

Gee, Laura Katherine

Working Paper

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

IZA Discussion Papers, No. 10372

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Gee, Laura Katherine (2016) : The More You Know: Information Effects on Job Application Rates in a Large Field Experiment, IZA Discussion Papers, No. 10372, Institute for the Study of Labor (IZA), Bonn

This Version is available at:

<http://hdl.handle.net/10419/149231>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 10372

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

Laura K. Gee

November 2016

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

Laura K. Gee

Tufts University and IZA

Discussion Paper No. 10372
November 2016

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

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

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 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.

JEL Classification: C93, D01, D83, J21, J22

Keywords: field experiment, labor search, social information, big data, gender, risk aversion, ambiguity aversion, uncertainty, herding, competition

Corresponding author:

Tufts University
Department of Economics
Tufts University
305 Braker Hall
8 Upper Campus Rd.
Medford, MA 02155-6722
USA
E-mail: laura.gee@tufts.edu

* A huge thanks to Joe Adler and Jim Baer at LinkedIn. Additionally, this work would not have been possible without the assistance of Mathieu Bastian, Joi Deaser, Vitaly Gordon, Abishek Gupta, Kuo-Ning Huang, Sachit Kamat, Navneet Kapur, Shreya Oswal, and others at LinkedIn who shared their knowledge and expertise. Thank you to Kevin Morsony, product counsel at LinkedIn, for reviewing this paper. This research was approved by the Tufts IRB (1308011 and 1405012). Thanks to Darcy Covert for research assistance. Thanks to Yan Chen, Catherine Eckel, Dorothea Kubler, John List, Ioana Marinescu, Gratiana Pol, Lones Smith, Chris Stanton, Sonny Tambe, 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, 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.

1 Introduction

Labor market matching is important. Both firms and job seekers expend a large amount of time and money attempting to match. Understanding how job seekers decide to apply is useful. This paper presents the results from a large (over two 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. However, the treatment group is shown one extra piece of information: the current number of current applicants. Specifically for job postings that are internal to the LinkedIn website jobs seeker see the number of finished applications. While for job postings that route to an external website job seekers see the number of started applications, since LinkedIn would not know if a person had finished an application.

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, age, and experience, although those differences are not always statistically significant. So, 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). Also, firms and policy makers want to increase workforce diversity.¹ Yet, a firm cannot hire more female or black engineers, for example, if there is a lack

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

of female or black applicants. So, knowing how to encourage a wide range of individuals to apply could increase both the quantity and the diversity of the applicant pool.²

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.³ 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 fictitious CVs to actual job opening.⁴ 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.⁵ Samek (2015) advertises a temporary administrative position she created and finds that a non-profit framing increases applications, and that 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 adding information about the actions of others to the offer letter for a position as a Teach for America (TFA) teacher (specifically stating that 84% of applicants accept their TFA offer) significantly increases an applicant’s likelihood of accepting the offer as

²It is also possible that increasing the number of applicants could lead to too much congestion Roth (2008). I explore that topic in the Robustness Checks section.

³See 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).

⁴See 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).

⁵In a related paper Leibbrandt and List (2014) find that women are less likely to negotiate their wage unless explicitly told the wage is negotiable.

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 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 is 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.⁶

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. Such information can be a positive cue on the one hand (e.g., the presence of several other applicants may signal a good job) or a negative one on the other hand (e.g., the presence of many other applicants may signal high competition). The reverse should hold true if a scarcity of applicants for a position is revealed. Given the contradictory psychological effects of this particular type of social 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

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

rates.

LinkedIn ran this particular experiment as part of their normal business practices. While the experiment wasn't designed to reveal the underlying mechanisms for why more people applied, 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 larger treatment effect for higher numbers of applicants shown (e.g., a job seeker sees 100 applicants and believes this is to be a positive signal of job quality, and is hence more likely to apply).⁷ Conversely, if I observe a smaller treatment effect for higher numbers of applicants shown, that would be consistent with a competition aversion mechanism (e.g., a job seeker sees 100 applicants and does not want to compete with so many applicants).⁸ Since female and older individuals tend to be more competition averse, finding a smaller treatment effect for higher numbers in these sub-populations would be further evidence of a competition aversion mechanism.⁹ I do not observe heterogeneous treatment effects that are consistent with either a herding or competition aversion 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 a more-information mechanism is the main driver, then one would expect experienced job seekers to be less affected by the treatment. Additionally, if I do observe a positive treatment effect, this effect would not vary by the number of applicants shown. This mechanism could also be called an ambiguity or risk aversion mechanism.¹⁰ Since women and older

⁷See Bougheas et al. (2013); Smith and Sørensen (2011); Yechiam et al. (2008) and Anderson and Holt (1997).

