

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

Laura K. Gee<sup>1</sup>

<sup>1</sup>Tufts University, Economics (laura.gee@tufts.edu)

October 22, 2017

(Keywords: field experiment, labor search, social information, big data,  
gender, uncertainty, competition)

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

## Abstract

This paper presents the results from a 2.3-million-person field experiment that varies whether or not a job seeker sees the number of applicants for a job posting on a large job-posting website, LinkedIn. This intervention increases the likelihood that a person will finish an application by 3.5%. Women have a larger increase in their likelihood of

---

\*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 and Thomas Tagliaferro for research assistance. Thanks to Yan Chen, Catherine Eckel, Lisa Kahn, Dorothea Kubler, John List, Ioana Marinescu, Gratiana Pol, Maria Recalde, Heather Royer, Olga Stoddard, Lones Smith, Chris Stanton, Sonny Tambe, Isabel Trevino, Myrna Wooders, participants at the American Economic Association Conference, Bay Area Behavioral and Experimental Economics Workshop, Behavioral & Experimental Economics of the Mid-Atlantic Workshop, Bureau of Labor Statistics seminar, Economics Science Association Conference, CeMENT Workshop, Experimental Methods in Policy Conference, George Mason seminar ICES, MIT Sloan, MIT, NBER Summer Institute, Northwestern University NU-Lab seminar, University of Michigan seminar, University of Massachusetts Amherst seminar, and Tufts University seminar for their helpful comments. I would also like to thank three anonymous referees, the associate editor and the editor for their thoughtful and helpful comments.

finishing an application than men. Overall, adding this information to a job posting may offer a light-touch way to both increase application rates and alter the diversity of the applicant pool.

# 1 Introduction

Both firms and job seekers expend a large amount of time and money attempting to match with each other. Understanding how job seekers decide to apply is useful. This paper presents the results from a large (over 2-million-person) field experiment run at the popular business-networking website LinkedIn. The experiment randomly varies whether a job seeker viewing a job posting sees the number of people who have clicked a button to start a job application. Both the control and the treatment group contain LinkedIn members who, collectively, are actively searching through over 100,000 job postings. Job seekers in the control and treatment conditions see identical real job postings.

I find that adding information about others' actions raises the likelihood of application by 1.9%–3.6%. That represents a potential increase of 1,500 started applications each day. There are differences in the increase from the information by observable characteristics like gender and experience, although those differences are not always statistically significant. Adding this kind of information to a job posting may offer a light-touch way to both increase applications and alter the diversity of the applicant pool.

Increasing the applicant pool can be beneficial because vacancies are filled faster when there is a larger applicant pool (Van Ours and Ridder, 1992). Firms and policy makers may want to increase workforce diversity,<sup>1</sup> but a firm cannot hire, for example, more female or black engineers if there is a lack of female or black applicants. Knowing how to encourage a wide range of individuals to apply could increase both the quantity and the diversity of the applicant pool.<sup>2</sup>

Most theoretical models assume that people rely on the most pertinent pieces of information—like the probability of an offer or the utility of the job—when they decide whether to apply.<sup>3</sup> But in reality job seekers may not pay attention to or have access to such information when they make their decision to apply. This paper begins to bridge between the theoretical assumptions of full information and the reality of very little information.

There is a rich history of using field experiments in labor economics. Many field experiments have explored how the demand side of the market decides who to interview by sending

---

<sup>1</sup>See Weber and Zulehner (2014, 2010) and Hellerstein et al. (2002). As an example, in May 2014, Google, noting that only 30% of its overall workforce and 17% of its “tech” workforce is female, acknowledged wanting to increase the diversity of its workforce. See <http://www.forbes.com/sites/jaymcgregor/2014/05/29/2-of-google-employees-are-black-and-just-30-are-women/>.

<sup>2</sup>It is also possible that increasing the number of applicants could lead to too much congestion (Roth, 2008). I explore that topic in the Further Analysis section.

<sup>3</sup>See Lazear et al. (2017); 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).

fictitious resumes/CVs to actual job openings.<sup>4</sup> Yet there are relatively few supply-side field studies. For example, Flory et al. (2015) set up an office assistant position and advertise it in 16 different cities. They find that women are less likely to apply for a job if its description includes more “male”-oriented wording, or alludes to a more competitive pay structure or greater pay uncertainty.<sup>5</sup> Samek (2015) advertises a temporary administrative position she created and finds that a nonprofit framing increases applications, while a more competitive pay scheme deters women from applying more than it deters men. Both Flory et al. (2015) and Samek (2015) vary the description of the position but do not offer any information about the actions of others. In contrast, Coffman et al. (2014) find that stating that 84% of applicants accept their Teach for America (TFA) offer significantly increases an applicant’s likelihood of accepting the offer as well as her commitment to the teaching position.

Many of the previous labor market field experiments have rather limited generalizability because they either rely on researcher-created resumes/CVs or study only one specific type of position. In contrast, this study presents evidence about the behavior of a broad set of real professional job seekers in the context of a wide range of career-oriented job postings. Over the duration of the experiment, these job seekers view over 100,000 different job postings from over 23,000 firms. The experimental results are hence likely to be applicable across various other professional labor markets.

The current number of applicants for a job can be thought of as a piece of social information because it describes the actions of others. Showing social information can increase the likelihood that a person would engage in a variety of behaviors such as applying to college, accepting a job offer, staying at a job, donating to charity, paying taxes, taking environmentally friendly actions, and more.<sup>6</sup>

What separates the present study from the aforementioned experiments is that in most of the previous settings the information represents a clearly positive signal that should increase a person’s likelihood to engage in the desired action. For example, the information that 84% of potential teachers accept an offer is likely an unambiguously good signal about the quality of a TFA teaching position (Coffman et al., 2014). In contrast, in the present experiment it is not obvious if knowing the number of applicants creates a positive or negative signal.

---

<sup>4</sup>See Deming et al. (2016); Eriksson and Rooth (2014); Kroft et al. (2013); Oreopoulos (2011); Lahey (2008); Petit (2007); Riach and Rich (2006); Bertrand and Mullainathan (2004); Neumark et al. (1996).

<sup>5</sup>In a related paper Leibbrandt and List (2014) find that women are less likely to negotiate their wage unless explicitly told the wage is negotiable. Also, Mas and Pallais (2016) use a field experiment with a specific firm for a specific position to explore how people value alternative work arrangements.

<sup>6</sup>See Cialdini et al. (1990); Frey and Meier (2004); Shang and Croson (2006); Martin and Randal (2008); Croson and Shang (2008); Chen et al. (2010); Allcott (2011); Anik et al. (2014); Hallsworth et al. (2014); Mobius and Rosenblat (2014); Smith et al. (2015); Hoxby et al. (2013); Allcott and Rogers (Allcott and Rogers); and Chen et al. (2016).

Such information can be a positive cue on the one hand (more applicants may signal a good job) or a negative one on the other hand (more applicants may signal high competition). The reverse should hold true if a scarcity of applicants for a position is revealed. Given the contradictory effects of this particular type of information—which have the potential to cancel each other out—it cannot be clearly predicted whether learning about the number of applicants generally raises, lowers, or has no impact on overall job application rates.

LinkedIn ran this particular experiment as part of their normal business practices without explicitly making website users aware they were in an experiment, so it can be thought of as a “found” natural field experiment (Alubaydli and List, 2015). Job seekers in the treatment group who saw the number of previous applicants were 1.9%-3.9% more likely to start an application to a position. While the experiment wasn’t designed to reveal the underlying mechanisms for why those in the treatment group were more likely to apply, certain candidate mechanisms, which have heterogeneous treatment effect predictions, can be examined and potentially ruled out. For example, if a herding mechanism is the main driver, then I should observe a positive gradient in the treatment effect for higher numbers of applicants shown (e.g., a job seeker sees 100 applicants and believes this to be a positive signal of job quality, and is hence more likely to apply than if she/he saw only 10 applicants).<sup>7</sup> Conversely, if I observe a negative gradient in the treatment effect for higher numbers of applicants shown, that would be consistent with a competition-aversion mechanism.<sup>8</sup> Since women tend to be more competition averse, finding a more pronounced negative gradient in the treatment effect for higher numbers for women would be evidence of a competition-aversion mechanism.<sup>9</sup> The theoretical prediction for a herding mechanism is a positive gradient by the number shown in the treatment effect, while the theoretical predication of a competition-aversion mechanism is a negative gradient by the number shown in the treatment effect. When I test for the theoretical predictions of a herding or competition-aversion mechanism, the results are not consistent with either mechanism.

There could also be what I will call a “more information” mechanism, whereby simply knowing the number of other applicants reduces information uncertainty and makes job applicants more comfortable with the idea of applying (Gunasti and Ross, 2009). If this mechanism is the main driver, then one would expect experienced job seekers and those

---

<sup>7</sup>See Bougheas et al. (2013); Smith and Sørensen (2011); Yechiam et al. (2008); and Anderson and Holt (1997). An extreme version of competition aversion would be when a job seeker thinks the number of applicants signals that the position has already been filled (Fradkin, 2017).

<sup>8</sup>It is also possible there is a nonlinear relationship, but I find no evidence of that, as shown in Figure 2. Additionally, it is possible that the two effects are washing out, which I discuss in the Further Analysis section.

<sup>9</sup>See Garratt et al. (2013); Dohmen and Falk (2011); Vandegrift and Yavas (2009); Niederle and Vesterlund (2007); and Gneezy et al. (2003).

viewing job postings from well-known firms to be less affected by the treatment of showing the number of applicants. Additionally, if I do observe an overall positive treatment effect, there may be no gradient by the number of applicants shown. This mechanism could also be called an ambiguity- or risk-aversion mechanism.<sup>10</sup> Women are more ambiguity and risk averse, so finding a larger treatment effect for women would be further evidence of a more-information mechanism.<sup>11</sup> The theoretical predictions of a more-information mechanism are no gradient by the number shown in the treatment effect, and a larger treatment effect for those viewing less known firms, the less experienced, and female job seekers. When I test the theoretical predictions of a more-information mechanism, some of the results are consistent with a more-information mechanism.

In sum, I find that adding information about the number of applicants increases the likelihood that job seekers will apply. This illustrates that companies can employ light-touch and low-cost ways to influence the behavior of job seekers in real-stakes situations.

