The More You Know: Information Effects on Job Application Rates by Gender in a Large Field Experiment*

Laura K. Gee[1]

[1]Tufts University, Economics (laura.gee@tufts.edu)

June 25, 2015

(Keywords: field experiment, gender, labor search, big data, ambiguity aversion, herding, competition)

(JEL: C93, D01, D83, J21, J22)

*A huge thanks to Joe Adler and Jim Baer at LinkedIn. Additionally this work would not have been possible without the assistance of Mathieu Bastian, Joi Deaser, Vitaly Gordon, Abishek Gupta, Kuo-Ning Huang, Sachit Kamat, Navneet Kapur, Shreya Oswal and others at LinkedIn who shared their knowledge and expertise. Thank you to Kevin Morsony, Product Counsel at LinkedIn, for reviewing this paper. This research was approved by the Tufts IRB (1308011 and 1405012). Thanks to Darcy Covert for research assistance. Thanks to Yan Chen, Dorothea Kubler, Sonny Tambe, participants at the NBER Summer Institute, Economics Science Association Conference, Bay Area Behavioral and Experimental Economics Workshop, Behavioral & Experimental Economics of the Mid-Atlantic, Northwestern University NU-Lab seminar, George Mason seminar ICES, Bureau of Labor Statistics seminar, University of Michigan seminar, University of Massachusetts Amherst seminar and Tufts seminar for their helpful comments.

Abstract

This paper presents the results from a 2.3 million person field experiment that varies whether a job seeker is shown the number of applicants for a job posting on a large job posting website, LinkedIn. This intervention increases the likelihood a person will start/finish an application by 0.6%-1.9%, representing an economically significant potential increase of over a thousand applications per day. This increase is greater for female applicants. Firms in industries that are highly represented on this job posting website may be particularly interested in this low cost, light touch intervention that potentially increases the number of female applicants.

1 Introduction

There is a documented wage gap in the U.S. with women earning about 30% less than men (Goldin, 2014). To study this issue, previous research has focused on differences in human capital accumulation and firm side discrimination. However, a more recent stream of laboratory produced research has found that women tend to be more ambiguity/risk averse than men, and that women dislike competition more than men do (Azmat and Petrongolo, 2014; Bertrand, 2011; Croson and Gneezy, 2009; Eckel and Grossman, 2008). These behavioral differences observed in the lab present another possible explanation for differences in occupation choices and competitive performance in the real world. To understand the extent to which these laboratory results translate to real-world behavior, several large scale field studies have been conducted (Chen and Konstan, forthcoming; Samek, 2015; Flory et al., 2015). This paper fits within this field of research by examining behavioral motivations across genders in a large scale field study with real labor market implications.

In many theoretical models job seekers are generally modeled as facing decisions with both known risks, such as the likelihood of a job offer, and known utilities regarding prospective positions.[1] However, in reality, job seekers face many unknown risks or "ambiguity" about the probability of an offer or the utility of a position. If job seekers are ambiguity averse or if they use signals to update their beliefs, then these theoretical models lose some of their predictive power. This paper is intended to bridge the gap between theoretical assumptions and real-life behavior by analyzing the job search behavior of over 2.3 million job seekers viewing over 100,000 job postings on the website LinkedIn in March 2012. Specifically, I compare the behavior of job seekers based on the information they receive. This study varies whether job seekers are shown the number of people who previously began an application for a viewed posting.

Intuitively, adding extra information may change either the cost of applying or the expected benefit in terms of obtaining the position. Behavioral

[1] See Galenianos and Kircher (2009); Mortensen (1970); Das and Tsitsiklis (2010); Chade and Smith (2006); Weitzman (1979); Kohn and Shavell (1974); Telser (1973); Nachman (1972) and Stigler (1961).

economics offers several theories from which we can derive predictions about the effect of knowing the previous number of started applications on the likelihood of application. Specifically, I focus on the following three in this study: (1) ambiguity aversion, (2) competition avoidance, and (3) herding. Ambiguity aversion suggests having more information will reduce ambiguity and in turn increase application likelihood. Competition avoidance suggests job seekers may try to avoid competition when there is a high number of started applications. On the other hand, herding suggests seekers will apply to more popular postings. All three behaviors may be observed in the data, but from a policy perspective it is important to know which one has the largest average effect on job application rates, and in particular if there are heterogeneous treatment effects for men and women.

Understanding the interaction between behavioral factors and job search behavior could be used to create a welfare gain from a better functioning labor market. If ambiguity aversion dominates, then adding more information to the job posting may increase the likelihood of application, especially for women. In turn, this may enhance welfare by both increasing the thickness of the market and decreasing the gender occupation gap.[2] By contrast, if competition avoidance dominates there may be a welfare gain from decreased congestion, but also a decrease in the number of female applicants. Last, if herding behavior dominates, the resulting congestion may create a welfare loss.

The results from this experiment show no strong pattern of either competition avoidance or herding for either gender. However, interestingly, the results show that the addition of information increases the likelihood of starting or finishing an application by 0.6-1.9%, representing a potential increase of about a thousand applications per day for posting on the site. My analysis shows that this increase is largely driven by female job seekers being induced to apply. For example, showing this information results in an almost 6% increase in

[2]Theoretically, having a larger applicant pool will increase the expected value of the final match (Barron et al., 1985). Empirically, Van Ours and Ridder (1992) find that vacancies are filled more quickly when there is a larger applicant pool. Thus, increasing the number of applicants may result in welfare gains as long there is not too much congestion (Roth, 2008) and as long as it does not exacerbate differences in occupational choices across genders.

the likelihood a woman will finish an application, while the effect is not measurable for men. This finding is consistent with research that shows women are more ambiguity averse than men (Eckel and Grossman, 2008). Additionally, I find that the treatment increases the likelihood a female job seeker will apply to a job traditionally perceived as "male" by 0.7-1.7 percentage points. The findings from this study have both academic and policy applications. Specifically, the results suggest that providing more information can increase female applicants in industries like high tech and finance that have higher male participation rates. Overall, this paper finds that showing more information on job postings could mitigate the male-female occupation gap by exploiting gender differences in behavioral factors to increase the thickness of the female applicant pool.

The rest of the paper proceeds as follows. Section 2 discusses the literature in more detail. Section 3 describes the field experiment. Section 4 discusses the empirical strategy and results and Section 5 concludes with suggestions for further research.

2 Literature Review

Research has shown that one reason for the gender wage gap is that men are concentrated in higher paying occupations than women. However, it is unclear how much of this occupational segregation is driven by the supply side choices of women to seek lower-paying occupations rather than demand side discrimination. Previous studies have tended to focus on the demand side factors. For example Petit (2007) and Neumark et al. (1996) manipulate the name on a resume and find that men are more likely to be called for interviews than women.[3] Similarly, Goldin and Rouse (2000) find that blind auditions increase the likelihood that a woman is hired for a position by 50%. Finally, Bohnet et al. (Forthcoming) find that female applicants are evaluated differently than male applicants. However, other studies have not found evidence of gender bias in the hiring process. For example, Kuhn and Shen (2013) find that

[3] A notable exception is Bertrand and Mullainathan (2004).

higher skilled jobs are actually less likely to show a gender preference in their job postings. In addition, large employers of high skilled workers in the US have recently explicitly stated they would like to close the gender gap in their firms.[4] Finally, a set of studies has shown that increased gender diversity in the workforce has positive results for a firm (Weber and Zulehner, 2014, 2010; Hellerstein et al., 2002). Together these studies show that although some of the occupation gender gap may be driven by demand side discrimination, this does not seem to be the full story.

Regarding supply side factors, Fernandez and Friedrich (2011) find that female job seekers state a preference for a more "female" receptionist position versus a more "male" computer programmer position. This implies that women's underlying preferences are driving the occupation gap; thus from a policy perspective, we would need to change women's preferences to close the gap. However, in a recent field study Flory et al. (2015) find women are less likely to apply to postings that include more "male" wording, a more competitive pay structure, or greater pay uncertainty. Gaucher et al. (2011) find a similar result in a laboratory setting. In another study Samek (2015) finds a more competitive pay scheme deters both men and women from applying, but that the effect is larger for women. With the exception of Fernandez and Friedrich (2011), supply side studies find that women are deterred from applying by the specific information in the job posting or advertised pay structure. In related work on financial disclosures, two studies find that psychological factors affect take up behavior (Bertrand and Morse, 2011; Bertrand et al., 2010). These findings imply that changing how job positions are advertised could decrease the occupation gender gap.