⁸It is also possible there is a non-linear relationship, but I find no evidence of that as described in the Appendix Figure 5. Additionally it is possible that the two effects are washing out which I discuss in the Robustness Checks section.

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

¹⁰Note 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.

people are more ambiguity and risk averse, finding a larger treatment effect in these sub-populations would be further evidence of a more-information mechanism.¹¹ I do find some evidence that less experienced, female, and younger individuals have larger treatment effects. This finding provides some support for a more-information mechanism driving the results.

This paper presents the results from an experiment studying 2.3 million job seekers looking at over 100,000 job postings across a wide range of domains. 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 countries.¹² 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.¹³ These industries have not traditionally had a very diverse workforce.¹⁴ A first step toward a more diverse workforce is a more diverse applicant pool. Because the industries represented on LinkedIn often

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

¹²See <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.

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

¹⁴For example only 32.5% of US professionals in STEM related fields (Science Technology Engineering and Mathematics) are female. See <http://dpeaffcio.org/programs-publications/issue-fact-sheets/women-in-stem/>.

struggle with diversity, LinkedIn represents an ideal setting for exploring how job seekers decide to apply to job postings.

To use the job postings on LinkedIn, a member can either use the search bar or access the Jobs landing page (see Appendix Figure 2 and Appendix Figure 3), where she can see a number of job postings that are pre-selected by LinkedIn based on information the member has listed on her profile such as education, industry, and previous employment. Then the member can click on one of the postings listed. After clicking on a posting, the member sees a full page description of the posting.

LinkedIn has two types of job postings (see Appendix Figure 4): 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.¹⁵ In the case of exterior postings I can observe only if a user starts an application. During the experiment 61% 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.¹⁶ 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 1.¹⁷ The

¹⁵I 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.

¹⁶To maintain security, LinkedIn has a policy that all analysis is done at LinkedIn using secured computers and networks. This means that to run any new analysis I need to go back to LinkedIn.

¹⁷For 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

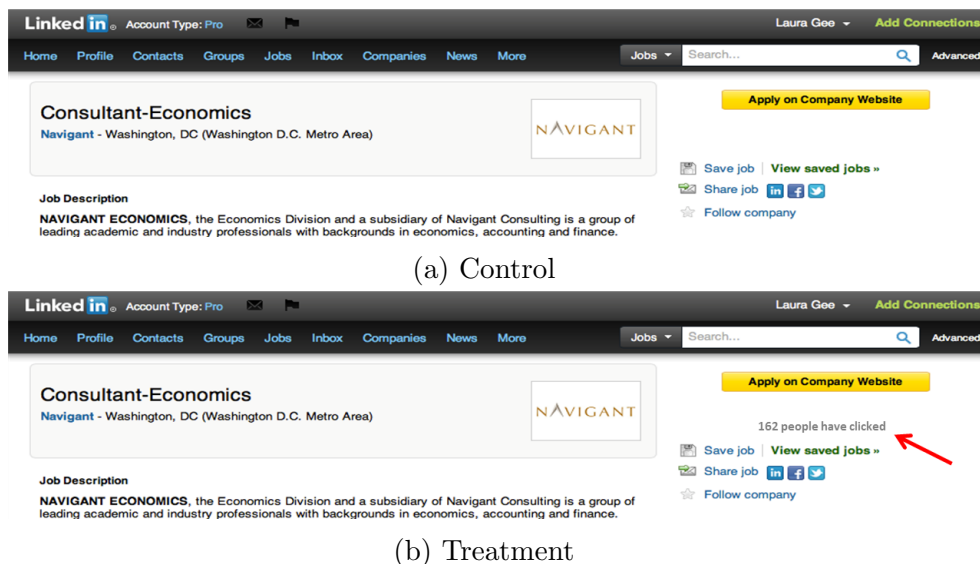


Figure 1: Job Posting as Seen in Control and Treatment

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.

content of the job postings did not differ between the control and treatment conditions, and in fact 95% of viewed job postings were the same across the two conditions. As a robustness check I show that the main results hold if I only include jobs that were seen in both the control and the treatment group (Appendix Table 7).