## 2 Field Experiment

### 2.1 Setting

LinkedIn is a large worldwide business networking website with over 350 million members from over 200 regions.<sup>12</sup> LinkedIn has been hosting job postings since 2005, only 2 years after its original launch in 2003.

LinkedIn members are a large and particularly interesting portion of the labor market to study. However, LinkedIn is not representative of the total worldwide labor force. The high-tech and finance industries are heavily represented on this site.<sup>13</sup> These industries have not traditionally had a very diverse workforce.<sup>14</sup> A first step toward a more diverse workforce is a more diverse applicant pool. Because the industries represented on LinkedIn often struggle with diversity, LinkedIn represents an ideal setting for exploring how job seekers decide to apply to job postings.

---

<sup>10</sup>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.

<sup>11</sup>See Garratt et al. (2013); Bertrand (2011); Croson and Gneezy (2009); Eckel and Grossman (2008); Moore and Eckel (2003); Schubert et al. (2000).

<sup>12</sup>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.

<sup>13</sup><http://www.linkedinppc.com/target-by-industry-company-category/>.

<sup>14</sup>For example, only 32.5% of U.S. professionals in STEM-related fields (science, technology, engineering, and mathematics) are female. See <http://dpeaficio.org/programs-publications/issue-fact-sheets/women-in-stem/>.

To use the job postings on LinkedIn, a logged-in member can either use the search bar or access the Jobs landing page (see Appendix Figure 4 and Appendix Figure 5), where she can see a number of job postings that are preselected by LinkedIn based on information the member has listed on her profile, such as education, industry, and previous employment.<sup>15</sup> Then the member can select one of the postings listed. After selecting a posting, the member sees a full-page description of the posting. Every line of the data set represents a logged-in member who has selected to look at a full-page description of the posting. At this point a person may choose to click on a button that reads “Apply on company website” or “Apply now” to start an application, and this is coded as starting an application.

The reason there are two different buttons on the job-posting pages is that LinkedIn has two types of job postings (see Appendix Figure 6): interior postings, which entail LinkedIn collecting the finished application and forwarding it to the firm, and exterior postings, which link a job seeker to an external website. With interior job postings, I can observe if a user both starts and finishes an application.<sup>16</sup> In the case of exterior postings I can observe only if a user starts an application by clicking on the “Apply on company website” button. During the experiment, 60% of the job postings viewed were external.

## 2.2 Experimental Design

LinkedIn designed and ran the field experiment for 16 consecutive days in March 2012 as part of its regular business practices. LinkedIn members who were actively searching through job postings were randomly assigned to either the treatment or the control condition. For the duration of the experiment, each time a member of the treatment group visited a job posting, she saw the number of current applicants for that job, as pictured in Figure 1.<sup>17</sup> The content of the job postings did not differ between the control and treatment conditions. And in fact 95% of viewed job postings were viewed by at least one person in the control and one person in the treatment condition.

LinkedIn chose to randomly assign one-fourth of the active logged-in job seekers to the

---

<sup>15</sup>Note, that job seekers who are not logged-in to a LinkedIn account can also view job postings, but they will not be shown suggestions based on their background. This type of user was not part of the experiment and I do not report on their actions in this paper.

<sup>16</sup>I have the timestamps for when a job seeker clicks “Apply” and for when she submits the 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.

<sup>17</sup>For an exterior job posting, the button read “Apply on Company Website.” For exterior postings the treatment group was shown the number of started applications. For an interior job posting the button read “Apply Now,” and those in the treatment group saw the number of finished applications.

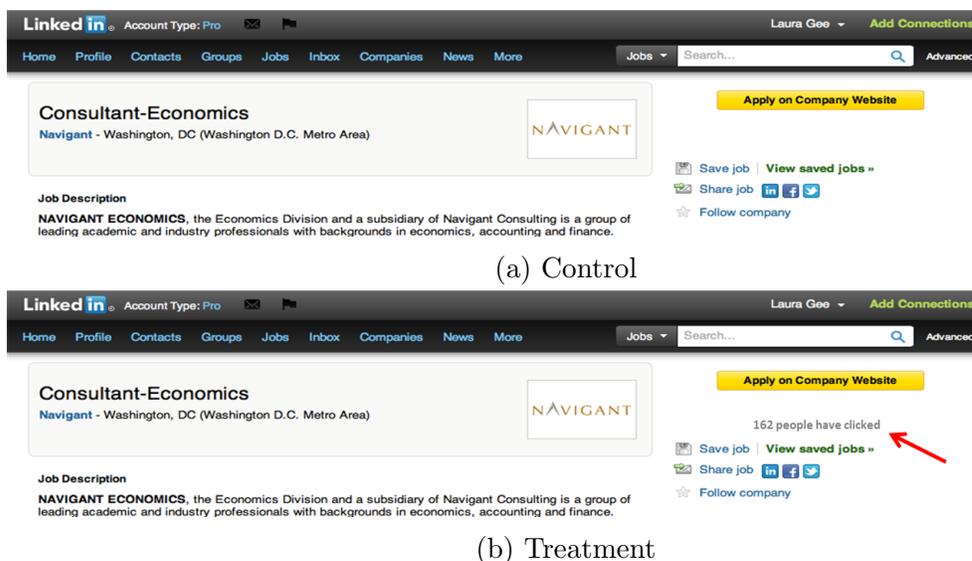


Figure 1: Job Posting as Seen in Control and Treatment

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) serves to highlight the treatment for the reader and was not shown to subjects in the experiment. Those in the treatment group 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.

treatment group and the remaining three-fourths to the control group.<sup>18</sup> This is 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 current number of applicants for that job. 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.

Overall, the sample includes about 2.3 million registered members from 235 countries or areas.<sup>19</sup> There are about 580,000 job seekers in the treatment and 1.7 million job seekers in the control group. During the experiment, those job seekers viewed a total of over 100,000 job postings from 23,000 companies. On average, each job posting was viewed 80 times during the 16 days of the experiment and each firm had about 4.7 jobs posted.<sup>20</sup>

<sup>18</sup>I exclude members who were included in a previous pilot study that took place in the two weeks before the main experiment. I also exclude people who were searching without being logged-in to a LinkedIn account, as they were not randomized into the treatment condition.

<sup>19</sup>There are only 193 UN-recognized countries, but there are about 245 ISO alpha-2 country codes designating different areas.

<sup>20</sup>The minimum number of views during the 16-day period was one and the maximum was 6,740, with 44 being the median number of views. The minimum number of job postings from a firm was one and the maximum was 2,568, with the median number of postings from a firm being one. Only 78 companies had 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 of both the control and the treatment group had an average of 17–18 current applicants at the beginning of the experiment.

The two main outcome variables 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 somebody spent applying or even if the click on the “Apply” button 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.

## 2.3 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 (63.5% male). Age is identified for 79% of the sample (mean = 35).<sup>21</sup> The average year when a person became a LinkedIn member is 2009. About 42% of participants are from the U.S., with an average of 315–316 LinkedIn connections as of spring 2013.<sup>22</sup> The subjects are very well educated, with 2% listing an associate’s degree, 52% listing a bachelor’s degree, and 46% listing a postbachelor’s degree as the highest education level attained. Overall, subjects in the control and treatment groups are similar on observable characteristics. There is a statistically significant difference in the proportion of U.S.-based subjects between the two groups, but the magnitude of this difference is extremely small. Beyond country, I do not know more details of the subjects’ locations.

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	0.635	2,092,347	0.635	1,568,690	0.635	523,657	0	1	0.454
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
U.S.	0.419	2,326,207	0.419	1,743,880	0.418	582,327	0	1	2.233
connections (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	797,023	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.

<sup>21</sup>Members do not provide gender, but it is imputed from their country and name by LinkedIn (e.g., Laura in the U.S. is coded female, and Miroslav is coded male in Slovakia). Also, while members do not provide their age it can be imputed based on the year the person graduated from college or high school.

<sup>22</sup>A “link” is a connection between two LinkedIn members that must be approved by both members. For example, a person may ask to be “connected” to a coworker, and then that coworker must approve the connection before it appears on the website. LinkedIn may keep records of an individual’s number of connections at the time of viewing, but I did not have access to this information.

## 3 Analysis

This study examines how varying the information that job applicants see impacts their subsequent application choices.

### 3.1 Main Results

I can observe starting an application for both exterior and interior job postings, but I can observe finishing an application only for interior job postings. For that reason, I will conduct the analysis over two groups: those who saw an exterior posting, and those who saw an interior posting.

It would be interesting to know whether a person takes a job as well as her tenure at the position. However, since fewer than 3,000 job seekers can be matched to a job at the firm to which they applied, it is impossible to draw any conclusions.

The average number of job postings viewed by both the control and the treatment group was 3.8 (control mean 3.825, treatment mean 3.835,  $t = 0.91$ ) over the 16 days of the experiment. The average number of days a person visits the website by both the control and treatment group was 1.6 days over the 16 days of the experiment (control mean 1.601, treatment mean 1.604,  $t = 1.676$ ). So the treatment does not seem to have a measurable effect on search intensity. This is surprising; one might expect a person in the treatment to search more job postings because each posting contains more information. Recall randomization is at the member level, and 95% of job postings are seen by at least one person in the control and at least one person in the treatment. So by design the job postings seen by the control and the treatment group have the same mean number of current applicants (control mean 71.5, treatment mean 71.6,  $t = -0.38$ ).

For the main analysis I restrict the dataset to the first job posting a person views during the experimental period.<sup>23</sup> If I were to look at all the job postings viewed by members in both the control and treatment, one may worry that for those in the treatment group there would be some path dependence. For example, a person who sees job postings with 100, 15, and 10 current applicants may act very differently than a person who sees two, 15, and 10 current applicants. When I restrict the dataset to the first job posting a person views, that leaves a total of 2,326,207 members for analysis.<sup>24</sup> Those are split roughly evenly between

---

<sup>23</sup>Note that people may have viewed postings before the experimental period, so this is likely not the first posting ever viewed by a person on LinkedIn. However, I discuss the results for those who are newly joined members during the experimental period, and thus may be viewing a job posting on LinkedIn for the very first time, in the Further Analysis section.

<sup>24</sup>The results using all views by all members are quite similar to those reported in the text. However, a summary of the differences would be quite lengthy; if a reader is interested, these results are available from the author upon request. One may also wonder what was the average total number of applications started

first viewing an external job posting (1,134,109) and first viewing an internal job posting (1,192,098).