This paper contributes to the gender gap research by studying how the information provided to applicants impacts their decision to apply. Specifically, being shown the number of previous applicants may help a job seeker weigh the costs of application against the benefits of a possible job offer. Application

[4]For example in May 2014 Google announced that only 30% of its workforce is female, and only 17% of its "tech" workforce is female. Google also acknowledged that they would like to increase diversity in their workforce. See http://www.forbes.com/sites/jaymcgregor/2014/05/29/2-of-google-employees-are-black-and-just-30-are-women/

costs can be quite high in terms of time cost.[5] The benefits of applying are related to actually obtaining a job offer. In this study, I examine the relative importance of the following information effects to gain insight into applicant behaviors: (1) ambiguity aversion, (2) competition avoidance, and (3) herding.

Laboratory studies find that both women and men are affected by all three of these behavioral factors. These studies further find that women systematically differ in the extent to which they exhibit these behaviors (see Bertrand (2011); Croson and Gneezy (2009); Eckel and Grossman (2008) for extensive literature reviews). Most of these studies find that women are more likely to choose a piece-rate versus competitive tournament style payment scheme (Dohmen and Falk, 2011; Vandegrift and Yavas, 2009; Niederle and Vesterlund, 2007; Gneezy et al., 2003). Applying this finding to the job search process, being shown a higher number of applicants on a job posting may discourage women from applying if they prefer to avoid competition. However, it is also possible that herding toward more popular jobs may offset this reduction. Experiments on herding find that people are more likely to make the same choice they observe others making (Bougheas et al., 2013; Yechiam et al., 2008; Anderson and Holt, 1997). Although these studies do no break down results by gender, they suggest in general, that herding would lead to more job seekers applying to positions which are already over-subscribed even though the overall effect is welfare dis-enhancing. Finally, it is possible that ambiguity aversion dominates in the job application process. Ambiguity refers to situations in which the distribution of the random variable is unknown, whereas in contrast, risk refers to situations for which the distribution is known.[6] The job search process contains a number of random variables that determine the likelihood of an offer, the quality of the position, etc. When job seekers receive

[5] See the online Appendix available at http://laurakgee.weebly.com/index.html, for survey results finding most people estimate the time cost of an application at over an hour. The survey includes 188 respondents and a snow ball sampling method.

[6] Note that ambiguity aversion can be modeled as a specific form of risk aversion following the work of Halevy and Feltkamp (2005) who show that behavior indicative of ambiguity aversion could also be explained by risk aversion over correlated risks. See the appendix for details. Women have been shown to be more risk averse than men in many lab experiments (Bertrand, 2011; Croson and Gneezy, 2009; Eckel and Grossman, 2008).

information regarding the number of other applicants some of this ambiguity is reduced. As a result, ambiguity averse job seekers may be more likely to apply. Laboratory experiments on gambling choices find that subjects prefer options where the distribution of risks is known over gambles where the distribution is less well known (Halevy and Feltkamp, 2005; Ellsberg, 1961). Additional studies have found that women are more ambiguity averse than men over gains in non-abstract environments (Moore and Eckel, 2003; Schubert et al., 2000); such as the job application process. In the next section I elaborate on the setting for my field experiment and the experimental procedures used to test which of these effects dominates in the job search setting.

3 Field Experiment

The field experiment took place on the professional social networking website LinkedIn in March 2012. LinkedIn was launched in 2003. By April 2015, the website had 350 million members from over 200 countries.[7] LinkedIn is well known for its professional social networking functionality. However, it also acts as a job posting website. This paper concentrates on the job posting functionality of LinkedIn.

Although the population on LinkedIn is not a representative sample of the total worldwide labor force, it is particularly well-suited for a study of gender differences in the labor force. The largest industries represented on LinkedIn are "High Tech" and "Finance".[8] Industries like this tend to have lower levels of female labor force participation. For example only 32.5% of US professionals in STEM related fields (Science Technology Engineering and Mathematics) are female.[9]

To use the job postings on LinkedIn, a member first navigates to the Jobs

[7]See https://press.linkedin.com/about-linkedin. As there are about 3.5 billion people in the worldwide labor force (https://www.cia.gov/library/publications/the-world-factbook/rankorder/2095rank.html), the LinkedIn population would represent about 10% of the total labor force.

[8]http://www.linkedinppc.com/target-by-industry-company-category/

[9]http://dpeaflcio.org/programs-publications/issue-fact-sheets/women-in-stem/

landing page (Figure 1) where she is shown some pre-selected job postings.[10] At this point the member can click on one of the postings listed, or can enter a term into the search bar, which will return results like those shown in Figure 2. After clicking on a posting, a member will see a full page description of the posting. In the field experiment, the treatment and control groups receive different descriptions, with the treatment group receiving information on the number of previous applicants for the posted position.

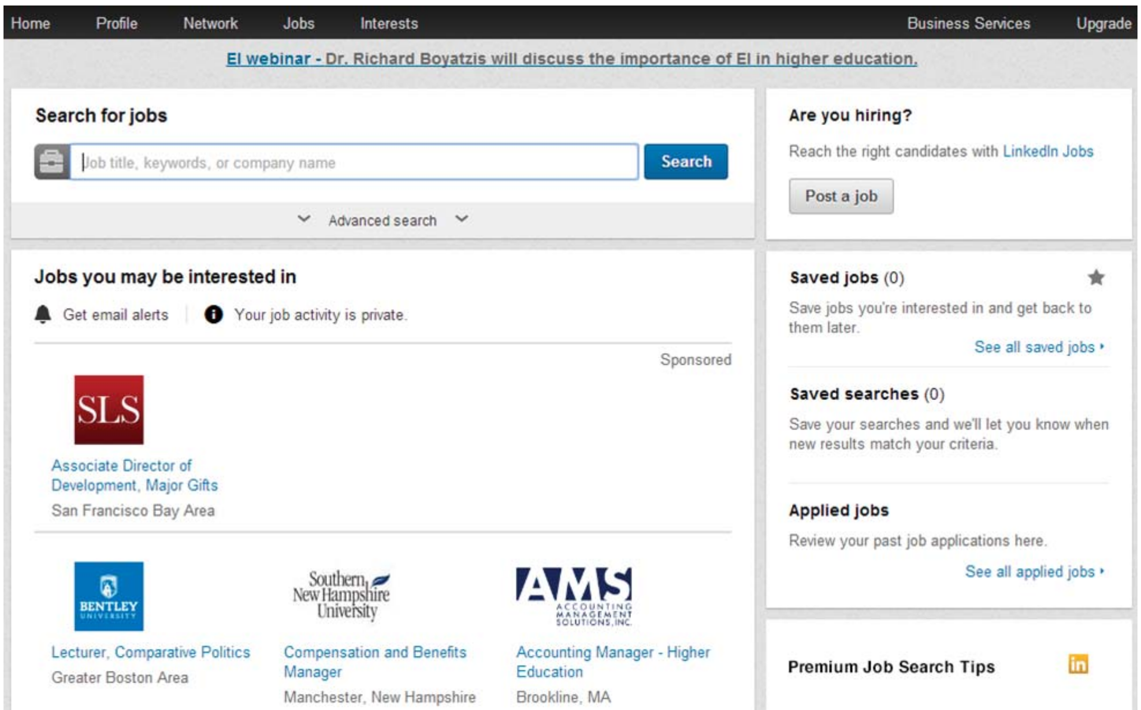

Figure 1: Jobs Landing Page

Note: This figure shows the Jobs landing page a LinkedIn user might see when she logs on to the website.

[10] Jobs are generally selected by LinkedIn based on information the member has listed on his/her profile like education, industry, and previous employment.