LinkedIn chose to randomly assign one-fourth of the active job seekers to the treatment group and the remaining three-fourths to the control group.¹⁸ 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 applicants for that job. Additionally, because the information is exogenously assigned, I can rule out the possibility that those who seek out job posting the button read “Apply Now”, and those in the treatment group saw the number of finished applications.

¹⁸I exclude members who were included in a previous pilot study that took place in the two weeks before the main experiment.

more information are already more likely to apply for a position.

Overall, the sample includes about 2.3 million registered members from 235 countries. There are about 580,000 job seekers in the treatment and 1.7 million job seekers in the control 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.¹⁹

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).²⁰ 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 2013.²¹

¹⁹The 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.

²⁰Members do not provide gender, but it is imputed from their country and name by LinkedIn (e.g., Laura in the US 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.

²¹A “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 co-worker, and then that

The subjects are very well educated, with 2% listing an Associates degree, 52% listing a Bachelors degree, and 46% listing a post-Bachelors degree as their highest education level attained. Overall, subjects in the control and treatment groups are similar on observable 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' location.

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
US	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.

3 Analysis

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

3.1 Main Results

The outcome variables are (1) whether a person starts an application and (2) whether a person finishes an application. While I can observe starting an application for both exterior and interior job postings, I can observe finishing an application only for interior job postings. I will conduct the analysis over two groups: those who saw an exterior posting, and those who saw an interior posting.

co-worker 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.

Whether a person takes a job, as well as her tenure at the position, would also be an interesting outcome. Since less than 3,000 job seekers can be matched to a job at the firm they applied to, there are too few observations to draw any conclusions.

The average number of job postings viewed by both the control and the treatment group was 3.8 postings ($t = 0.91$) over the 16 days of the experiment. So the treatment does not seem to have a measurable effect on search intensity. However, to mitigate path dependence for those in the treatment group, I restrict the dataset to the first job posting a person views, which leaves a total of 2,326,207 members for analysis.²² Those are split roughly evenly between first viewing an external job posting (1,134,109) and first viewing an internal job posting (1,192,098).

My simple model is:

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

where $A_{i,j}$ takes the value 100 if member i starts/finishes application to job j and zero otherwise. If the member is in the treatment group that sees the number of previously started applications, then $T_i = 1$ and zero otherwise. By having $A_{i,j}$ take the value 100 when a person applies the coefficient β can be interpreted as the percentage point (without having to multiply by 100) difference in likelihood of application from being in the treatment group.

One may worry 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 firm. I can however include a job posting fixed effect P_j in the analysis to control for all time-invariant attributes of the job. Since 95% of jobs were seen in both the control and 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

²²For 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 52 for more details.

the treatment group.

I do, however, control for variables that are likely used by the background search and suggestion 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 $NumCurrApply100_{i,j}$ (even when this is not seen in the control). $NumCurrApply100_{i,j}$ is the true number of current applicants divided by 100.²³ 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.²⁴ This leaves me with the following preferred model:

$$A_{i,j} = \beta T_i + P_j + \gamma D_{i,j} + \alpha NumCurrApply100_{i,j} + \epsilon_{i,j}. \quad (2)$$

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.²⁵

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. A back-of-the-envelope calculation suggests that the treatment would result in an extra 1,500 started and 250 finished applications per day.²⁶ It could also

²³I chose to divide the number of applicants by 100 because the coefficients on the non-scaled variable are extremely close to zero.

²⁴A logit model yields similar results and is presented in Appendix Table 7.

²⁵See 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 only members who started searching during the experiment (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).

²⁶First, 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

Table 2: Likelihood of Starting/Finishing an Application

Simple: Without Controls or Fixed Effects			
	1	2	3
	Start Ext	Start Int	Finish Int
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%
Preferred: With Controls and Job Fixed Effects			
	1	2	3
	Start Ext	Start Int	Finish Int
Treatment β	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)
NumCurrApply100	-0.716*** (0.124)	-1.221*** (0.112)	-1.075*** (0.070)
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. The model includes job posting fixed effects, days posted during experiment, and the current number of applicants at time of viewing divided by 100 (NumCurrApply100). Legend: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$			

change the final pool of applicants, which I will explore later.