A simple model might have the left-hand-side variable as  $A_{i,j}$ , which takes the value 100 if member  $i$  starts or finishes an application to job  $j$ , and zero if she does not. In this case, the right-hand-side variable would be  $T_i = 1$  for treatment-group members who see the number of previously started applications, and  $T_i = 0$  for those who do not. By having  $A_{i,j}$  take the value 100 when a person applies the coefficient  $\beta$  can be interpreted as the percentage-point difference in likelihood of application from being in the treatment group.

$$A_{i,j} = \beta T_i + c + \epsilon_{i,j} \quad (1)$$

In such a simple model, one may believe that attributes of the job posting might interact with the treatment; however, I do not observe details of the job posting like industry, title, or job description. I can, however, include a job-posting fixed effect  $P_j$  in the analysis to control for all time-invariant attributes of the job.<sup>25</sup> Since 95% of jobs were seen by at least one person in the control and at least one person in the treatment condition, this does not reduce the sample substantially. I do not know how a member came to view the posting (e.g., through suggestion or via search), but there is no reason to assume that LinkedIn’s background algorithm for suggestions and search would differ between the control and the treatment group. I cannot include a member  $i$  fixed effect because each member is either always in the control or always in the treatment group.

LinkedIn would not reveal the details of their background search and suggestion algorithms, so I control for variables that are likely used by these algorithms, like the numbers of days posting  $j$  has been online during the experiment when viewed by person  $i$  ( $D_{i,j}$ ), and the current number of applicants  $NumCurrApply_{i,j}$  (even when this is not seen in the control).  $NumCurrApply_{i,j}$  is the true number of current applicants. LinkedIn chose to never display if a job posting had zero applicants, so those views of postings with zero current applicants

---

for all job seekers. Those in the treatment group start 0.548 applications versus only 0.539 in the control group ( $t = 2.29$ ). See footnote 55 for more details. One of the main concerns that is alleviated by restricting to the first posting seen is how the path of numbers seen might affect the gradient of the treatment effect by number shown. For example, imagine two people who on the second job posting they view see 15 current applicants. The person who sees job postings with first 100 and then 15 current applicants may act very differently than a person who first sees two and then 15 current applicants. Restricting the data to the first posting seen omits the effects of such path dependence.

<sup>25</sup>For example, imagine a job posting that states, “The typical number of applicants for this position is 40 people.” It seems likely that on such a posting there would be a very atypical effect of the treatment of showing the current number of applicants that might make the overall results difficult to interpret. A job-posting fixed effect controls for such statements within the text of the posting by comparing the behavior of those looking at the same exact posting. I will use the preferred specification with job-posting fixed effects for the rest of the analysis, but results without a job-posting fixed effect are quite similar and are available from the author upon request.

are omitted from the analysis. The number of current applicants ranges from 1 to 4,125.<sup>26</sup>

My dependent variables take two values, so a logit model would be appropriate. However, since I am most interested in the average probability of applying, I present the results from a linear probability model in the main text.<sup>27</sup> This leaves me with the following preferred model:

$$A_{i,j} = \beta T_i + P_j + \gamma D_{i,j} + \alpha \text{NumCurrApply}_{i,j} + \epsilon_{i,j}. \quad (2)$$

Using either the simple model without controls or the more complex model I find that showing the number of applicants significantly increases the likelihood that a person will start or finish an application as presented in Table 2. This increase holds up to a number of robustness checks.<sup>28</sup>

The absolute magnitude of the observed effect ranges between a 0.089- and a 0.355-percentage-point increase, meaning a proportional increase between 1.964% and 3.707%. This may seem small, but given that during the experimental period an average of 500,000 job postings were viewed each day it could lead to a large increase in applications. A back-of-the-envelope calculation suggests that the treatment would result in an extra 1,500 started and 250 finished applications per day.<sup>29</sup> It could also change the final pool of applicants, which I will explore later.

---

<sup>26</sup>Concentrating on the first posting seen, the mean number of current applicants is 71.5 with standard deviation 181.7. The distribution is 25th percentile, 9; 50th percentile, 26; 75th percentile, 68. The variation in number of current applicants is both across all job postings (standard deviation 181.7) and within a given job posting (standard deviations range from zero to 651.8, with the average standard deviation within a job posting being 4.81).

<sup>27</sup>A logit model yields similar results and is presented in Appendix Table 7.

<sup>28</sup>See Appendix Table 7. This table shows that the treatment is robust to using a conditional logit model (Panel C), clustering standard errors at the job-posting level (Panel D), using only jobs seen in both the control and the treatment group (Panel G), and using all the jobs viewed rather than the first job viewed (Panel H). The one robustness check that yields different results entails splitting the sample into U.S. and non-U.S. applicants; here, the coefficients remain positive but lose significance for the U.S. interior job postings (Panel E and F).

<sup>29</sup>First, I assume that those who apply are not substituting this application for another, which seems to be the case given that changes seem to be on the extensive rather than intensive margin, as I will explain later. Second, there are about 275,000 exterior and 280,000 interior postings viewed per day. A 0.349-percentage-point increase in started exterior applications and a 0.208 increase in started interior applications would be a total of about 1,500 started applications. A 0.090-percentage-point increase would be an extra 250 in finished applications. Another way to think about this is the number of extra applications per job. The experiment was 16 days long, so using my back-of-the-envelope calculations that would be an extra 24,000 (1,500 applications per day\*16 days) started and 4,000 finished (250 applications per day\*16 days) applications during the 16 days of the experiment. These would be spread out over the 109,233 job postings (43,291 of these are interior postings) viewed during the 16 days of the experiment. So on average I would predict an increase of 0.219 started applications per posting (24,000 started / 109,233 postings), and 0.092 finished applications per posting (4,000 finished / 43,291 postings). Note these are small increases in the number of applications to any specific job, which implies that hiring managers for individual positions were not adversely overloaded with applications as a result of this experiment. I discuss this further in section 4.2 of the paper.

Table 2: Likelihood of Starting/Finishing an Application

<b>Simple: Without Controls or Fixed Effects</b>			
	<b>1</b>	<b>2</b>	<b>3</b>
	<b>Start Ext</b>	<b>Start Int</b>	<b>Finish Int</b>
Treatment	0.355*** (0.065)	0.225*** (0.065)	0.094** (0.034)
Adj R2	0.000	0.000	0.000
Pct Increase $\frac{\beta}{\bar{A}_{T=0}}$	3.689%	2.125%	3.707%
<b>Preferred: With Controls and Job Fixed Effects</b>			
	<b>1</b>	<b>2</b>	<b>3</b>
	<b>Start Ext</b>	<b>Start Int</b>	<b>Finish Int</b>
Treatment $\beta$	0.349*** (0.065)	0.208** (0.065)	0.089** (0.034)
Days Posted	-0.064*** (0.007)	-0.064*** (0.008)	-0.064*** (0.004)
NumCurrApply	-0.007*** (0.001)	-0.012*** (0.001)	-0.011*** (0.000)
Adj R2	0.049	0.052	0.013
Pct Increase $\frac{\beta}{\bar{A}_{T=0}}$	3.626%	1.964%	3.508%
N	1,134,109	1,192,098	1,192,098
Control Mean $\bar{A}_{T=0}$	9.623	10.589	2.536

Notes: The dependent variable takes the value 100 if a job seeker starts an exterior application (column 1) or interior application (column 2), and zero otherwise. In column 3 the dependent variable takes the value 100 if a job seeker finishes an interior application (not conditional on starting), and zero otherwise. Legend: +  $p < 0.10$ ; \*  $p < 0.05$ ; \*\*  $p < 0.01$ ; \*\*\*  $p < 0.001$

Although the treatment effect is not very sensitive to the inclusion of the control variables, the coefficients on the control variables are statistically significant. The coefficient on NumCurrApply implies that increasing the number of applicants by one decreases the likelihood of application by 0.007–0.010 percentage points. This is a relatively small decrease given that the median number of applicants is 21, and the 90th percentile is 133. While the coefficient on Days Posted shows that for each extra day the posting has been online there is a 0.064 percentage point decrease in the likelihood of application, which is also a small decrease.<sup>30</sup>

## 3.2 Candidate Mechanisms

Understanding the mechanisms behind the increased application rate could allow firms to target those from whom they want additional applications. Unfortunately, the experiment was not designed to trace out the mechanism for why a person is more likely to apply, but some candidate mechanisms have heterogeneous-treatment-effects predictions that I will explore in this section.

### 3.2.1 Competition-Aversion Mechanism

Seeing many applicants could signal that a job is highly competitive and may deter a competition-averse person from applying.<sup>31</sup> In this case, the treatment effect should decrease as the number shown rises. Conversely, if a herding mechanism is the driver, then the treatment effect should be positive for larger versus smaller numbers of applicants shown.

---

<sup>30</sup> Although these coefficients are small in magnitude, the fact they are statistically significant hints at the power that design choices have on people’s actions. Let us suppose that LinkedIn’s background algorithm preferentially pushed jobs with fewer applicants and jobs that had been more recently posted higher in the search results. This would make it more likely for a job posting to be pushed lower in the search results when either NumCurrApply or Days Posted are high. This would be true for both members in the control and members in the treatment group. This would mean both those in the control and the treatment may be less likely to see a job posting with a high value of NumCurrApply or Days Posted, but this should not bias the treatment effect overall or by number shown because it would change the composition of job postings seen in both the treatment and control group in the same way. This could, however, explain the negative coefficient seen on both NumCurrApply and Days Posted when not interacted with the treatment. These postings would be lower in the search results and may therefore be postings that were viewed later in the search process, making a person less likely (in either control or treatment) to apply to them. This could still be true of the first posting viewed during the experimental period, as I don’t have pre-experiment data (so, for example, the first posting viewed during the experiment may have been the fourth posting viewed during this particular job search). Also, I should note that in a previous version of the paper I chose to divide the number of applicants by 100 because the coefficients on the nonscaled variable are extremely close to zero. Those results are available from the author upon request.

<sup>31</sup> Here I use the term *competition averse* to mean someone who, with everything else being equal, is less likely to apply to a job posting with more applicants. Someone would be more competition averse the more she decreased her likelihood of application in response to a single extra applicant.

While both can conceivably take place, from a policy perspective the overall effect is most important. In the Further Analysis section I will show that there are indeed some people who seem competition averse and some that are herding. In this section I will show that under a number of specifications there is no consistent pattern of either competition aversion or herding in the data.