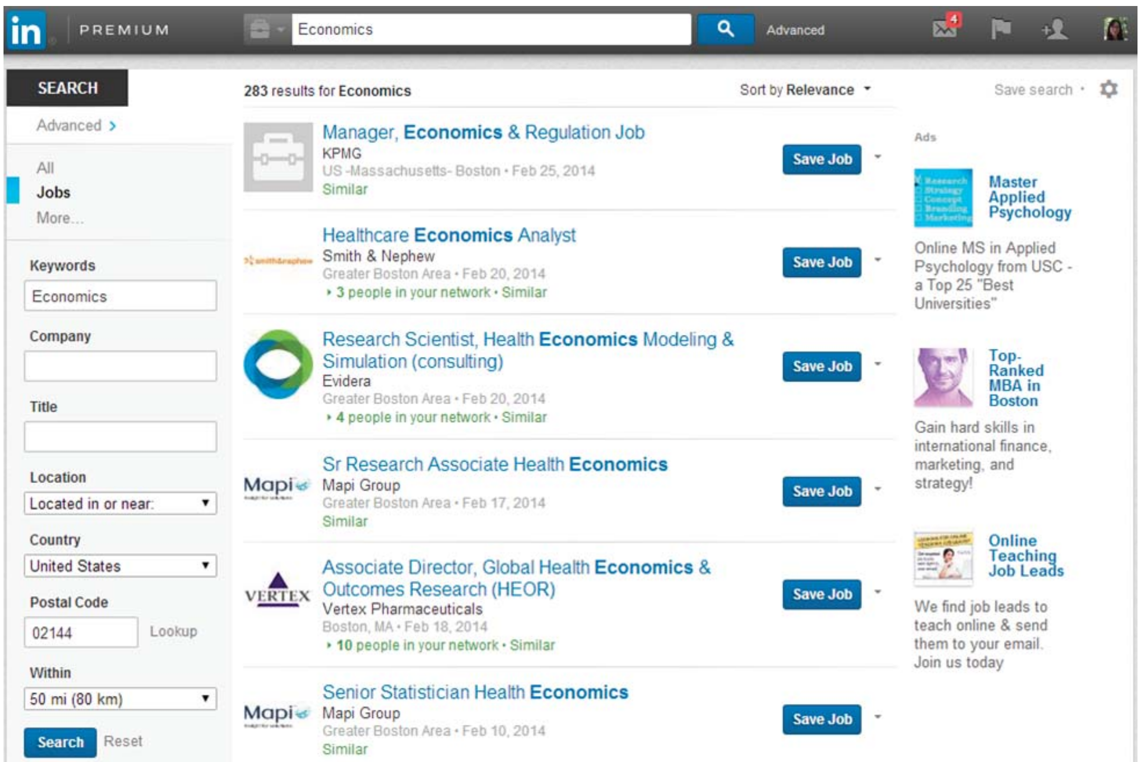

Figure 2: Job Search Landing Page

Note: This figure shows the results from a search for the term "Economics"

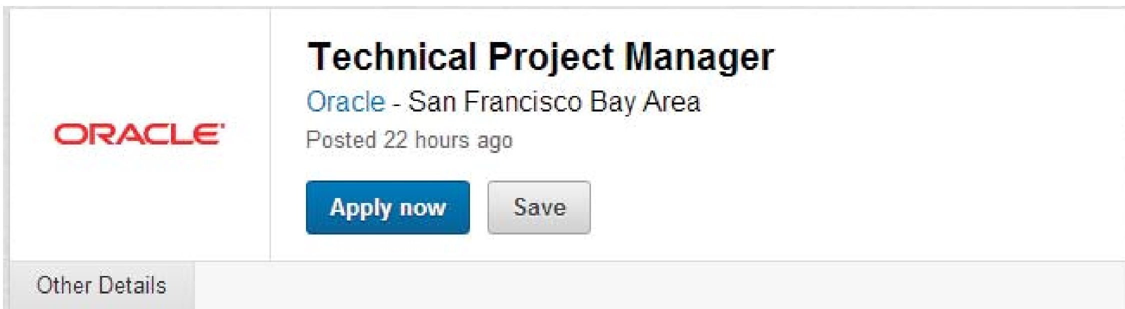

(a) Interior Job Posting

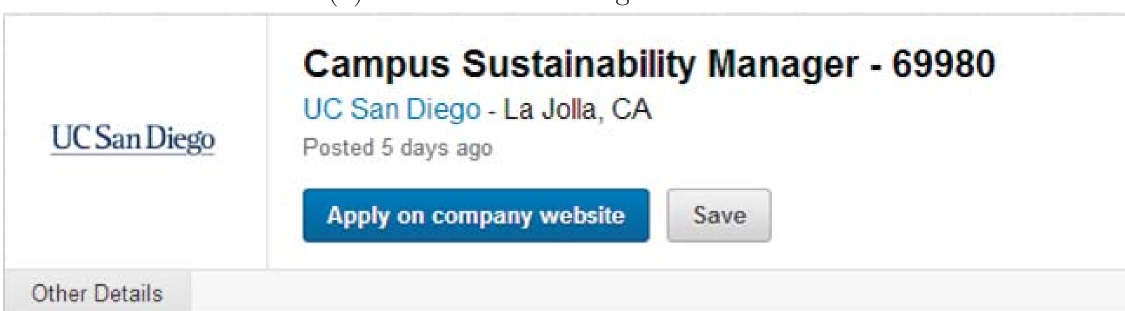

(b) Exterior Job Posting

Figure 3: Types of Job Postings on LinkedIn

Note: This figure shows an example of the two types of job postings on LinkedIn. Panel (a) shows an interior posting, which means that LinkedIn collects applications for a third party (Oracle). For these, I can observe if a person both begins and finishes an application. Panel (b) shows an exterior posting, which means that a person is directed to an external website to begin an application and thus I can only observe if someone begins the application. These screenshots were taken in February 2013, which is why they differ very slightly from the formatting seen in the example of the treatment vs. control screenshots in Figure 4.

LinkedIn provides two types of job postings (Figure 3). Interior postings are those where LinkedIn collects the finished application and forwards it to the company. With interior job postings, I can observe if a member both starts and finishes an application.[11] Exterior postings, on the other hand, link a job seeker to an external website. In this case I can observe only if a user starts an application.

The two main outcome variables in my experiment are the dummy variables "Start Application" and "Finish Application". For exterior postings, I can tell only if someone clicks on the "Apply" button. I cannot determine the time spent applying or even if the click was intentional. This limited information makes Start Application a noisy measure of interest in the position. By contrast, I can measure the outcome Finish Application for interior postings making it a more accurate measure of investment in applying for the job.

In this experiment, users were randomized into groups at the member level, so a member in the control group would see no information on any postings he visits during the 16 days of the experiment. On the other hand, a member in the treatment group looking at the same job postings would see the number of job seekers who had previously started an application as pictured in Figure 4.[12]

This design presents a unique experiment because I can observe how two people looking at the exact same posting change their behavior based on whether they know the number of other people who have already started an application. Additionally, because the information is exogenously assigned, I can rule out the possibility that those who seek out more information are already more likely to apply for a position.

The groups were determined from the set of all active LinkedIn members

[11] I have the timestamp of when a job seeker clicks "Apply" and also the timestamp for when the user submits an application. If a person submits an application within one day of viewing the posting, then I code this as a finished application. This restriction is likely to bias the number of total finished applications downward since some people may take more than a day to finish an application or may come back at a later date to finish the application. However, I have no reason to believe this bias will be different across the control and treatment groups.

[12] For an exterior job posting, the button reads "Apply on Company Website," while for an interior job posting the button simply reads "Apply Now".

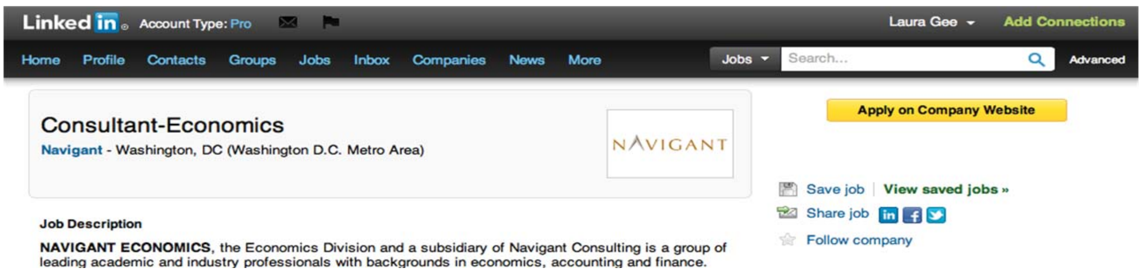

(a) Control

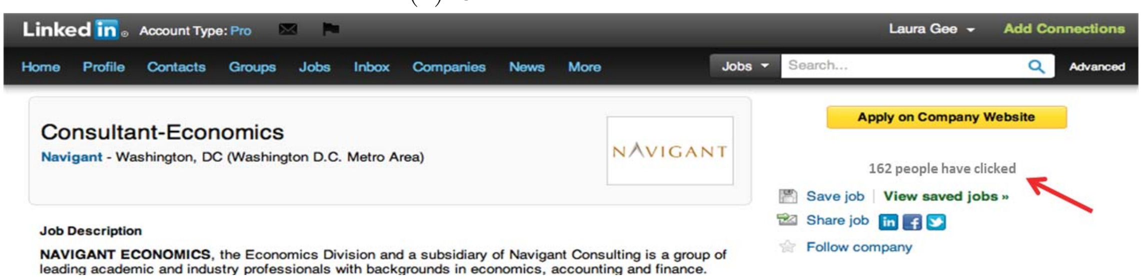

(b) Treatment

Figure 4: Job Posting As Seen In Control and Treatment

Note: This figure shows the way a job posting would be seen by those in the control (Panel (a)) and the treatment (Panel (b)) groups. The arrow in Panel B is to highlight the treatment for the reader, and was not shown to subjects in the experiment. The difference is those in the treatment see that "162 people have clicked" on this job posting to begin an application on the exterior website. Apart from this difference, the job posting is displayed identically to those in the control and treatment groups.

who viewed a job posting during a 16 day window in March 2012. One-fourth of these were randomly assigned to the treatment group and the remaining three-fourths were assigned to the control group.[13]

Overall, the sample includes about 2.3 million registered members from 235 countries. There are about 580,000 job seekers in the treatment and 1.7 million job seekers in the control. During the experiment, these job seekers viewed a total of over 100,000 job postings from 23 thousand companies. On average each job posting was viewed 80 times during the 16 days of the experiment and each company had about 4.7 jobs posted.[14]

3.1 Summary and Balance Statistics

The summary statistics for the subjects in the experiment are provided in Table 1. Gender is identified for 90% of the sample (57% male and 32% female).[15] For the subjects, the average age is 35, and the average year when she became a LinkedIn member is 2009. Furthermore the statistics show that 42% of the subjects are from the US, with an average of 315-316 links as of Spring 2013.[16]

[13] I exclude members who were included in a previous pilot study that took place in the two weeks before the main experiment.