Although the treatment effect is not very sensitive to the inclusion of the control variables, the coefficients on the control variables are statistically significant. For example, increasing the number of applicants by 100 decreases the likelihood of application by 0.716-1.221 percentage points. This is a relatively small decrease given that the median number of applicants is 21, and the 90th percentile is 133.

3.2 Candidate Mechanisms

Understanding the mechanisms behind the increased application rate could allow firms to target those who they want the additional applications from. 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.²⁷ 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. While both can conceivably take place, from a policy perspective the overall effect is most important. In the Robustness Checks 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.

of about 1,500 started applications. A 0.090 percentage point increase would be an extra 250 in finished applications.

²⁷Here I use the term competition averse to mean someone who all else equal is less likely to apply to a job posting with more applicants. Someone would be more competition averse, the more they decreased their likelihood of application in response to a single extra applicant.

The number of applicants seen 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 path dependence I begin by using only the first posting seen. Using this approach, the interaction of the treatment with the number of applicants ($Treatment * NumCurrApply100$) is not consistent in sign, and is statistically insignificant (Table 3, Panel A).²⁸ Non-linear models show a similarly noisy relationship (Appendix Figure 5).²⁹

Next, I expand the dataset beyond the first job posting viewed, so that the same 2.3 million members now view over 8 million job postings. I find that the coefficients on $Treatment * NumCurrApply100$ 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 -0.058 percentage points for every 100 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).³⁰

The noise seen in the data could result from job seekers' inability to interpret the absolute number seen.³¹ For example, Eleanor might think that 25 applicants is a high number, while Dan may perceive the same number to be low. To benchmark a number as low or high I use the difference in the number of applicants between the currently and the previously viewed posting ($DiffNumCurrApply100 = NumCurrApply100_t - NumCurrApply100_{t-1}$).

²⁸The coefficients on $Treatment * NumCurrApply100$ are 0.028 to 0.019. Even if these point estimates were more exact they imply a weak relationship, because 100 applicant increases are uncommon ($NumCurrApply100$ median = .26, mean = .71).

²⁹A quadratic model also yields noisy estimates. Results available from author upon request.

³⁰The relationship is weak because a 100 applicant increases is uncommon using all 8 million views ($NumCurrApply100$ median = 0.26, mean = .71, coincidentally the same as for first view sub-sample).

³¹To 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 for details.

Table 3: Heterogeneous Treatment Effects by Number Shown

	1	2	3
	Start Ext	Start Int	Finish Int
A. First View Only			
Treatment	0.366*** (0.067)	0.211** (0.073)	0.072+ (0.039)
Treatment*NumCurrApply100	-0.028 (0.043)	-0.003 (0.052)	0.019 (0.032)
NumCurrApply100	-0.709*** (0.125)	-1.220*** (0.113)	-1.080*** (0.071)
Adj R2	0.049	0.052	0.013
N (members)	1,134,109	1,192,098	1,192,098
B. All Views			
Treatment	0.355*** (0.049)	0.205*** (0.051)	0.050+ (0.026)
Treatment*NumCurrApply100	-0.058+ (0.030)	-0.015 (0.034)	0.019 (0.020)
NumCurrApply100	-0.501*** (0.091)	-0.724*** (0.078)	-0.887*** (0.046)
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
C. Current Num - Prev Num			
Treatment	0.281*** (0.062)	0.142* (0.063)	0.027 (0.034)
Treatment*DiffNumCurrApply100	-0.031 (0.032)	-0.008 (0.030)	0.003 (0.016)
DiffNumCurrApply100	-0.249*** (0.022)	-0.273*** (0.018)	-0.091*** (0.010)
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 and the number of applicants at the time of viewing divided by 100 (NumCurrApply100). Legend: + $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$ Panel B includes 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 C includes data for all but the first job posting viewed by the 1,248,289 members with 2+ views, $DiffNumCurrApply100 = NumCurrApply100_t - NumCurrApply100_{t-1}$, observations are weighted so that each member's weights sum to 1, and standard errors are clustered at the member level.