The exact number of current applicants shown on the posting can be thought of as pseudo-random because it is largely determined by when a person is searching on LinkedIn (Smith et al., 2015). To avoid issues of the order of viewing postings affecting the treatment effect, I begin by using only the first posting seen.

I begin by using a nonlinear model to plot the treatment effect by the number shown in Figure 2.<sup>32</sup> On the vertical axis of Figure 2 is the percentage-point difference in the likelihood of applying between the treatment and the control groups. On the horizontal axis is the number of applicants shown in the treatment group. The error bars show the 95% confidence interval around each predicted difference. If competition avoidance is the dominant effect, one would expect a downward sloping trend in the panels of Figure 2. On the other hand, if herding is the dominant effect, one would expect to see an upward-sloping trend.<sup>33</sup> However, there is no clear or consistent pattern of either competition aversion or herding, especially when one takes into account the wide error bars.

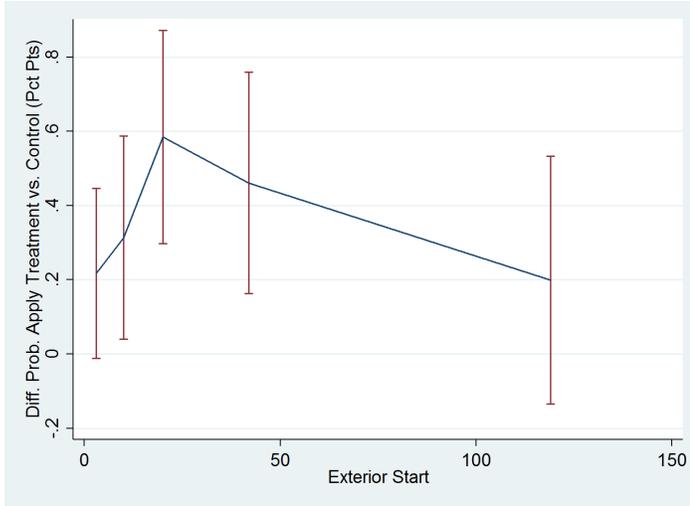
---

<sup>32</sup>I created a model with quantile bins for the number of current applicants. The number of applicants was broken into five equal-sized quantiles  $QNumCurrApply_{i,d,j}$  and then interacted with the treatment as in the equation below:

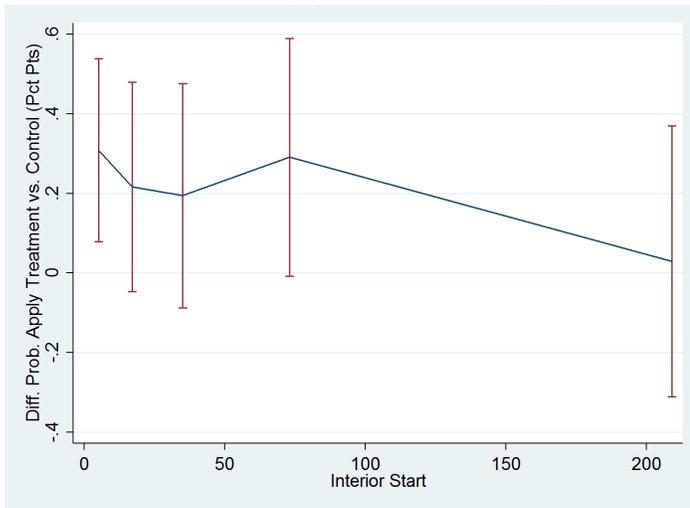
$$A_{i,d,j} = \beta T_i + \lambda T_i * QNumCurrApply_{i,d,j} + \alpha QNumCurrApply_{i,d,j} + P_j + D_d + \epsilon_{i,d,j}.$$

I have also used bins of the numbers 0–25, 26–49, ... 200+, or bins of numbers 0–49, 50–99, ... 200+. The graphs show a similar lack of pattern. Figure 2 graphically represents the results from this model.

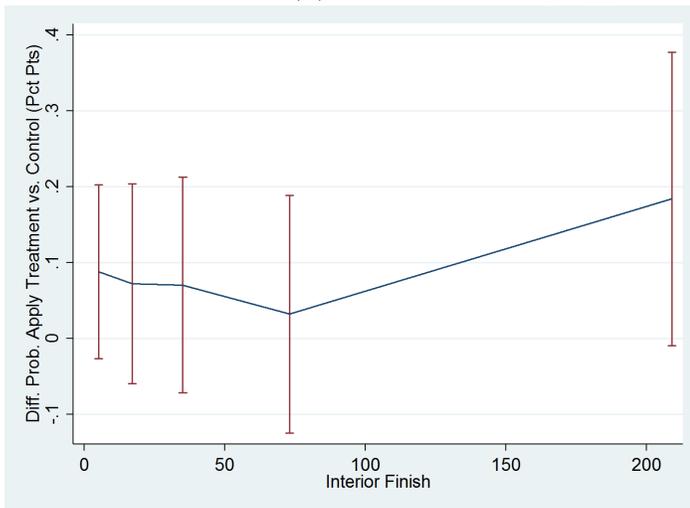
<sup>33</sup>If one uses a model without job-posting fixed effects or the variable days posted the graphs show a similar lack of a pattern. Also, if one uses all the views, the graphs show a similar lack of a pattern. Additionally, if one uses the difference between the number of applicants seen on a current job posting and the number seen on a previously viewed posting, the graphs show a similar lack of a pattern. These graphs and underlying regressions are available from the author upon request.



(a) Exterior: Starting Application



(b) Interior: Starting Application



(c) Interior: Finishing Application

Figure 2: Plots of Coefficients on Treatment Dummy Variable by Number of Applicants Shown

Notes: The coefficients are plotted at the median of each quantile.

Next, I interact the treatment with the number of applicants ( $Treatment * NumCurrApply$ ), but the coefficient on this interaction is not consistent in sign and is statistically insignificant (Table 3, Panel A).<sup>34</sup>

Table 3: Heterogeneous Treatment Effects by Number Shown

	1	2	3
	Start Ext	Start Int	Finish Int
<b>A. First View Only</b>			
Treatment	0.366*** (0.067)	0.211** (0.073)	0.072+ (0.039)
Treatment*NumCurrApply	-0.000 (0.000)	-0.000 (0.001)	0.000 (0.000)
NumCurrApply	-0.007*** (0.001)	-0.012*** (0.001)	-0.011*** (0.001)
Adj R2	0.049	0.052	0.013
N (members)	1,134,109	1,192,098	1,192,098
<b>B. All Views</b>			
Treatment	0.355*** (0.049)	0.205*** (0.051)	0.050+ (0.026)
Treatment*NumCurrApply	-0.001+ (0.000)	-0.000 (0.000)	0.000 (0.000)
NumCurrApply	-0.005*** (0.001)	-0.007*** (0.001)	-0.009*** (0.000)
Adj R2	0.056	0.053	0.019
N (members)	1,134,109	1,192,098	1,192,098
Member-view observations	4,499,007	4,405,032	4,405,032
<b>C. Current Num - Member Specific Avg Num Apply</b>			
Treatment	0.317*** (0.047)	0.192*** (0.047)	0.066** (0.024)
Treatment*(NumCurrApply-MemAvgNumApply)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
(NumCurrApply-MemAvgNumApply)	-0.003*** (0.000)	-0.006*** (0.000)	-0.001*** (0.000)
Adj R2	0.056	0.053	0.018
N (members)	1,134,109	1,192,098	1,192,098
Member-view observations	4,499,007	4,405,032	4,405,032
<b>D. Current Num - Prev Num</b>			
Treatment	0.281*** (0.062)	0.142* (0.063)	0.027 (0.034)
Treatment*(NumCurrApply <sub>t</sub> -NumCurrApply <sub>t-1</sub> )	-0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)
(NumCurrApply <sub>t</sub> -NumCurrApply <sub>t-1</sub> )	-0.002*** (0.000)	-0.003*** (0.000)	-0.001*** (0.000)
Adj R2	0.060	0.055	0.022
N (members)	940,289	932,591	932,591
Member-view observations	3,364,898	3,212,934	3,212,934

Notes: The dependent variable takes the value 100 if a job seeker starts an exterior application (column 1) or interior application (column 2). In column 3 the dependent variable takes the value 100 if a job seeker finishes an interior application (not conditional on starting). Includes job posting fixed effects and days posted. Panel B and C include data for every job posting viewed by the 2.3 million members; observations are weighted so that each member's weights sum to 1, and standard errors are clustered at the member level. Panel D includes data for all but the first job posting viewed by the 1,248,289 members with 2+ views, observations are weighted so that each member's weights sum to 1, and standard errors are clustered at the member level. See Appendix Table 11 for reproductions of these models where the coefficients on the interaction between treatment and measures of number who applied multiplied by 100 and for a model without control variables. Legend: +  $p < 0.10$ ; \*  $p < 0.05$ ; \*\*  $p < 0.01$ ; \*\*\*  $p < 0.001$

So far I have concentrated on the first job viewed so now I expand the dataset, so that

<sup>34</sup>The coefficients on  $Treatment * NumCurrApply$  are  $-0.001$  to  $0.000$  and are very noisy. This implies a weak relationship because the actual range of applicants seen is not that wide ( $NumCurrApply$  median = 26, mean = 71). Also, a quadratic model yields noisy estimates. Results available from author upon request.

it includes the same 2.3 million members but now uses all their views of over 8 million job postings. I find that the coefficients on  $Treatment * NumCurrApply$  are still neither consistent in sign nor statistically significant for starting or finishing an interior job application. Yet for those viewing an exterior job posting, the treatment effect decreases by 0.001 percentage points for every extra applicant increase—a result statistically significant at the 10% level. This provides preliminary evidence in favor of a weak competition-aversion mechanism (Table 3, Panel B).

The noise seen in the data could result from job seekers' inability to interpret the absolute number seen.<sup>35</sup> For example, Ellie might think that 25 applicants is a high number, while Charlie may perceive the same number to be low. I compute the average number of current applicants for all postings viewed by a member ( $MemAvgNumApply$ ).<sup>36</sup> I then use the difference between the number of applicants for the currently viewed posting and the average. The coefficients for the interaction  $Treatment * (NumCurrApply - MemAvgNumApply)$  are zero and not statistically significant (Table 3, Panel C).

