability of beliefs. Panel (a) reports the coefficient on the regression with men only and panel (b) reports the same for women only. In each graph, the first estimate includes only the vector of treatment variables and the belief of interest. The second estimate adds measures of alternative beliefs, the third estimate adds preferences, and the fourth estimate includes controls for the wave, city, and job board in which the worker was recruited and whether the worker reports actively searching for a job.

Figure 11: Regression Coefficient on Belief about First Manager by Gender, Other Specifications




Notes: 95% confidence intervals are shown. Standard errors are robust. Coefficient is scaled to a one unit increase in the favorability of beliefs. Panel (a) reports the coefficient on the regression with men only and panel (b) reports the same for women only. In each graph, specification (1) interacts treatment with first manager type. Specification (2) adds first-order beliefs. Specification (3) replaces the most favorable belief of the two facing managers with the belief about the second manager type. Specification (4) replaces the most favorable belief of the two facing managers and the average belief of all other managers, except the first, with the most favorable belief of all managers, excepting the first. Specification (5) replaces the average belief of all other managers, except the first, with the most favorable belief of the non-facing managers. Specification (6) replaces the most favorable belief of the two facing managers and the average belief of all other managers, except the first, with the belief about the second manager type and the average belief of the non-facing managers.


H Cardinal Belief Distributions

Figure 12: Cardinal Distributions of Second-Order Beliefs


(a) Female Manager in Female Sector




(b) Female Manager in Male Sector



(c) Male Manager in Female Sector




(d) Male Manager in Male Sector



Notes: Distributions are fitted with lines estimated using kernel densities.

Figure 13: Cardinal Distribution of First-Order Beliefs



Notes: Distribution is fitted with a line estimated using a kernel density.