This requires restricting the data to the 1.2 million members who view at least two job postings.³² The coefficients for the interaction between *DiffNumCurrApply100* and the treatment are neither consistent in sign nor statistically significant (Table 3, Panel C).

Since previous work finds that female and older individuals are more competition averse, competition aversion may be more pronounced in these sub-samples.³³ The coefficient for the interaction between *Treatment* and *NumCurrApply100* is neither consistently negative nor statistically significant for female job seekers.³⁴ Also, the coefficient for the interaction is neither consistently negative nor statistically significant for job seekers over the median age of 33.³⁵

The original experiment was not designed to test for competition aversion, though there are some heterogeneous treatment effects predictions that would be consistent with competition aversion. However, there is no consistent pattern of a decline in the treatment effect for more applicants shown, as summarized in Table 4. 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. If more-information is the main driver, then the treatment

³²This results in losing about half the sample. This sub-sample is balanced on observables across the control and treatment. The sub-sample is similar to the full sample with the exception of having 20 more LinkedIn connections.

³³See 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. Moreover, 94% of jobs are seen by both older and younger job seekers, where older is defined as over the median age of 33.

³⁴See Appendix Table 6, Panel B. The coefficient for women (*Treatment * NumCurrApply100*) is -0.063 to 0.066 but never statistically significant. The coefficient for men is -0.003 to 0.023 but never statistically significant.

³⁵See Appendix Table 5, Panel B. The coefficient for those over 33 on *Treatment * NumCurrApply100_t* is -0.002 to 0.094 , but never statistically significant, and neither is the triple interaction *Age * Treatment * NumCurrApply100_t*.

Table 4: Candidate Mechanisms

Competition Aversion		More Info	
$Treat * NumPrvApply100 < 0$	No ¹	$Treat * NumPrvApply100 \sim 0$	Yes ¹
$Treat_{Female} * NumPrvApply100 < 0$	No	$Treat_{Female} > Treat_{Male}$	Maybe ²
$Treat_{Older} * NumPrvApply100 < 0$	No	$Treat_{Older} > Treat_{Younger}$	No ³
		$Treat_{Inexperienced} > Treat_{Experienced}$	Maybe ⁴

Notes: 1. I use nine models to test for the coefficient on $Treat * NumPrvApply100$ (Table 3 shows three outcome variables over three sub-sets of the data, for a total of nine models). It is only statistically significant at the 10% level for starting an exterior application if I use all the postings viewed. 2. For two of three outcomes the treatment effect is zero for men, but positive and statistically significant for women. 3. For two of three outcomes the treatment effect for finishing an interior application does not differ by age for finishing an interior application $Treat_{Older} < Treat_{Younger}$. 4. For two of three outcomes the coefficient on $Treat * NumPrvApply100$ is statistically significantly larger for inexperienced job seekers.

effect should be more pronounced for the risk/ambiguity averse (e.g., women and older job seekers) and for inexperienced job seekers. Additionally, unlike the results described in the previous section, the specific number of applicants shown would not moderate the magnitude of the treatment effect.³⁶ The tests for these heterogeneous treatment effects are summarized in Table 4.

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 and older people are generally more ambiguity or risk averse, finding a larger treatment effect in these sub-populations would be evidence of a more-information mechanism.³⁷ The treatment effect is directionally larger for women than men; however, the difference is only statistically significant for finishing an interior application (see Appendix Table 6).³⁸ Also, it is important to note that the treatment effect for men is only statistically greater than zero for one of the three outcomes variables.³⁹ The treatment effect for

³⁶For 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.

³⁷See Garratt et al. (2013); Bertrand (2011); Croson and Gneezy (2009); Eckel and Grossman (2008); Moore and Eckel (2003); Schubert et al. (2000).

³⁸The 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).

³⁹In Panel A of Appendix Table 6 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 Appendix Table 6 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.

older job seekers is not larger than for younger ones.⁴⁰ This provides mixed evidence in support of a more-information mechanism.

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.⁴¹ The treatment effect is larger for inexperienced job seekers for two of the three outcomes.⁴²

This experiment was not designed to test for underlying mechanisms, but the lack of a pattern by number shown, and some evidence of higher treatment effects for women and inexperienced job seekers, are consistent with a more-information mechanism as summarized in Table 4.