Last, an alternative way to benchmark a number as low or high is to use the difference in the number of applicants between the currently and the previously viewed posting ( $DiffNumCurrApply = NumCurrApply_t - NumCurrApply_{t-1}$ ). This requires restricting the data to the 1.2 million members who view at least two job postings.<sup>37</sup> The coefficients for the interaction between  $DiffNumCurrApply$  and the treatment are neither consistent in sign nor statistically significant (Table 3, Panel D).

Since previous work finds that women are more competition averse, one might expect competition aversion to be greater for women than men.<sup>38</sup> The coefficient for the interaction between  $Treatment$  and  $NumCurrApply$  is neither consistently negative nor statistically significant for female job seekers (see Table 4).

The original experiment was not designed to test for competition aversion, though there are some heterogeneous-treatment-effects predictions that would be consistent with competi-

---

<sup>35</sup>To gain further insight into the findings, in June 2014 I administered an online survey meant to uncover how job seekers interpret the number of applicants. This survey presented respondents with a hypothetical job-posting scenario that included the number of applicants. The results show that 50% of respondents use this information to avoid competition, 22% to herd toward more popular jobs, and 27% to avoid ambiguity. While the majority of respondents indicate being competition averse, they differ in what number constitutes high competition. See Appendix Section A.2 for details.

<sup>36</sup>A single average is computed for each person over all postings viewed (pooling exterior and interior) because it seems likely members would not keep a separate average in their head for internal and external postings.

<sup>37</sup>This results in losing about half the sample. This subsample is balanced on observables across the control and treatment. The subsample is similar to the full sample with the exception of having 20 more LinkedIn connections.

<sup>38</sup>See Garratt et al. (2013); Dohmen and Falk (2011); Vandegrift and Yavas (2009); Niederle and Vesterlund (2007); and Gneezy et al. (2003). Note that 94% of jobs are seen by both male and female job seekers.

tion aversion. However, there is no consistent pattern of a decline in the treatment effect for more applicants shown. Another candidate mechanism is a more-information mechanism, which I explore in the next section.

### 3.2.2 More-Information Mechanism

I will use the term “more information” to refer to a mechanism by which simply providing additional information about the job posting increases one’s likelihood of applying. This could be because it is difficult to determine that a posting is legitimate, so seeing the number of current applicants legitimizes the posting. If more information is the main driver, then the treatment effect should be more pronounced for the risk/ambiguity averse (e.g., women), for inexperienced job seekers, and for those viewing postings from lesser-known firms. Additionally, unlike the predictions described in the previous section, the specific number of applicants may not moderate the magnitude of the treatment effect.<sup>39</sup>

Job seekers who are ambiguity averse may experience stronger benefits from more information (Ellsberg, 1961). Ambiguity aversion can be modeled as a specific form of risk aversion (Halevy and Feltkamp, 2005). Since women are generally more ambiguity or risk averse, finding a larger treatment effect in this subpopulation would be evidence of a more-information mechanism.<sup>40</sup> The treatment effect is directionally larger for women than men; however, the difference is only statistically significant for finishing an interior application (see Table 4).<sup>41</sup> Also, it is important to note that the treatment effect for men is only statistically greater than zero for one of the three outcome variables.<sup>42</sup> Last, the treatment effect does not vary by the number shown as evidenced by the insignificant coefficients on *Treatment\*NumCurrApply* and *Treatment\*Male\*NumCurrApply* in Panel B of Table 4. This provides some evidence in support of a more-information mechanism.

In addition to being supportive of a more-information mechanism, the finding that women are more affected than men could be used to increase the number of female applicants. Indeed, large employers of highly skilled workers in the U.S. have recently explicitly stated

---

<sup>39</sup>For example, think of a badge that states “someone has applied” rather than “X people have applied.” Knowing that someone applied still increases the level of information and doesn’t require knowing the specific number of applicants.

<sup>40</sup>See Garratt et al. (2013); Bertrand (2011); Croson and Gneezy (2009); Eckel and Grossman (2008); Moore and Eckel (2003); Schubert et al. (2000).

<sup>41</sup>The coefficients for women are statistically larger if I use all 8 million views and do not cluster standard errors at the member level (results available from author by request and reported in a previous draft of this paper).

<sup>42</sup>In Panel A of Table 4 the linear combination of *Treatment* and *Treatment \* Male* is 0.127  $t = 1.48$  for starting and 0.041  $t = 0.91$  for finishing an interior application. In Panel B of Table 4 the linear combination of *Treatment* and *Treatment \* Male* is 0.120  $t = 1.26$  for starting and 0.023  $t = 0.45$  for finishing an interior application.

they would like to close the gender gap in their firms.<sup>43</sup> Also, previous research finds that increased gender diversity in the workforce has positive results for the firm (Weber and Zulehner, 2014, 2010; Hellerstein et al., 2002). So this may be a light-touch low-cost intervention to increase the number of female applicants, and perhaps eventually the gender balance in some firms. I will speak more about this in the Further Analysis section below.

Table 4: Heterogeneous Treatment Effects by Gender

	<b>1</b>	<b>2</b>	<b>3</b>
	<b>Start Ext</b>	<b>Start Int</b>	<b>Finish Int</b>
<b><u>A. Gender</u></b>			
Treatment	0.383***	0.302**	0.212***
	(0.111)	(0.112)	(0.058)
Treatment*Male	-0.033	-0.174	-0.170*
	(0.141)	(0.141)	(0.074)
Male	1.102***	1.426***	0.498***
	(0.073)	(0.074)	(0.038)
NumCurrApply	-0.007***	-0.012***	-0.011***
	(0.001)	(0.001)	(0.001)
Adj R2	0.049	0.052	0.013
N	1,020,017	1,072,330	1,072,330
<b><u>B. Gender + Number Seen</u></b>			
Treatment	0.421***	0.242+	0.162*
	(0.114)	(0.125)	(0.066)
Treatment*NumCurrApply	-0.001	0.001	0.001
	(0.001)	(0.001)	(0.001)
Treatment*Male	-0.068	-0.122	-0.138+
	(0.144)	(0.155)	(0.083)
Treatment*Male*NumCurrApply	0.001	-0.001	-0.000
	(0.001)	(0.001)	(0.001)
Male	1.102***	1.425***	0.498***
	(0.073)	(0.074)	(0.038)
NumCurrApply	-0.007***	-0.012***	-0.011***
	(0.001)	(0.001)	(0.001)
Adj R2	0.049	0.052	0.013
N	1,020,017	1,072,330	1,072,330
Notes: The dependent variable takes the value 100 if a job seeker starts an exterior application (column 1) or interior application (column 2), zero otherwise. In column 3 the dependent variable takes the value 100 if a job seeker finishes an interior application (not conditional on starting), zero otherwise. Includes job-posting fixed effects and days posted. Legend: + $p < 0.10$ ; * $p < 0.05$ ; ** $p < 0.01$ ; *** $p < 0.001$			

Intuitively, novices need more information than experienced job seekers, so the treatment effect should be larger for novices. I use the number of years a person has been a LinkedIn member as a proxy for job-search experience (mean: 3.047, standard deviation: 2.108, min:

<sup>43</sup>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/>.

0, max: 9).<sup>44</sup> The treatment effect is larger for inexperienced job seekers for two of the three outcomes.<sup>45</sup>

Last, the treatment effect does not vary by the number shown as evidenced by the insignificant coefficients on  $Treatment * NumCurrApply$  and  $Treatment * YearsMem * NumCurrApply$  for two of the three outcome variables, and the significant coefficients are quite close to zero in Panel B of Table 5. These results are supportive of a more-information mechanism.

Table 5: Heterogeneous Treatment Effects by Experience

	1	2	3
	Start Ext	Start Int	Finish Int
<b>A. Experience (Years LinkedIn Member)</b>			
Treatment	0.508***	0.313**	0.216***
	(0.121)	(0.117)	(0.058)
Treatment*YearsMem	-0.051+	-0.036	-0.043**
	(0.030)	(0.030)	(0.015)
YearsMem	-0.255***	-0.378***	0.130***
	(0.016)	(0.016)	(0.008)
NumCurrApply	-0.007***	-0.012***	-0.011***
	(0.001)	(0.001)	(0.001)
Adj R2	0.049	0.053	0.013
N	1,134,109	1,192,098	1,192,098
<b>B. Experience (Years LinkedIn Member)+Number Seen</b>			
Treatment	0.546***	0.284*	0.283***
	(0.125)	(0.130)	(0.064)
Treatment*YearsMem	-0.057+	-0.024	-0.070***
	(0.031)	(0.033)	(0.017)
Treatment*NumCurrApply	-0.001	0.000	-0.001+
	(0.001)	(0.001)	(0.000)
Treatment*YearsMem*NumCurrApply	0.000	-0.000	0.000*
	(0.000)	(0.000)	(0.000)
NumCurrApply	-0.007***	-0.012***	-0.011***
	(0.001)	(0.001)	(0.001)
YearsMem	-0.255***	-0.378***	0.130***
	(0.016)	(0.016)	(0.008)
Adj R2	0.049	0.053	0.013
N	1,134,109	1,192,098	1,192,098

Notes: The dependent variable takes the value 100 if a job seeker starts an exterior application (column 1) or interior application (column 2). In column 3 the dependent variable takes the value 100 if a job seeker finishes an interior application (not conditional on starting). Includes job-posting fixed effects and days posted. Legend: +  $p < 0.10$ ; \*  $p < 0.05$ ; \*\*  $p < 0.01$ ; \*\*\*  $p < 0.001$

Just as inexperienced job seekers need more information than experienced, a person

<sup>44</sup>Age and membership years have a correlation coefficient of 0.31. Including age in the models in Table 5 does not change the sign or significance of the coefficients on  $Treatment * YearsMem$ , but does reduce the sample size (results available from author upon request).