[14] The minimum number of views during the 16 day period was 1 and the maximum was 6,740 with 44 being the median number of views. The minimum number of job postings from a company was 1 and the maximum was 2,568, with the median number of postings from a company being 1. Only 78 companies have 100 or more job postings up during the experiment, and the results are similar if I exclude postings from these companies in the analysis (results available from the author by request). Postings viewed by members in the control and treatment both started with an average of 17-18 previous applicants at the beginning of the experiment.

[15] Members do not provide gender, but it is imputed from their country and name (e.g. Laura in the US is coded female, and Miroslav is coded male in Slovakia). Since a large portion of the analysis concentrates on heterogeneous treatment effects by gender, a balance table by gender is provided in the online Appendix. All observable variables are similar across the control and treatment for both men and women. Also, members do not actually provide age, but it is imputed from the year the person graduated from college or high school.

[16] A "link" is a connection between two LinkedIn members that must be approved by both members. For example, a person may ask to be "linked" to a co-worker, and then that co-worker must approve that link before it appears on the website. LinkedIn keeps records of the number of connections at a company at the time of viewing, but they do not keep systematic records of the total number of links at the time of viewing.

The subjects are very well educated, with 2% listing an Associates degree, 52% listing a Bachelors, and 46% listing a post-Bachelors degree as their highest education level attained. Overall, subjects in the control and treatment groups are similar on observable variables. There is a statistically significant difference between the proportion of subjects from the US between the two groups, but the magnitude of this difference is extremely small. Finally, the statistics in Table 2 show that subjects in the control and the treatment groups view similar postings.

Table 1: Member-Level Summary Statistics

Variable	Mean (All)	N (All)	Mean (Control)	N (Control)	Mean (Treatment)	N (Treatment)	Min.	Max.	t-test for diff.
male	0.572	2,326,207	0.572	1,743,880	0.571	582,327	0	1	0.666
female	0.328	2,326,207	0.328	1,743,880	0.328	582,327	0	1	0.200
gender known	0.899	2,326,207	0.900	1,743,880	0.899	582,327	0	1	0.639
age	34.845	1,837,316	34.850	1,378,146	34.831	459,170	17	136	1.089
year membership	2008.938	2,304,683	2008.938	1,727,755	2008.939	576,928	2003	2012	0.041
US	0.419	2,326,207	0.419	1,743,880	0.418	582,327	0	1	2.233
total links (2013)	315.439	2,305,208	315.220	1,727,947	316.094	577,261	0	40,500	1.091
high school listed	0.002	1,058,647	0.002	797,023	0.002	261,624	0	1	0.408
assoc. listed	0.018	1,058,647	0.018	797023	0.018	261,624	0	1	0.183
BA listed	0.519	1,058,647	0.518	797,023	0.520	261,624	0	1	1.545
post BA listed	0.461	1,058,647	0.462	79,7023	0.460	261,624	0	1	1.562

Notes: In this table each observation is a single member. Each member occurs multiple times in the actual data set, once for each job posting the member views.

Table 2: Posting-Level Summary Statistics

Variable	Mean (All)	N (All)	Mean (Control)	N (Control)	Mean (Treatment)	N (Treatment)	Min.	Max.
start prev. apps	17.434	109,233	17.511	108,675	18.027	104,530	1	3,320
unixtime 1st seen	1332.517	109,233	1332.514	108,675	1332.502	104,530	1332	1334
firm total postings	4.726	23,115	4.727	23,107	4.756	22,926	1	2,568

Notes: "start prev. apps" is the average number of previously started applications before the experiment began. "unixtime 1st seen" is the timestamp coded into unixtime (a common measure of date used by Internet companies) when the job posting was first viewed during the experiment. Most job postings and companies are seen at least once by both the control and treatment groups. In rows 1-2 the observations are at the job posting level. In row 3, the observations are at the firm level.

4 Results

This study examines how varying the information job applicants are shown impacts their subsequent application choices. Showing job seekers the number of previous applicants may impact their choices through three possible mechanisms. First, if job seekers are ambiguity averse, showing them more information would decrease the overall ambiguity and thus increase the likelihood they will apply. Second, if job seekers avoid congested job postings, seeing a higher number of applicants would deter them from applying. Third, if job seekers herd toward more popular job postings, seeing a higher number of applicants would induce them to apply. To test which effect dominates, I make the following related set of predictions: (1) if ambiguity aversion dominates, the treatment will have a positive effect on applications, (2) if competition avoidance dominates, the treatment will have a negative effect on applications if the number shown is sufficiently high, (3) if herding dominates, the treatment will have a positive effect on applications if the number shown is sufficiently high.

There are differing welfare implications from each of these three effects. If ambiguity aversion is the dominant effect, then this could be welfare enhancing by increasing the thickness of the market and possibly decreasing the gender occupation gap. However, if instead job seekers are avoiding competition, then this could be welfare enhancing by decreasing congestion, but it could also lower the number of female applicants. Last, if the dominant effect is herding toward more popular jobs, there may be too much congestion, resulting in a welfare loss.

It is an empirical question which effect is dominant in the data. I begin by presenting the results for the overall treatment effect. I then proceed to provide the result of my tests for competition avoidance and herding by exploring the size of the treatment effect by the change in the number of applicants seen. Last, I show differential treatment effects by the type of job being applied to.

4.1 Overall Treatment Effect

Since each viewing is coded as a separate observation, I have 8,904,039 viewing/posting combinations. The outcome variables are (1) whether a person starts an application and (2) whether a person finishes an application. As explained, I can observe starting an application for both exterior and interior job postings, while I can observe finishing an application for only interior job postings. One can think of the outcome variables over two groups: those who saw an exterior posting (4,499,007 observations), and those who saw an interior posting (4,405,032 observations). The data include all the postings that a member views during the experiment, so the same member often shows up in both the Exterior and Interior sub-samples.

When a job seeker decides to apply to a job posting, she is faced with a number of unknown risks: the probability of an offer, the probability the position is a good fit, the probability of liking the corporate culture, and so on. Ambiguity aversion describes a preference for known versus unknown risks. So, for example, an ambiguity averse job seeker might prefer to apply to a job posting with a known 50% chance of an offer, rather than a posting where the odds are unknown.[17] This experiment varies the amount of information a job seeker receives and thus decreases the ambiguity and by consequence should change the behavior of ambiguity averse job seekers. In particular, I predict that it should change the behavior of female applicants because previous work has found that women are more ambiguity averse than men regarding gains in contextual environments (Moore and Eckel, 2003; Eckel and Grossman, 2008; Schubert et al., 2000). Furthermore, both Samek (2015) and Flory et al. (2015) present evidence from field experiments to show that compensation uncertainty either in the form of a tournament or an uncertain bonus, has a negative effect on women's application rates. An important distinction between this paper and the work of Samek (2015) and that of Flory et al. (2015) is that here the uncertainty is about the probability of an offer or attributes of the potential position, whereas in their studies the uncertainty is in the amount of compensation contingent on being hired. Given the findings of previous work, if ambiguity aversion is driving my results, one would expect the treatment to have a larger effect on female job seekers.

In the first three columns of Table 3, I present the results from a simple

[17]This pattern of decisions can be explained by a number of models, including max-min expected utility or bundled risky decision making (see the online Appendix for a short discussion of Ellsberg (1961); Gilboa and Schmeidler (1989); Halevy and Feltkamp (2005) and Halevy (2007)).

regression:

$$A_{i,d,j} = \beta T_i + \epsilon_{i,d,j}, \tag{1}$$

where each observation is a user i who viewed a job posting j on day d. In Panels A and B.i, the dependent variable $A_{i,d,j}$ takes the value of 1 if that user *started* that job application by clicking on the "apply" button. In Panel B.ii, the dependent variable $A_{i,d,j}$ takes the value of 1 if that user *finished* that job application by submitting all the requested materials. Note that the results in Panel B.ii indicate the unconditional likelihood of finishing an application, meaning that the dependent variable takes the value of 0 either if a person did not start the application or if the person started but did not finish the application. The reason that B.ii concentrates on the unconditional finish rate is that the randomization does not control for selection into starting an application.