4 Robustness Checks

Showing the number of applicants increases the likelihood of application and might be used to increase minority applications, but there are some robustness checks that I discuss in this section.

⁴⁰In Panel A and B of Appendix Table 5 the coefficients on *Treatment*over33* are not statistically significant for starting an exterior or interior job application. In fact, the only statistically significant interaction between the treatment and age indicates that those over the median age of 33 have a 0.166-0.198 percentage point decline (rather than increase) in the treatment effect on likelihood of finishing an interior application (see Appendix Table 5 Panel B).

⁴¹Age and membership years have a correlation coefficient of 0.31. Including age in the models in Appendix Table 5 Panel C 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).

⁴²An extra year of LinkedIn membership decreases the treatment effect on starting an exterior application by 0.063 percentage points, and finishing an interior application by 0.051 percentage points (though the interaction coefficient is not statistically significant for starting an interior application; see Appendix Table 5 Panel C). 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 Appendix Table 5 Panel C). 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, $0.108/0.043 = 2.5$.

4.1 Are women more likely to apply to masculine jobs?

On LinkedIn men are more likely to start an application than women.⁴³ 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.⁴⁴ Note, 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.⁴⁵ Using a model without job fixed effects, I interact the *MasculineJob* variable with the treatment for men and women. I find that showing the number of applicants increases the likelihood of a female job seeker starting or finishing a masculine

⁴³If 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.

⁴⁴To 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).

⁴⁵Restricting 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 sub-samples. The sub-samples are balanced on the observable characteristics across the control and the treatment, so these results should be internally valid. The members in these sub-samples 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.

job application.⁴⁶ 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, however, have data on the final number of applications started or finished for each job by the end of the experimental period.⁴⁷ 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.⁴⁸ 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.

⁴⁶See 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 (vs. 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 sub-sample.

⁴⁷This is different from the number shown, since that is a running tally of applicants.

⁴⁸Recall 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 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 (98,070 members). I then compute the correlation between starting an application and the number shown.⁴⁹ I find that the mean correlation is 0.043, quite close to zero. However, as shown in Appendix Figure 10, almost half of the members have a positive correlation (e.g. herding) and the other half a negative correlation (e.g. competition aversion).⁵⁰ 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.⁵¹

4.4 Is this encouraging new applicants?

Since each woman can ultimately take one only job, increasing the number of jobs she applies to may not actually increase workforce diversity. I therefore

⁴⁹I 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.

⁵⁰The 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.

⁵¹Note 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.

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. 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.⁵² Since many job searches last longer than 16 days, 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.

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

⁵²For 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 ($t = 0.810$).

may be a welfare gain from decreased congestion, but also a decrease in the number of minority applicants. I find that women (especially those who apply for male-dominated jobs) and inexperienced job seekers 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.

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. Because this change is quite simple at the time of writing LinkedIn allowed firms to choose whether they would like to show this information on their job postings.

Such behavioral “nudges” have become popular tools for policy makers and firms to influence short-term behavior. Recently Coffman et al. (2014) have shown that a 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. Forthcoming. the short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation.. *American Economic Review*.
- Anderson, L. R. and C. A. Holt (1997). Information cascades in the laboratory. *The American Economic Review* 87(5), pp. 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. *The 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.
- 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 and 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. *The Quarterly Journal of Economics* 75(4), pp. 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*.
- 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. *The 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).
- 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.
- 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.
- 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.
- 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. *The American Economic Review* 60(5), pp. 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?*. *The Economic Journal* 118(527), 285–310.
- 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), pp. 213–225.
- Telser, L. G. (1973). Searching for the lowest price. *The American Economic Review* 63(2), pp. 40–49.
- 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. *The 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), pp. 641–654.
- Yan Chen, F. Maxwell Harper, J. K. S. X. L. (2010). Social comparisons and contributions to online communities: A field experiment on movielens. *The American Economic Review* 100(4), 1358–1398.
- Yechiam, E., M. Druryan, and E. Ert (2008). Observing others’ behavior and risk taking in decisions from experience. *Judgment and Decision Making* 3(7), 493–500.