<sup>45</sup>A back-of-envelope calculation finds the treatment to be half as effective after 2.5 more years of LinkedIn membership (see column 3 of Table 5 Panel A). For a person who joined during the present year, the treatment increases the likelihood of finishing an application by 0.216. Since half that effect would be 0.108, each year of membership decreases the treatment effect by 0.043. That implies that after 2.5 years the treatment is half as effective,  $0.108/0.043 = 2.5$ .

viewing a listing from an unknown firm will need more information than a person viewing a listing from a well-known firm, so the treatment effect should be smaller for well-known firms. I identify well-known firms by matching firm name to the 2,000 biggest public firms from Forbes.<sup>46</sup> Because only 13% of well-known-firm job postings are interior, I will concentrate on the results for exterior postings.<sup>47</sup> For exterior job postings the treatment effect is smaller for well-known firms, as shown in column 1 of Panel A of Table 6. Interestingly, for well-known firms the treatment effect increases as the current number of applicants shown increases, yet for less-known firms the treatment effect decreases as the current number of applicants shown increases, as shown in column 1 of Panel B of Table 6. This might be because when one sees a high number of applicants at a firm like Google that signals a higher quality position, but because of Google's size one believes they will simply hire more workers if they interview more qualified candidates.<sup>48</sup> Whereas for a less-known and possibly smaller firm there is more likely to be only a single vacancy, so one views a higher number of applicants as a clearer signal of competition.

This experiment was not designed to test for underlying mechanisms, but evidence of higher treatment effects for women, inexperienced job seekers, and less-known firms are consistent with a more-information mechanism.

---

<sup>46</sup>The Forbes 2000 is a list of the 2,000 biggest public companies (<http://www.forbes.com/global2000/list/>). Research assistants were able to match 1.702% of the firms and 24.021% of job postings, and 19.835% of the observations in the main analysis, to a firm listed on the Forbes 2000 list. Research assistants also attempted to match firm names to those in the Reference USA database, but the match rate was two-thirds as large at 1.095%.

<sup>47</sup>In contrast, 51% of less-known firm job postings are exterior and 49% are interior.

<sup>48</sup>LinkedIn clearly states that each posting is meant to be for a single vacancy, but this policy is likely unknown to job seekers.

Table 6: Heterogeneous Treatment Effects by Firm Type

	1	2	3
	Start Ext	Start Int	Finish Int
<b>A. Known Firm</b>			
Treatment	0.431***	0.253***	0.093**
	(0.077)	(0.069)	(0.035)
Treatment*KnownFirm	-0.262+	-0.304	0.003
	(0.139)	(0.222)	(0.114)
KnownFirm	2.234***	1.336***	0.404***
	(0.070)	(0.111)	(0.057)
Adj R2	0.001	0.000	0.001
N	1,134,109	1,192,098	1,192,098
Postings from KnownFirm	347,918	113,487	113,487
<b>B. Known Firm + Number Seen</b>			
Treatment	0.536***	0.253**	0.075+
	(0.078)	(0.078)	(0.040)
Treatment*KnownFirm	-0.817***	-0.172	0.030
	(0.152)	(0.254)	(0.133)
Treatment*KnownFirm*NumCurrApply	0.009***	-0.002	-0.000
	(0.001)	(0.001)	(0.001)
Treatment*NumCurrApply	-0.002***	-0.000	0.000
	(0.000)	(0.001)	(0.000)
NumCurrApply	0.007***	0.017***	0.006***
	(0.000)	(0.000)	(0.000)
KnownFirm	2.198***	1.165***	0.348***
	(0.072)	(0.115)	(0.060)
Adj R2	0.004	0.009	0.004
N	1,134,109	1,192,098	1,192,098
Postings from KnownFirm	347,918	113,487	113,487

Notes: The dependent variable takes the value 100 if a job seeker starts an exterior application (column 1) or interior application (column 2), zero otherwise. In column 3 the dependent variable takes the value 100 if a job seeker finishes an interior application (not conditional on starting), zero otherwise. Includes days posted as a control variable. These models do not include a job-posting fixed effect. Legend: +  $p < 0.10$ ; \*  $p < 0.05$ ; \*\*  $p < 0.01$ ; \*\*\*  $p < 0.001$

## 4 Further Analysis

Showing the number of applicants increases the likelihood of an application and might be used to increase applications from women, the less experienced, and to less-known firms. As aforementioned, many large firms have publicly announced that they would like to increase gender balance in their hiring. There has been less media coverage of initiatives to hire inexperienced workers or to drive applicants to less-known firms. However, before I can recommend the intervention of showing the number of current applicants to increase applications overall or from any subgroup, there are some further analyses that I need to address in this section.

### 4.1 Are women more likely to apply to masculine jobs?

On LinkedIn men are more likely to start an application than women.<sup>49</sup> About 10% of the job postings have only women applying in the control condition, so increasing the number of female applicants for these jobs does nothing to increase female applications to male-dominated positions. I do not have access to the actual job-posting description, so I cannot use job attributes to categorize jobs as masculine. Instead I define a job as “masculine” if over 80% of those who started or finished an application in the control group were men.<sup>50</sup> Note that although behavior by those in the control group defines a job as “masculine,” individuals in both the control and treatment group view these job postings. I can only determine the proportion of male applicants for jobs that have at least one applicant with gender known in the control. To be consistent I restrict the data to those viewing a job that has at least one applicant of known gender in both the control and the treatment.<sup>51</sup> Using a model without job-fixed effects, I interact the MasculineJob variable with the treatment

---

<sup>49</sup>If I control for the type of job posting with a job-fixed effect, I find that the likelihood of starting an exterior application is 9.775% for men, but only 8.687% for women. Similarly, the likelihood of starting an interior application is 10.599% for men, but only 9.931% for women. Last, the likelihood of finishing an interior application is 2.674% for men, but only 2.179% for women.

<sup>50</sup>To be consistent with the main analysis I use only the first job viewed by those in the control group. About 40% of the jobs are “masculine” if the 80% cutoff is used. The results are similar using other thresholds (results available from author upon request).

<sup>51</sup>Restricting the sample to only those who viewed an application with at least one person of known gender who started an application in the control and treatment results in losing 49% to 58% of the sample. Restricting to only those who viewed an application with at least one person of known gender who finished an application in the control and treatment results in losing 76% of the sample for finishing an interior application. So these are highly selected subsamples. The subsamples are balanced on the observable characteristics across the control and the treatment, so these results should be internally valid. The members in these subsamples look similar to the whole sample, except for the proportion of U.S. members dropping by 15–28 percentage points, and the number of LinkedIn connections being higher by 30–50 (since the rate of starting/finishing an application is higher outside the U.S. and for those with more connections). Therefore these results are less externally valid.

for men and women. Women in the control and the treatment both search for jobs before they are treated with extra information. I find that women who were already searching for “masculine” jobs are more likely to apply to those jobs if they are shown the previous number of applicants. I do not find that being in the treatment group affects which jobs a person (of any gender) searches for, since the treatment takes place after a job posting has already been found.<sup>52</sup> So this light-touch intervention has the desirable effect of increasing the likelihood of a woman applying for a masculine job, which could help ameliorate the gender occupation gap.

## 4.2 Is there too much congestion from the increased application rate?

Even if more women apply for masculine jobs, diversity could still be hindered if hiring managers are being overloaded with too many applicants. If the treatment causes people to apply for jobs that ultimately end up with a large number of applicants, that could actually harm an applicant’s chances of receiving an offer. While the data do not record job offers, I do have data on the final number of applications started or finished for each job by the end of the experimental period.<sup>53</sup> I scale this by the number of days the job posting was online during the experiment to get a measure of job congestion, and create a *Congested* variable that takes the value 1 if a job has an above-average number of final applications per day.<sup>54</sup> I find that the interaction of the treatment with *Congested* is never statistically significant (results available from the author upon request) for either gender. This result means that showing the number of applicants does not cause people to apply for more congested jobs.

---

<sup>52</sup>See Appendix Table 8. Note the coefficient on *MasculineJob* is mechanically negative in this model because it represents the likelihood a female job seeker will apply to a *MasculineJob* in the control group, which by definition is lower than the likelihood of a male job seeker. However, in this model the effect of the treatment on female job seekers viewing a masculine job is the sum of the coefficients on *Treatment* and *Treatment\*StartMasculine*. So the mechanically negative coefficient on *MasculineJob* does not interfere with our interpretation. One could obtain the same results by running the model on only women viewing a so-called *MasculineJob*. Female job seekers in the treatment (versus the control) are 11.696 percentage points more likely to start an application for an exterior “masculine” job, 8.792 percentage points more likely to start an interior “masculine” job application, and 4.719 percentage points more likely to finish the interior application. All effects are statistically significant at the 1% level. Part of the reason why the coefficients are so large is that the starting/finishing rate is about 40% larger for this subsample.

<sup>53</sup>This is different from the number shown, since that is a running tally of applicants.

<sup>54</sup>Recall that over 90% of jobs are seen in both the control and treatment, hence I cannot use the number of applications started/finished in the control.

### 4.3 Is this encouraging new applicants?

Since each woman can ultimately take only one job, increasing the number of jobs she applies to may not actually increase workforce diversity. I therefore explore whether the observed increase reflects new applicants (extensive margin) rather than an increase in applications from current applicants (intensive margin) during the 16 days of the experiment. Keep in mind that all the people in the experiment were already actively looking at job postings, so this section is not about bringing people into the job-seeking process, but rather about whether those who are already seeking a job apply to at least one job or increase the total number of jobs they apply to. Across both genders, as well as for female job seekers alone, the treatment group starts more applications than does the control. However, when looking only at those with at least one application, that difference goes away. This means that the treatment increases the number of applications on the extensive margin.<sup>55</sup> In other words the treatment induces an active job seeker who otherwise would not have applied to any job to apply to one job, rather than inducing an active job seeker who would have applied to a single job to apply to two or more jobs. Since many job searches last longer than 16 days and job seeker in this experiment only visited LinkedIn 1.6 days during the experiment, this is suggestive but not conclusive evidence that the treatment is adding to the thickness of the applicant pool by encouraging those who otherwise would not have started an application to apply.

### 4.4 Is the novelty of the change driving the positive treatment effect?

I find that a change to the way job postings are displayed in the form of showing the number of previous applicants increases the likelihood a job seeker will start and finish a job application. Such a positive treatment effect could also be explained by LinkedIn members reacting positively to any novel change in the way job postings are displayed rather than specifically showing the number of previous applicants. Job seekers who became LinkedIn members during the experimental period would most likely have no previous experience with the old way job postings were displayed, and therefore could not be affected by the novelty of

---

<sup>55</sup>For all job seekers, those in the treatment group start 0.548 applications versus only 0.539 in the control group ( $t = 2.29$ ). Yet looking only at those who apply to at least one job, those in the treatment start 2.532 applications, while those in the control start 2.549—a statistically insignificant difference ( $t = 0.987$ ). The same pattern holds for finishing applications, though the differences are never statistically significant. For female job seekers, those in the treatment group start 0.493 applications versus only 0.481 for this in the control group ( $t = 2.1$ ). Yet for those who apply to at least one job, the difference is not statistically significant ( $t = 0.810$ ).

the change.<sup>56</sup> I create the variable `NewMem` that takes the value 1 if a person created their LinkedIn account during the two week experimental period. The interaction of `NewMem` and the treatment is never statistically significant (results available from author upon request). This implies there is no difference between the treatment effect for those who were never exposed to the old job-posting user interface and those who may have recently been exposed to the old job-posting interface. However, only 72 thousand of the 2.3-million-person sample became LinkedIn members during the experimental period, so the lack of significance could be driven by loss in the sample size.