Since my dependent variable takes the value of 0 or 1, a logit model would be appropriate. However, since I am most interested in the average probability of applying I use a linear probability model.[18] The independent variable T_i takes the value of one if a user is assigned to the treatment group and thus sees the number of previously started applications. All standard errors are clustered at the job posting j level.

Column 1 of Table 3 shows the results for all LinkedIn users, Column 2 shows the results for female users, and Column 3 shows the results for male users. Looking at Column 1, we can see that the treatment increases the likelihood a user will start and/or finish an application by 0.044 to 0.238 percentage points; representing a proportional increase above the control mean of between 0.855%-1.929%, representing an economically significant potential increase of over a thousand of applications per day during the 16 days of the experiment.[19] As a robustness check, I rerun the analysis with only the first

[18] A logit model yields similar results and those results are available from the author upon request.

[19] A back of the envelope calculation would be that the 2.3 million users viewed almost 9 million job postings. If they had all been in the treatment group, we would have expected an extra 16,699 applications to have been started over the 16 days of the experiment, assuming that those who apply are not substituting this application for another. This does not appear

job posting viewed and find larger effect sizes representing between a 2.124%-3.706% increase over the control (see Appendix).[20] This finding suggests there is no path dependent bias in the sequence of postings viewed.

I next compare the results for the female users in Column 2 to those of the male users in Column 3. This comparison shows the effect of the treatment is always larger for female job seekers. For example, the results in Panel B.ii indicate that the treatment increases the likelihood a female user will finish an application by 0.200 percentage points compared to an insignificant coefficient for male users of -0.033.[21] Furthermore when comparing the results in Column 2 to those in Column 3, we can see that the positive and significant effect of the treatment on starting and finishing applications is largely driven by female LinkedIn users being induced to apply.

The differences described so far may be driven by a number of factors including selection of job posting, order of viewing, and the actual number of applicants displayed. I next test these explanations using the following model:

$$A_{i,d,j} = \beta T_i + P_j + D_d + \alpha NumApply_{i,d,j} + \gamma O_{i,d,j} + \epsilon_{i,d,j}. \tag{2}$$

Note that the dependent variable $A_{i,d,j}$ still takes the value of 1 if a user decides to start or finish an application after viewing the posting. The independent variable T_i takes the value of 1 if the user was assigned to the treatment group which sees the number of previously started applications. I include a fixed effect P_j for each job posting j, so that the treatment identifies differences

to be the case, since those in the treatment start about 0.548 applications on average while those in the control start about 0.539 applications ($t = 2.293$). This difference in total applications is driven by a statistically significant increase for female job seekers, but a non-detectable effect of the treatment on males. Additionally, it is driven by more female job seekers being induced to apply (the extensive margin) rather than by women applying to more jobs (intensive margin), as discussed later.

[20] Since each observation is a user-job pair, users who look at many jobs, and jobs that are particularly popular have more observations in the data. One may worry that the results are being driven by these heavier users or the more popular jobs but if I weight the results so that either each user's weights sum to one or that each job posting's weights sum to one, the results are similar (see Appendix).

[21] The male and female coefficients are always statistically significantly different from each other with the exception of those in panel B.i Column 2 vs. those in Column 3.

Table 3: Likelihood of Starting/Finishing an Application

	Simple			With Fixed Effects		
	1	**2**	**3**	**4**	**5**	**6**
	A. Exterior: Likelihood Starting Application					
	All	**Female**	**Male**	**All**	**Female**	**Male**
Control Mean $\overline{A}_{T=0}$	12.333	11.312	12.471	12.333	11.312	12.471
Treatment β	0.238***	0.365***	0.095*	0.236***	0.409***	0.083
	(0.036)	(0.062)	(0.048)	(0.036)	(0.063)	(0.048)
Adj R2	0.000	0.000	0.000	0.040	0.044	0.039
N	4,499,007	1,477,866	2,562,137	4,499,007	1,477,866	2,562,137
Pct Increase $\frac{\beta}{\overline{A}_{T=0}}$	1.929%	3.226%	0.761%	1.913%	3.615%	0.665%
	B.i Interior: Likelihood Starting Application					
	All	**Female**	**Male**	**All**	**Female**	**Male**
Control Mean $\overline{A}_{T=0}$	15.901	14.198	16.383	15.901	14.198	16.383
Treatment β	0.136***	0.107	0.046	0.102**	0.193**	-0.028
	(0.040)	(0.067)	(0.054)	(0.039)	(0.067)	(0.053)
Adj R2	0.000	0.000	-0.000	0.055	0.060	0.056
N	4,405,032	1,414,655	2,554,216	4,405,032	1,414,655	2,554,216
Pct Increase $\frac{\beta}{\overline{A}_{T=0}}$	0.855%	0.753%	0.280%	0.641%	1.359%	-0.170%
	B.ii Interior: Likelihood Unconditional Finishing Application					
	All	**Female**	**Male**	**All**	**Female**	**Male**
Control Mean $\overline{A}_{T=0}$	4.422	3.681	4.740	4.422	3.681	4.740
Treatment β	0.044	0.200***	-0.033	0.030	0.208***	-0.041
	(0.023)	(0.037)	(0.031)	(0.023)	(0.038)	(0.031)
Adj R2	0.000	0.000	0.000	0.021	0.020	0.022
N	4,405,032	1,414,655	2,554,216	4,405,032	1,414,655	2,554,216
Pct Increase $\frac{\beta}{\overline{A}_{T=0}}$	0.994%	5.431%	-0.696%	0.575%	5.649%	-0.864%

Notes: The dependent variable takes the value of 1 if a job seeker started or finished an application. All coefficients are multiplied by 100 for ease of reading results. Columns 1, 2 & 3 are simple models that only use the treatment as the right hand side variable. Columns 4, 5, & 6 include job posting fixed effects, days posted (omitted category 1st day) fixed effects and categorical dummies for the previous number of people who started a job application at the time of viewing (omitted category is 1-24, other bins are 25-49, 50-74, 75-99, 100-124, 125-149, 150-175, 175+). All standard errors are clustered at the job posting level. The coefficients for male vs. female job seekers are statistically significantly different from each other for all comparisons except panel B.i Column 2 vs. 3. Details of the tests are as follows: panel A ($Prob > chi2 = 0.0006$ for Column 2 vs. 3 and $Prob > chi2 = 0.0000$ for Column 5 vs. 6); panel B.i ($Prob > chi2 = 0.4819$ for Column 2 vs. 3 and $Prob > chi2 = 0.0089$ for column 5 vs. 6); panel B.ii ($Prob > chi2 = 0.0000$ for column 2 vs. 3 and $Prob > chi2 = 0.0000$ for Column 5 vs. 6) Legend: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

in the likelihood of applying between two members viewing the exact same posting.[22] This posting fixed effect controls for time invariant attributes of a posting such as firm, industry, job description, pay range and job title. Additionally, to mitigate time trends in the raw data I use a fixed effect, D_d, for the number of days the posting has been live during the experiment. I also include the variable $O_{i,d,j}$ which controls for the order in which a job posting is seen by person i. In addition I created a set of categorical variables, $NumApply_{i,d,j}$, that divides the true number of previous applicants into eight bins: (1) 1-25, (2) 25-49, (3) 50-74, (4) 75-99, (5) 100-124, (6) 125-149, (7) 150-174 and (8) 175+. This variable controls for the true underlying number.[23]

Columns 4-6 in Table 3 present the results from this model, controlling for time invariant attributes of the job posting, the number of days the posting has been online, the order in which postings are seen, and the true number of previous applicants at the time of viewing. This analysis yields results similar to those in Column 1-3, specifically the treatment increases the likelihood a user will start or finish an application by 0.030-0.236 percentage points, representing a proportional increase above the control mean of between 0.575%-1.913%, or a potential increase of a thousand applications per day.

I next compare the results for female users in Column 5 to those for male users in Column 6. These results show that the coefficient for female job seekers is always statistically significantly larger than that for males. For example, from Panel B.ii, we see that the treatment increases the likelihood a female user will finish an application by 0.208 percentage points, compared to an insignificant coefficient for male users of -0.041.[24] Overall, the results indicate that being in the treatment group increases the likelihood a female job seeker will finish an application (Panel B.ii) by almost 6%. These results are summarized below.

[22]Only 1.4% of the job postings were seen by only a single person during the experiment so the fixed effects have a minimal effect on the effective sample size.

[23]I add this variable since LinkedIn may use this information to select which job postings to highlight for both the control and the treatment groups.