6 Appendix (For Online Publication)

6.1 Setting

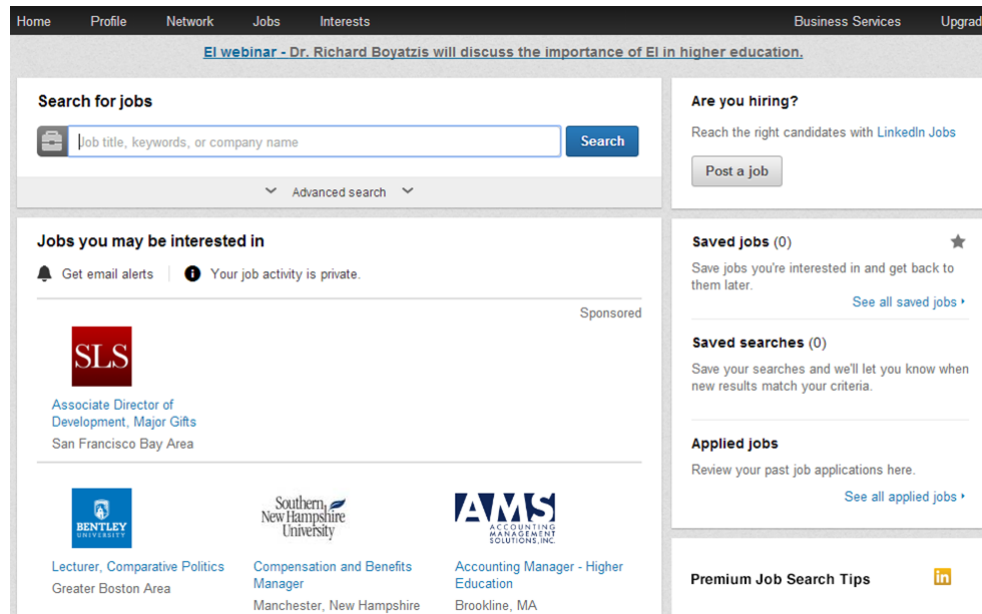


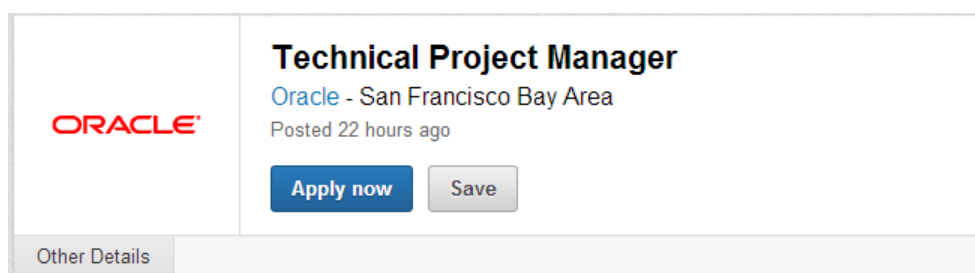
Figure 2: Jobs Landing Page

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

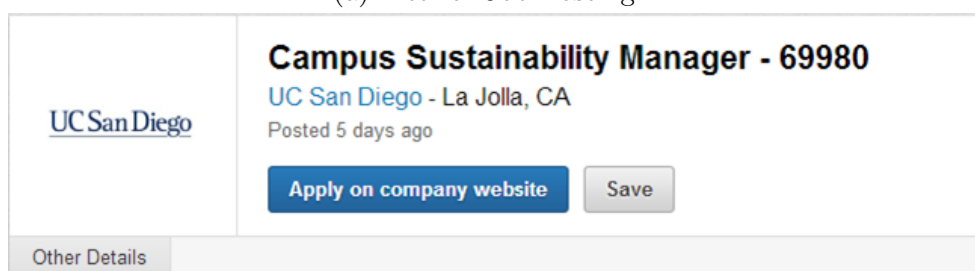
The screenshot shows the LinkedIn job search interface. At the top, the LinkedIn logo and 'PREMIUM' status are visible. The search bar contains the word 'Economics'. Below the search bar, the results are sorted by 'Relevance' and show 283 results. On the left side, there is a sidebar with filters for 'All', 'Jobs', and 'More...'. Below this, there are sections for 'Keywords' (with 'Economics' entered), 'Company', 'Title', 'Location' (set to 'Located in or near:'), 'Country' (set to 'United States'), 'Postal Code' (set to '02144'), and 'Within' (set to '50 mi (80 km)'). A 'Search' button and a 'Reset' link are at the bottom of the sidebar. The main content area displays five job listings, each with a company logo, job title, company name, location, date, and a 'Save Job' button. The job listings are: 1. Manager, Economics & Regulation Job at KPMG, US -Massachusetts- Boston, Feb 25, 2014. 2. Healthcare Economics Analyst at Smith & Nephew, Greater Boston Area, Feb 20, 2014. 3. Research Scientist, Health Economics Modeling & Simulation (consulting) at Evidera, Greater Boston Area, Feb 20, 2014. 4. Sr Research Associate Health Economics at Mapi Group, Greater Boston Area, Feb 17, 2014. 5. Associate Director, Global Health Economics & Outcomes Research (HEOR) at Vertex Pharmaceuticals, Boston, MA, Feb 18, 2014. On the right side, there are several advertisements, including 'Master Applied Psychology', 'Online MS in Applied Psychology from USC', 'Top-Ranked MBA in Boston', and 'Online Teaching Job Leads'.