An additional test that is less sensitive to sample-size issues is to consider that more experienced users in general (not those who specifically joined before the experimental period) may be more likely to be affected by any novel change because they are more used to the old job-posting user interface. However, as shown in Table 5 more experienced users, as measured by years of LinkedIn membership, are actually less affected by the treatment, which implies novelty is not driving the treatment effect.

Last, other researchers have found that displaying popularity information in other settings increases interaction with a web-page (Tucker and Zhang, 2011), yet displaying information about price in advertisements doesn't affect interaction with a web-page (Salganik et al., 2006). So not all novel changes to a web-page increase interaction, but specifically those that have to do with information about the actions of others. Also, Kohavi et al. (2014) report that novelty effects are uncommon in practice, all of which further implies that novelty alone is likely not driving the results.

## 4.5 Are competition aversion and herding both taking place?

Earlier I showed that there is no consistent pattern of competition aversion overall. However, if both competition aversion and herding are taking place simultaneously this could explain the lack of an overall pattern. From a policy perspective, the overall pattern matters, but one may also wonder if there are herding types and competition-averse types. The treatment was assigned at the individual member level, so a member sees the number of applicants either always or never. I restrict the data to those in the treatment group who have some variation in starting an application (97,858 members), because it is not possible to observe herding

---

<sup>56</sup>Note, it is possible for a person who was previously a LinkedIn member to create a new account. However, it seems unlikely that someone who had recent experience with the old way job postings were displayed would create a brand new account during the experimental period. Additionally, because LinkedIn is like an online resume/CV individuals tend to have a single account. It is also possible that someone who had been browsing job postings without being logged-in could then create an account during the experimental period, and so they would be familiar with the old way job postings were displayed. I have no ability to measure if this is the case since I cannot connect the actions of non-logged-in users to newly created accounts.

or competition aversion if someone never applies or always applies. I then compute the correlation between starting an application and the number shown.<sup>57</sup> I find that the mean correlation is 0.043, quite close to zero. However, as shown in Figure 3, almost half of the members have a positive correlation (e.g., herding) and the other half a negative correlation (e.g., competition aversion).<sup>58</sup> This is consistent with herding and competition aversion both occurring simultaneously, which would explain why there is no overall pattern by number seen. Being a herding or competition-averse type can only be determined for those in the treatment group, so I cannot make statements about differential treatment effects across these types.<sup>59</sup>

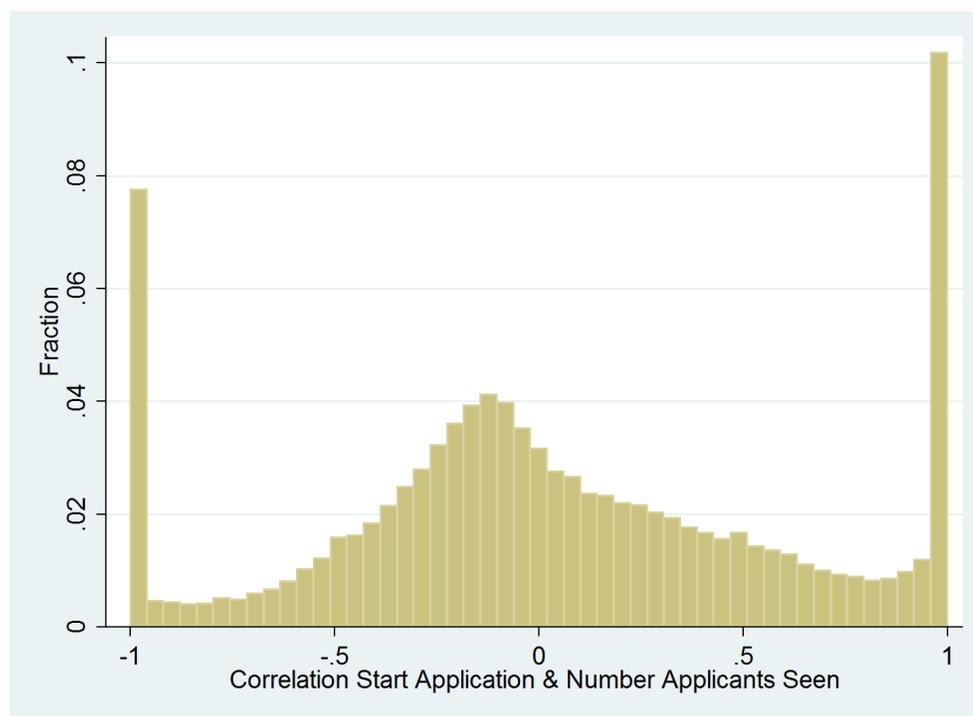


Figure 3: Distribution of Correlation Between Number Seen and Starting Application

Note: This figure shows the distribution of the correlation between number seen and starting an application for those in the treatment who have some variation in whether they apply (98,070). A correlation closer to  $-1$  is evidence of competition aversion. A correlation of 1 is evidence of herding.

<sup>57</sup>I use views of both external and internal job postings because if I restrict to certain types of postings I lose even more of the sample.

<sup>58</sup>The distribution of correlations is quite similar across genders (Appendix Figure 11). One might also wonder if those who saw higher numbers also tended to have more negative correlations. Appendix Figure 12 plots the average correlation by number seen in the treatment. There is no concentration of negative correlations for higher numbers shown.

<sup>59</sup>Note this same pattern of half herding, half competition aversion could also be explained by randomly generated data. But the distributions of types have peaks at  $-1$  (a person who didn't apply when they saw a higher number) and 1 (a person who only applied when they saw a higher number), and are centered around 0 with a smooth tail moving toward either extreme. One might expect purely random data to have a more uniform distribution.

## 5 Conclusion

Previous labor-market field experiments have concentrated on fictitious job seekers, or applicants to specific types of positions (administrative or teaching). This study is able to observe 2.3 million real job seekers on LinkedIn who look at over 100,000 real job postings. I find that showing the number of current applicants on the corresponding job posting increases a job seeker’s likelihood of applying by 1.9%–3.6%. Interestingly, job seekers in both the control and treatment group view about 3.8 job postings, so search intensity is not affected by the treatment.

Understanding the mechanism for increased applications could increase welfare through a better-functioning labor market. If a more-information mechanism dominates, then this may enhance welfare by increasing the thickness of the market and could be used as a strategic tool for targeting minorities. By contrast, if a competition-aversion mechanism dominates, there may be a welfare gain from decreased congestion, but also a decrease in the number of minority applicants. I find that women (including those who apply for male-dominated jobs), inexperienced job seekers, and those looking at less-known firms are more affected by the treatment. This implies that a more-information mechanism may be at play.

Importantly, showing the information does not simply push female applicants toward already female-dominated jobs, and it does not overly increase congestion so that hiring managers would be overloaded. Instead, it brings new job seekers into the applicant pool.

Managers and policy makers often want to increase the diversity of ideas in their workforce through hiring minorities (e.g., women) or less-experienced workers (who may bring novel ideas to the table). The results of this paper imply these managers and policy makers could include more information (like the number of current applicants, but also perhaps the gender or education of typical applicants) as a means to increase the diversity of their applicants. Managers and policy makers could include this in their online job postings, but could also offer this extra information to potential job applicants at job fairs and information sessions.

Previous labor-market field studies have shown that changing the pay structure can result in a large increase in minority applications (e.g., Flory et al. (2015) find that removing competitive pay halves the application gender gap). However, changing the pay structure is a relatively large change to the firm’s business practices. This study finds that simply showing the number of applicants increases the likelihood that a woman will finish an application by 0.162 percentage points, versus a zero effect for men. Although this increase is smaller than one that could be obtained from changing the pay structure, showing the number of applicants is likely a more easily implementable change.

I have presented the results for three outcome variables (starting exterior, starting in-

terior, and finishing interior job applications) over many sub-groups (e.g., male/female, by years of experiences, known/unknown firms). In total there are 157 hypotheses tested in this paper, so clearly multiple hypotheses testing is a concern (List et al., 2016). The Bonferroni correction calls for an adjustment of the acceptable significance level by dividing by the number of hypotheses. In this paper that would mean that results that had an uncorrected p-value of less than 0.0006 would still be significant at the 10% level after the correction. 71% of the hypotheses in this paper hold at the 10% level after correcting for multiple hypothesis testing. Appendix Table 9 lists all the hypotheses tested and whether they survive a multiple-hypotheses-testing correction in detail and discusses how this correction affects the main results. This is clearly a very conservative test, and also brings up the issue that p-values alone should not be the only thing researchers consider when they decide how much to update their priors based on the findings from a single study (Maniadis et al., 2014).

Indeed, another concern about this study may be that the data set contains 2.3 million observations and some may worry with such a large data set that finding statistical significance is almost assured. Again, this points to the greater issue that statistical significance alone should not be the only thing researchers consider when deciding how a single study should help them to update their prior beliefs about whether the findings from that study are indeed true. In Appendix Table 10 I use the methods described in Maniadis et al. (2014) to show the range of the post-study probability that the results are indeed true as a function of low (0.10), medium (0.50), and high (0.90) prior beliefs about the effect of the treatment.<sup>60</sup> Given the large sample size of this study even with a low prior probability about the results, one's updated post-study probability that the results are indeed true falls between 89% to 91%.

A shortcoming of the current study is that it may not generalize to nonprofessional labor markets. LinkedIn is primarily used by those holding a bachelor's degree or higher, so it is not clear that showing the number of current applicants would have the same effect in a labor market for less-educated workers.