[24]The male and female coefficients in Panel B.ii are statistically significantly different from each other ($Prob > chi2 = 0.000$).

Result 1*: Showing job seekers the number of previously started applications increases the likelihood they will start or finish an application by about 0.6% to 2%; this represents a potential increase of a thousand applications per day. The increase caused by the treatment is similar with or without controls for time invariant attributes of the job posting, the number of days the posting has been online, the order in which postings are seen, or the true number of previous applicants at the time of viewing.*

Result 2*: The increase in applications due to the treatment is largely driven by female job seekers being induced to start or finish an application. The size and significance of the coefficient on the treatment is almost always larger for female vs. male job seekers. For example, being in the treatment group increases the likelihood a female job seeker will finish an application by almost 6%, whereas the effect on men is not statistically significantly different from 0.*

These results suggest that providing job seekers with the number of previous applicants may be a means of increasing the overall number of female applicants to a posting. This increase would reduce the occupation gender gap without putting an undue burden on hiring managers, as the average number of applicants for an interior (exterior) posting is 4.5 (8.4).[25]

I further explore whether the observed increase reflects new applicants (extensive margin) rather than an increase in applications from current applicants (intensive margin). For women who have submitted at least one application, women in both the control and the treatment group start an average of 1.71 exterior applications and finish an average of 2.1 interior applications.[26] This

[25]Recall that the randomization takes place at the user level, not the job posting level. Thus, each job posting appears in both the control and the treatment groups. As a result comparison of the total number of applications from the control vs. the treatment groups would not be useful.

[26]These averages are not statistically significantly different from each other for women who have submitted at least one application ($t = 0.198$ and $t = 0.493$). However, when looking at all women (including those who did not submit at least one application), then the control group starts 0.318 exterior applications and finishes 0.068 interior applications

finding shows that the number of applications on the extensive margin for women in the treatment increases, suggesting that the treatment seems to be adding to the thickness of the female applicant pool by encouraging women who would not have otherwise started an application to apply.

4.2 Treatment Effects By Number

Intuitively it seems plausible that the actual number of previous applicants seen makes a difference in how a subject responds to this information. On the one hand, if job seekers want to avoid applying to postings with greater competition, we should see a decrease in the treatment effect if the number shown is perceived as larger. On the other hand, if job seekers herd toward more popular postings, we should see an increase in the treatment effect as the number shown is perceived as larger (see the online Appendix for a short discussion of herding models (Banerjee, 1992; Anderson and Holt, 1997)).

The reason that I concentrate on the perception of the number shown (rather than the number itself) is that survey evidence finds that people viewing the exact same number may have different opinions on whether it signals high or low competition (see the online Appendix).[27] To measure perceived magnitude, I compare the number being currently viewed and the number seen previously. For example, I might compare the difference between the number of applicants seen for the 2nd job to the number of applicants seen for the 1st job.[28] The number of applicants seen for the previous posting acts as a reference point with which to compare the current posting. If the current posting applicant number is higher, then we would expect a person avoiding competition to be less likely to apply. However, if a person is herding, then

on average vs. women in the treatment group who start 0.329 exterior applications and finish 0.072 interior applications on average ($t = 3.179$ and $t = 2.03$).

[27] As a robustness check I have rerun the analysis using the absolute rather than relative number for the first job posting seen during the experiment. The results are available in the Appendix.

[28] For example, imagine Laura looks at two job postings, and the 1st posting has 10 applicants while the second has 20; here 20 is perceived as a higher number because $20 > 10$. However, if Dan looks at two job postings, and the 1st has 30 applicants while the second has 20, then 20 is perceived as a lower number since $20 < 30$.

a higher current posting applicant number should increase her likelihood of applying. It is likely that both types of individuals exist in the data, but from the perspective of a hiring manager or policymaker it is important to figure out which effect dominates and how it affects the size and composition of the applicant pool.[29]

To test if the magnitude of the treatment effect changes with the number of applicants shown, I use the following model:

$$\begin{aligned} A_{i,d,j_{order=t+1}} &= \beta T_i \\ +\lambda T_i * DIFFAPP_{i,d,j_{order=t+1}-j_{order=t}} &+ \alpha DIFFAPP_{i,d,j_{order=t+1}-j_{order=t}} \\ +P_{j_{order=t+1}} + D_d &+ \gamma O_{i,d,j_{order=t+1}} + \epsilon_{i,d,j_{order=t+1}}. \end{aligned} \tag{3}$$

Note that the dependent variable $A_{ijd,order=t+1}$ takes the value of 1 if a user decides to start or finish an application for the posting seen in order $t+1$. Thus, this analysis excludes the first posting seen. The independent variable T_i takes the value of 1 if a user is assigned to the treatment group. The treatment dummy T_i is interacted with a categorical variable $DIFFAPP_{i,d,j_{order=t+1}-j_{order=t}}$ that represents the difference between the number of applicants for the previous posting and the number for the current posting. Specifically it is a set of categorical variables with the following sixteen bins based on the difference between the number of applicants for the posting being viewed now ($order = t+1$) and the number for the posting last viewed ($order = t$): (1) -176 or lower, (2) -175 to -151, (3) -150 to -126, ... (14) 125-149, (15) 150-174, (16) 175+. The model includes job posting fixed effects $P_{j_{order=t+1}}$ to control for time invariant attributes of the job posting, as well as days posted fixed effects, D_d. In the model, combination of the coefficients β and λ represents the effect of the treatment while holding constant the effect of the job posting and the effect of the numerical difference as measured by

[29] Because the treatment is at the individual level I cannot control for whether a person is a competition avoiding type, or a herding type. Varying the treatment within individuals might address this issue but would likely yield results that would be difficult to interpret.

its bin.

Figure 5 graphically represents the results from this model. On the vertical axis of Figure 5 is the percentage point difference in the likelihood of applying between the treatment and the control groups. On the horizontal axis is the difference in the number of applicants shown in the treatment. The error bars show the 95% confidence interval around each predicted difference. If competition avoidance is the dominating effect, one would expect a downward sloping trend in the panels of Figure 5. On the other hand, if herding is the dominating effect, one would expect to see an upward sloping trend in the panels of Figure 5.[30]

Continuing with the Figure, the top left-hand graph shows the change in the effectiveness of the treatment based on the relative number of applicants shown for all users who view an exterior job posting. The first bar on the far left shows that the treatment increases the likelihood of applying by about 0.25 percentage points above the control group when the job posting being currently viewed ($order = t + 1$) has at least 176 fewer applicants than the job posting the user last viewed ($order = t$). The second bar shows that the treatment increases the likelihood of applying by 0.30 percentage points when a user sees between -175 to -149 fewer applicants than viewed for the previous posting. Neither single point estimate is significantly different from 0.

Overall, the bars in Figure 5 do not illustrate either a strong upward or downward trend as the relative number of applicants shown increases, especially when we consider the noise in our estimates. This noise in the estimates increases as the difference becomes more extreme (either positive or negative). However, this increase may reflect the lower number of observations where users see differences of more or fewer than 100 applicants.[31]

[30]To test whether viewing order creates path dependence, I also analyze the data from only the first job posting viewed. Here, the dependent variable is whether a person applies to the first job posting viewed as a function of the number of current applicants to that first job posting. The results from this analysis do not exhibit a strong pattern of competition avoidance or herding (see the Appendix.

[31]There are actually a large number of observations in each bin. The bin with the fewest observations is 150 to 174 applicants with $N = 24,880$ for an exterior posting and -175 to -151 for interior postings where $N = 32,764$ for an interior posting. Although the bars do

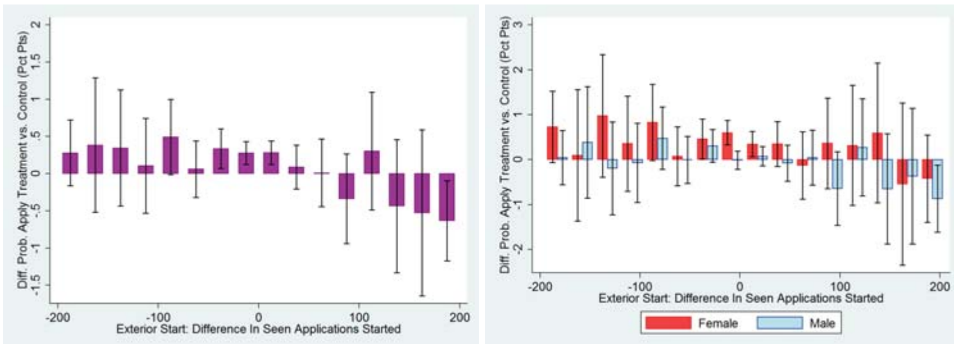

(a) Exterior: Starting Application

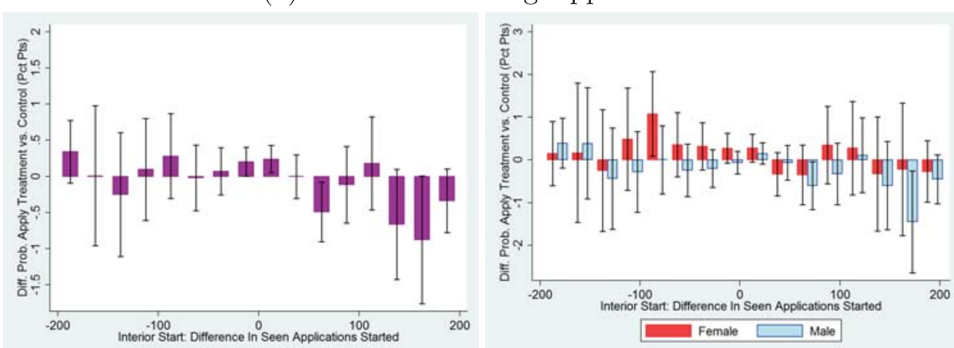

(b) Interior: Starting Application

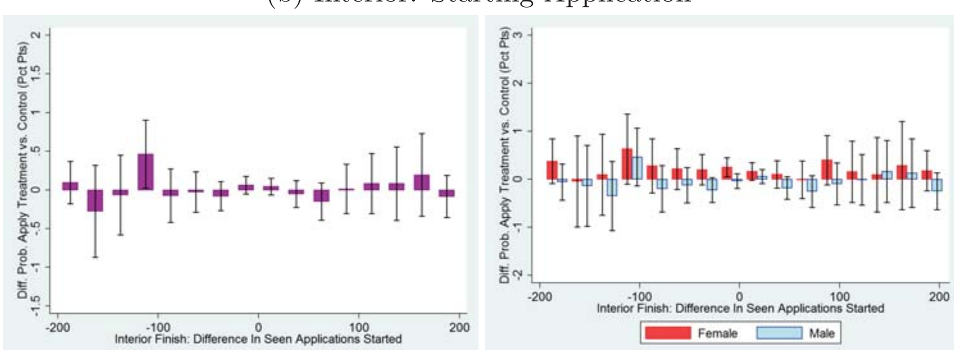

(c) Interior: Finishing Application

Figure 5: Plots of Coefficients on Treatment Dummy Variable By Difference in Number Applicants Shown

In the Figure the top right-hand graph shows the change in the effectiveness of the treatment based on the relative number of applicants shown for female vs. male users who view an exterior job posting. Again, the bars show no strong upward or downward pattern as the relative number shown increases. I find similar results for those starting and finishing an application for an interior job posting (panels (b) and (c), respectively). Recall that competition avoidance would suggest a downward trend while herding would suggest an upward trend as the relative applicant number increases. My finding of no trend could mean that neither effect is present, that the two effects balance each other out, or that my measure does not capture how each individual interprets the number she sees. It is beyond the scope of this study to determine which is the correct interpretation of these results. However, since both competition avoidance and herding may be welfare dis-enhancing, the lack of evidence for these behaviors lends encouragement that the treatment causes no harm.

To gain further insight into the findings in June 2014, I administered an online survey (details available from author upon request). The survey has 188 respondents who were recruited using snowball sampling. This survey presents respondents with a hypothetical scenario to understand how job seekers interpret the number of previous applicants. The survey shows that 50% of respondents use the information to avoid competition, 22% to herd toward more popular jobs, and 27% to avoid ambiguity. While the majority of respondents indicate they use the information to avoid competition, they differ in what number constitutes high competition. Respondents indicate they are more likely to advance to the next stage of the interview process if there are 10 previous applicants versus 100 previous applicants. They also indicated they believe they are more likely to enjoy the position if there are 100 versus 10 previous applicants. These survey results, combined with findings regarding

---

not suggest a linear model, I have run one with the treatment and the interaction of the treatment with the raw difference (available upon request), and find that the interaction with the treatment is always insignificant with two exceptions: (1) the full group and (2) men starting an external application. Even in these cases the coefficient is quite small, implying a one unit increase in the difference seen results in a -0.001 percentage point decrease in the likelihood of application.

treatment effect changes by difference in number seen, lead me to Result 3.

Result 3: *There is no strong evidence that either competition avoidance or herding is the dominant effect from showing a job seeker the number of previous applicants for either male or female job seekers.*

4.3 Treatment Effect By Job Type

Thus far, we have seen that showing the number of previous applicants increases the likelihood of a job seeker starting or finishing an application, and that this increase is larger for female job seekers than male job seekers. These findings have implications for firms actively seeking more female applicants. If a firm is interested in increasing the pool of female applicants it is important to know if the treatment is simply increasing the number of female applicants for "female jobs" or if it raises the likelihood women will apply to traditionally perceived "male" jobs.

For the purposes of this study a "male job," $M_{i,d,j}$, is defined as a job where over 80% of those who start an application in the control group are male. $M_{i,d,j}$ is defined for only those jobs which have at least one person who starts an application in the control group. Consequently, I restrict the sample in this analysis to those jobs with at least one male or female user who starts an application for the job posting in both the treatment and control groups.[32] I use $M_{i,d,j}$ as the dependent variable to test if the treatment increases female applications for these "male" positions. The model is shown below:

$$M_{i,d,j} = \beta T_i + P_j + D_d + \alpha NumApply_{i,d,j} + \gamma O_{i,d,j} + \epsilon_{i,d,j}. \tag{4}$$

Table 4 reports the results from this model. The results in Column 1 of Table 4 show that overall the treatment has a positive effect on the likelihood

[32]This definition is for the outcome variable of *starting* an application. For *finishing* an application, I define $M_{i,d,j}$ as a job with at least 80% males among those who finish the application. In this case, I restrict the sample to those jobs with at least one male or female user who finishes the application in the control and treatment groups.

that any person (male or female) will apply to a "male job." Furthermore, the results in Columns 2 and 3 show that this effect is largely driven by an increase in the likelihood of female applicants applying to "male jobs." This finding provides further evidence of the effectiveness of the treatment in increasing the number of female applicants in industries which are actively seeking to diversify their workforce, and leads to Result 4.[33]

Result 4: *The treatment increases the number of female applicants to "male jobs."*

[33] The proportional gains for the treatment group are also quite large (e.g. a 1.180 percentage point increase from a mean of 1.197 for female users in Panel A), but this is largely driven by the definition of the outcome variable as a job with greater than 80% male applicants in the control group.

Table 4: Likelihood of Applying to a "Male" Job

	With Fixed Effects		
	1	**2**	**3**
A. Exterior: Likelihood Start App			
	All	**Female**	**Male**
Control Mean $\overline{M}_{T=0}$	4.295	1.197	6.0216
Treatment β	0.606***	1.180***	0.340***
	(0.025)	(0.035)	(0.038)
Adj R2	0.129	0.111	0.124
N	3,004,335	1,024,128	1,686,593
B.i Interior: Likelihood Start App			
	All	**Female**	**Male**
Control Mean $\overline{M}_{T=0}$	5.944	1.798	8.148
Treatment β	0.419***	1.026***	0.112**
Adj R2	0.156	0.140	0.148
N	3,508,031	1,153,665	2,016,025
B.ii Interior: Likelihood Finish App			
	All	**Female**	**Male**
Control Mean $\overline{M}_{T=0}$	2.445	0.668	3.426
Treatment β	0.395***	0.741***	0.267***
	(0.025)	(0.034)	(0.038)
Adj R2	0.058	0.054	0.056
N	2,009,987	660,717	1,155,056

Notes: The dependent variable takes the value of 1 if a job seeker started or finished an application to a "male" job. A position is a "male" job if over 80% of the applicants in the control group are male. All the female coefficients are statistically significantly different from the male coefficients (Panel A $Prob > chi2 = 0.0000$; Panel B.i $Prob > chi2 = 0.0000$; Panel B.ii $Prob > chi2 = 0.0000$). All coefficients are multiplied by 100 for ease of reading results. The results include job posting fixed effects, days posted (omitted category 1st day) fixed effects and categorical dummies for the previous number of people who started a job application at time of viewing (omitted category is 1-24, other bins are 25-49, 50-74, 75-99, 100-124, 125-149, 150-175, 175+). All standard errors are clustered at the job posting level. Legend: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

5 Conclusion

This paper uses LinkedIn to conduct a large scale field experiment with about 2.3 million real world job seekers. The results of this experiment show that providing information about the number of previous applicants causes more people to apply to a job posting and that this effect is greater for female applicants. These findings are especially relevant for firms looking to increase the number of female applicants. In short, this paper illustrates a low cost, light touch intervention to reduce the occupation gender gap.