Additionally, LinkedIn did not design this study with the hopes of isolating the mechanism for why showing the number of current applicants might change behavior. Partnering earlier in the research process with large firms during field experiments could be beneficial to both furthering our knowledge and the firm's bottom line. However, using "found" experiments like this one is a first step toward showing firms the value that academics can bring to their business practices. It is my understanding that shortly after the initial experiment LinkedIn began to show the number of current applicants on all job postings, then later it made this an optional feature that firms could opt in or opt out of. As of June 2017

---

<sup>60</sup>The power calculations were done using the following tool: <http://clincalc.com/Stats/Power.aspx>.

it appears that LinkedIn no longer allows firms to show the number of applicants on job postings. Behavioral “nudges” have become popular tools for policy makers and firms to influence short-term behavior. Recently, Coffman et al. (2014) have shown that an information nudge has large long-term effects on job applicants accepting and staying at a teaching job. Understanding how light-touch nudges can be used to affect both long-term and short-term behavior in the job market is an important area for continued research.

## References

- Allcott, H. (2011). Social norms and energy conservation. *Journal of Public Economics* 95(9), 1082–1095.
- Allcott, H. and T. Rogers. The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*.
- Alubaydli, O. and J. List (2015). On the generalizability of experimental results in economics. *Handbook of Experimental Economic Methodology*.
- Anderson, L. R. and C. A. Holt (1997). Information cascades in the laboratory. *American Economic Review* 87(5), 847–862.
- Anik, L., M. I. Norton, and D. Ariely (2014). Contingent match incentives increase donations. *Journal of Marketing Research* 51(6), 790–801.
- Bertrand, M. (2011). New perspectives on gender. *Handbook of Labor Economics* 4, 1543–1590.
- Bertrand, M. and S. Mullainathan (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American Economic Review*.
- Bougheas, S., J. Nieboer, and M. Sefton (2013). Risk-taking in social settings: Group and peer effects. *Journal of Economic Behavior & Organization* 92, 273–283.
- Chade, H. and L. Smith (2006). Simultaneous search. *Econometrica* 74(5), 1293–1307.
- Chen, Y., F. M. Harper, J. Konstan, and S. Xin Li (2010). Social comparisons and contributions to online communities: A field experiment on movielens. *American Economic Review* 100(4), 1358–1398.
- Chen, Y., F. Lu, and J. Zhang (2016). Social comparisons, status and driving behavior. *Working Paper*.
- Cialdini, R. B., R. R. Reno, and C. A. Kallgren (1990). A focus theory of normative conduct: recycling the concept of norms to reduce littering in public places. *Journal of personality and social psychology* 58(6), 1015.
- Coffman, L. C., C. R. Featherstone, and J. B. Kessler (2014). Can social information affect what job you choose and keep?
- Croson, R. and U. Gneezy (2009). Gender differences in preferences. *Journal of Economic Literature*, 448–474.
- Croson, R. and J. Y. Shang (2008). The impact of downward social information on contribution decisions. *Experimental Economics* 11(3), 221–233.

- Das, S. and J. N. Tsitsiklis (2010). When is it important to know youve been rejected? a search problem with probabilistic appearance of offers. *Journal of Economic Behavior & Organization* 74(12), 104–122.
- Deming, D. J., N. Yuchtman, A. Abulafi, C. Goldin, and L. F. Katz (2016). The value of postsecondary credentials in the labor market: An experimental study. *American Economic Review* 106(3), 778–806.
- Dohmen, T. and A. Falk (2011). Performance pay and multidimensional sorting: Productivity, preferences, and gender. *The American Economic Review*, 556–590.
- Eckel, C. C. and P. J. Grossman (2008). Chapter 113 men, women and risk aversion: Experimental evidence. Volume 1 of *Handbook of Experimental Economics Results*, pp. 1061–1073. Elsevier.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics* 75(4), 643–669.
- Eriksson, S. and D.-O. Rooth (2014). Do employers use unemployment as a sorting criterion when hiring? evidence from a field experiment. *The American Economic Review* 104(3), 1014–1039.
- Flory, J. A., A. Leibbrandt, and J. A. List (2015). Do competitive workplaces deter female workers? a large-scale natural field experiment on job-entry decisions. *Review of Economic Studies*.
- Fradkin, A. (2017). Search, matching, and the role of digital marketplace design in enabling trade: Evidence from airbnb. *Working Paper*.
- Frey, B. S. and S. Meier (2004). Social comparisons and pro-social behavior: “Testing” conditional cooperation” in a field experiment. *The American Economic Review* 94(5), 1717–1722.
- Galenianos, M. and P. Kircher (2009). Directed search with multiple job applications. *Journal of Economic Theory* 144(2), 445–471.
- Garratt, R. J., C. Weinberger, and N. Johnson (2013). The state street mile: Age and gender differences in competition aversion in the field. *Economic Inquiry* 51(1), 806–815.
- Gneezy, U., M. Niederle, and A. Rustichini (2003). Performance in competitive environments: Gender differences. *The Quarterly Journal of Economics* 118(3), 1049–1074.
- Gunasti, K. and W. T. Ross (2009). How inferences about missing attributes decrease the tendency to defer choice and increase purchase probability. *Journal of Consumer Research* 35(5), 823–837.
- Halevy, Y. and V. Feltkamp (2005). A bayesian approach to uncertainty aversion. *Review of Economic Studies* 72(2), 449–466.

- Hallsworth, M., J. List, R. Metcalfe, and I. Vlaev (2014). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. Technical report, National Bureau of Economic Research.
- Hellerstein, J. K., D. Neumark, and K. R. Troske (2002). Market forces and sex discrimination. *The Journal of Human Resources* 37(2), 353–380.
- Hoxby, C., S. Turner, et al. (2013). Expanding college opportunities for high-achieving, low income students. *Stanford Institute for Economic Policy Research Discussion Paper* (12-014).
- Kohavi, R., A. Deng, R. Longbotham, and Y. Xu (2014). Seven rules of thumb for web site experimenters. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1857–1866. ACM.
- Kohn, M. and S. Shavell (1974). The theory of search. *Journal of Economic Theory*.
- Kroft, K., F. Lange, and M. J. Notowidigdo (2013). Duration dependence and labor market conditions: Evidence from a field experiment. *The Quarterly Journal of Economics* 1123, 1167.
- Lahey, J. N. (2008). Age, women, and hiring an experimental study. *Journal of Human Resources* 43(1), 30–56.
- Lazear, E. P., K. L. Shaw, and C. T. Stanton (2017). Who gets hired? the importance of competition among applicants. *Working Paper*.
- Leibbrandt, A. and J. A. List (2014). Do women avoid salary negotiations? evidence from a large-scale natural field experiment. *Management Science* 61(9), 2016–2024.
- List, J. A., A. M. Shaikh, and Y. Xu (2016). Multiple hypothesis testing in experimental economics.
- Maniadis, Z., F. Tufano, and J. A. List (2014). One swallow doesn't make a summer: New evidence on anchoring effects. *The American Economic Review* 104(1), 277–290.
- Martin, R. and J. Randal (2008). How is donation behaviour affected by the donations of others? *Journal of Economic Behavior & Organization* 67(1), 228–238.
- Mas, A. and A. Pallais (2016). Valuing alternative work arrangements. *National Bureau of Economic Research*.
- Mobius, M. and T. Rosenblat (2014). Social learning in economics. *Annual Review of Economics* 6(1), 827–847.
- Moore, E. and C. Eckel (2003). Measuring ambiguity aversion. *Working Paper*.
- Mortensen, D. T. (1970). Job search, the duration of unemployment, and the phillips curve. *American Economic Review* 60(5), 847–862.

- Nachman, D. (1972). On risk aversion and optimal stopping. *Working Paper*.
- Neumark, D., R. J. Bank, and K. D. Van Nort (1996). Sex discrimination in restaurant hiring: An audit study. *The Quarterly Journal of Economics* 101(3), 915–941.
- Niederle, M. and L. Vesterlund (2007). Do women shy away from competition? do men compete too much? *The Quarterly Journal of Economics*, 1067–1101.
- Oreopoulos, P. (2011). Why do skilled immigrants struggle in the labor market? a field experiment with thirteen thousand resumes. *American Economic Journal: Economic Policy*, 148–171.
- Petit, P. (2007). The effects of age and family constraints on gender hiring discrimination: A field experiment in the french financial sector. *Labour Economics* 14(3), 371–391.
- Riach, P. A. and J. Rich (2006). An experimental investigation of sexual discrimination in hiring in the english labor market. *Advances in Economic Analysis & Policy* 5(2).
- Roth, A. E. (2008). What have we learned from market design?\*. *Economic Journal* 118(527), 285–310.
- Salganik, M. J., P. S. Dodds, and D. J. Watts (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *science* 311(5762), 854–856.
- Samek, A. S. (2015). A university-wide field experiment on gender differences in job entry decisions. *Working Paper*.
- Schubert, R., M. Brown, M. Gysler, and H. W. Brachinger (2000). *Gender specific attitudes towards risk and ambiguity: an experimental investigation*. Institut für Wirtschaftsforschung, Eidgenössische Technische Hochschule.
- Shang, J. and R. Croson (2006). The impact of social comparisons on nonprofit fundraising. *Research in Experimental Economics* 11, 143–156.
- Smith, L. and P. N. Sørensen (2011). Observational learning. *The New Palgrave Dictionary of Economics Online Edition*, 29–52.
- Smith, S., F. Windmeijer, and E. Wright (2015). Peer effects in charitable giving: Evidence from the (running) field. *The Economic Journal* 125(585), 1053–1071.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy* 69(3), 213–225.
- Telser, L. G. (1973). Searching for the lowest price. *American Economic Review* 63(2), 40–49.
- Tucker, C. and J. Zhang (2011). How does popularity information affect choices? a field experiment. *Management Science* 57(5), 828–842.

- Van Ours, J. and G. Ridder (1992). Vacancies and the recruitment of new employees. *Journal of Labor Economics*, 138–155.
- Vandegrift, D. and A. Yavas (2009). Men, women, and competition: An experimental test of behavior. *Journal of Economic Behavior & Organization* 72(1), 554–570.
- Weber, A. and C. Zulehner (2010). Female hires and the success of start-up firms. *American Economic Review*, 358–361.
- Weber, A. and C. Zulehner (2014). Competition and gender prejudice: are discriminatory employers doomed to fail? *Journal of the European Economic Association* 12(2), 492–521.
- Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica* 47(3), 641–654.
- Yechiam, E., M. Druyvan, and E. Ert (2008). Observing others' behavior and risk taking in decisions from experience. *Judgment and Decision Making* 3(7), 493–500.