Specifically I find that showing a job seeker the number of previous applicants for a job posting increases the likelihood of application by 0.6-1.9%. Since millions of job seekers view job postings each week on websites like LinkedIn, this translates to an increase in the number of applications of at least a thousand per day.

I also find that the relative number of previous applicants shown does not lead to an increase or decrease in the applications when the relative number shown is high. I interpret this finding as evidence that the dominant effect applicants exhibit is neither competition avoidance nor herding behavior. I thus conclude that the overall positive treatment effect can be explained by models of ambiguity aversion, especially the larger effect observed for female job seekers. Overall, the results indicate that this intervention should not be welfare dis-enhancing since it increases the thickness of the female applicant pool to jobs that particularly need more female applicants.

This paper has focused on the short-term effects of providing applicants with more information during the application process. Research about the long-term effects of providing more information on both unemployment duration and job tenure is an important avenue for future research.

References

Anderson, Lisa R. and Charles A. Holt, "Information Cascades in the Laboratory," *The American Economic Review*, 1997, *87* (5), pp. 847–862.

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

Babcock, Linda and George Loewenstein, "Explaining Bargaining Impasse: The Role of Self-Serving Biases," *The Journal of Economic Perspectives*, 1997, *11* (1), pp. 109–126.

Banerjee, Abhijit V., "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics*, 1992, *107* (3), 797–817.

Barron, John M, John Bishop, and William C Dunkelberg, "Employer search: The interviewing and hiring of new employees," *The Review of Economics and Statistics*, 1985, pp. 43–52.

Bertrand, Marianne, "New perspectives on gender," *Handbook of labor economics*, 2011, *4*, 1543–1590.

_ **and Adair Morse**, "Information disclosure, cognitive biases, and payday borrowing," *The Journal of Finance*, 2011, *66* (6), 1865–1893.

_ **and Sendhil Mullainathan**, "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *The American Economic Review*, 2004.

_ **, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman**, "What's psychology worth? A field experiment in the consumer credit market," Technical Report, The Quarterly Journal of Economics 2010.

Bohnet, Iris, Alexandra Vivien Van Geen, and Max H. Bazerman, "When Performance Trumps Gender Bias: Joint Versus Separate Evaluation," *Management Science*, Forthcoming.

Bougheas, Spiros, Jeroen Nieboer, and Martin Sefton, "Risk-taking in social settings: Group and peer effects," *Journal of economic behavior & organization*, 2013, *92*, 273–283.

Bradley, Gifford W., "Self-serving biases in the attribution process: A reexamination of the fact or fiction question," *Journal of Personality and Social Psychology*, 1978, *36* (1).

Chade, Hector and Lones Smith, "Simultaneous Search," *Econometrica*, 2006, *74* (5), 1293–1307.

Chen, Yan and Joseph Konstan, "Online Field Experiments: A Selective Survey of Methods," *Journal of the Economics Science Association*, forthcoming.

Croson, Rachel and Uri Gneezy, "Gender differences in preferences," *Journal of Economic literature*, 2009, pp. 448–474.

Das, Sanmay and John N. Tsitsiklis, "When is it important to know youve been rejected? A search problem with probabilistic appearance of offers," *Journal of Economic Behavior and Organization*, 2010, *74* (12), 104 – 122.

Dohmen, Thomas and Armin Falk, "Performance pay and multidimensional sorting: Productivity, preferences, and gender," *The American Economic Review*, 2011, pp. 556–590.

Eckel, Catherine C. and Philip J. Grossman, "Chapter 113 Men, Women and Risk Aversion: Experimental Evidence," in Charles R. Plott and Vernon L. Smith, eds., , Vol. 1 of *Handbook of Experimental Economics Results*, Elsevier, 2008, pp. 1061 – 1073.

Eil, David and Justin M. Rao, "The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself," *American Economic Journal: Microeconomics*, 2011, *3* (2), 114–38.

Ellsberg, Daniel, "Risk, Ambiguity, and the Savage Axioms," *The Quarterly Journal of Economics*, 1961, *75* (4), pp. 643–669.

Fernandez, Roberto M. and Collette Friedrich, "Gender Sorting at the Application Interface," *Industrial Relations: A Journal of Economy and Society*, 2011, *50* (4), 591–609.

Flory, Jeffrey A., Andreas Leibbrandt, and John A. List, "Do Competitive Workplaces Deter Female Workers? A Large-Scale Natual Field Experiment on Job-Entry Decisions," *Review of Economic Studies*, 2015.

Galenianos, Manolis and Philipp Kircher, "Directed search with multiple job applications," *Journal of Economic Theory*, 2009, *144* (2), 445 – 471.

Gaucher, Danielle, Justin Friesen, and Aaron C. Kay, "Evidence That Gendered Wording in Job Advertisements Exists and Sustains Gender Inequality," *Journal of Personality and Social Psychology*, 2011, *1* (101).

Gilboa, Itzhak and David Schmeidler, "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics*, 1989, *18* (2), 141 – 153.

Gneezy, Uri, Muriel Niederle, and Aldo Rustichini, "Performance in Competitive Environments: Gender Differences," *The Quarterly Journal of Economics*, 2003, *118* (3), 1049–1074.

Goldin, Claudia, "A grand gender convergence: Its last chapter," *The American Economic Review*, 2014, *104* (4), 1091–1119.

_ **and Cecilia Rouse**, "Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians," *The American Economic Review*, 2000, *90* (4), 715–741.

Halevy, Yoram, "Ellsberg Revisited: An Experimental Study," *Econometrica*, 2007, *75* (2), 503–536.

_ **and Vincent Feltkamp**, "A Bayesian Approach to Uncertainty Aversion," *The Review of Economic Studies*, 2005, *72* (2), 449–466.

Hellerstein, Judith K, David Neumark, and Kenneth R Troske, "Market Forces and Sex Discrimination," *The Journal of Human Resources*, 2002, *37* (2), 353–380.

Kohn, M. and S. Shavell, "The Theory of Search," *Journal of Economic Theory*, 1974.

Kuhn, Peter and Kailing Shen, "Gender discrimination in job ads: Evidence from China," *The Quarterly Journal of Economics*, 2013, *128* (1), 287–336.

Moore, Evan and Catherine Eckel, "Measuring Ambiguity Aversion," *Working Paper*, 2003.

Mortensen, Dale T., "Job Search, the Duration of Unemployment, and the Phillips Curve," *The American Economic Review*, 1970, *60* (5), pp. 847–862.

Nachman, David, "On Risk Aversion and Optimal Stopping," *Working Paper*, 1972.

Neumark, David, Roy J Bank, and Kyle D Van Nort, "Sex Discrimination In Restaurant Hiring: An Audit Study," *The Quarterly Journal of Economics*, 1996, *101* (3), 915–941.

Niederle, Muriel and Lise Vesterlund, "Do Women Shy Away from Competition? Do Men Compete Too Much?," *The Quarterly Journal of Economics*, 2007, pp. 1067–1101.

Ours, Jan Van and Geert Ridder, "Vacancies and the recruitment of new employees," *Journal of Labor Economics*, 1992, pp. 138–155.

Petit, Pascale, "The effects of age and family constraints on gender hiring discrimination: A field experiment in the French financial sector," *Labour Economics*, 2007, *14* (3), 371 – 391.

Roth, Alvin E, "What Have We Learned from Market Design?*," *The Economic Journal*, 2008, *118* (527), 285–310.

Samek, Anya Savikhin, "A University-Wide Field Experiment on Gender Differences in Job Entry Decisions," *Working Paper*, 2015.

Schubert, Renate, Martin Brown, Matthias Gysler, and Hans Wolfgang Brachinger, *Gender specific attitudes towards risk and ambiguity: an experimental investigation*, Institut für Wirtschaftsforschung, Eidgenössische Technische Hochschule, 2000.

Stigler, George J., "The Economics of Information," *Journal of Political Economy*, 1961, *69* (3), pp. 213–225.

Telser, L. G., "Searching for the Lowest Price," *The American Economic Review*, 1973, *63* (2), pp. 40–49.

Vandegrift, Donald and Abdullah Yavas, "Men, women, and competition: An experimental test of behavior," *Journal of Economic Behavior & Organization*, 2009, *72* (1), 554–570.

Weber, Andrea and Christine Zulehner, "Female hires and the success of start-up firms," *The American Economic Review*, 2010, pp. 358–361.

_ **and** _ , “Competition and gender prejudice: are discriminatory employers doomed to fail?,” *Journal of the European Economic Association*, 2014, *12* (2), 492–521.

Weitzman, Martin L., “Optimal Search for the Best Alternative,” *Econometrica*, 1979, *47* (3), pp. 641–654.

Yechiam, Eldad, Meir Druyan, and Eyal Ert, “Observing others’ behavior and risk taking in decisions from experience,” *Judgment and Decision Making*, 2008, *3* (7), 493–500.