Figure 3: Job Search Landing Page

Note: This figure shows the results from a search for the term “Economics.”



(a) Interior Job Posting



(b) Exterior Job Posting

Figure 4: Types of Job Postings on LinkedIn

Note: This figure shows an example of the two types of job postings on LinkedIn. Panel (a) shows an interior posting, which means that LinkedIn collects applications for a third party (Oracle). For these, I can observe if a person both begins and finishes an application. Panel (b) shows an exterior posting, which means that a person is directed to an external website to begin an application and thus I can only observe if someone begins the application. These screenshots were taken in February 2013, which is why they differ very slightly from the formatting seen in the example of the treatment versus control screenshots in Figure 1.

6.2 Heterogeneous Treatment Effects

I created a model with quantile bins for the number of current applicants. The number of applicants was broken into five quantiles $QNumCurrApply100_{i,d,j}$ and then interacted with the treatment as in the equation below ⁵³


$$A_{i,d,j} = \beta T_i + \lambda T_i * QNumCurrApply100_{i,d,j} + \alpha QNumCurrApply100_{i,d,j} + P_j + D_d + \epsilon_{i,d,j}.$$

Figure 5 graphically represents the results from this model. On the vertical axis of Figure 5 is the percentage point difference in the likelihood of applying between the treatment and the control groups. On the horizontal axis is the 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 5. On the other hand, if herding is the dominant effect, one would expect to see an upward sloping trend in the panels of Figure 5.

Figure 5 uses only data from the first posting viewed to eliminate path dependence. However, 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 current job posting and the number seen on the 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.


⁵³I 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.



(a) Exterior: Starting Application



(b) Interior: Starting Application



(c) Interior: Finishing Application

Figure 5: Plots of Coefficients on Treatment Dummy Variable by Number Applicants Shown

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

Table 5: Heterogeneous Treatment Effects: Age and Experience

	1	2	3
	Start Ext	Start Int	Finish Int
A. Age			
Treatment (Under33)	0.206+	0.103	0.198**
	(0.108)	(0.107)	(0.061)
Treatment*Over33	0.202	0.169	-0.172*
	(0.149)	(0.149)	(0.084)
Over33	-0.033	-0.099	0.004
	(0.079)	(0.081)	(0.046)
NumCurrApply100	-0.871***	-1.246***	-1.116***
	(0.143)	(0.124)	(0.081)
Years member	-0.299***	-0.453***	0.076***
	(0.017)	(0.017)	(0.010)
Adj R2	0.051	0.055	0.014
N	897,696	939,620	939,620
B. Age + Number Seen			
Treatment (Under33)	0.264*	0.161	0.184**
	(0.113)	(0.119)	(0.068)
Treatment*Over33	0.142	0.093	-0.166+
	(0.154)	(0.164)	(0.094)
Over33	-0.032	-0.097	0.003
	(0.079)	(0.081)	(0.046)
Treatment*Over33*NumCurrApply100	0.094	0.080	-0.002
	(0.089)	(0.109)	(0.070)
Treatment*NumCurrApply100 (Under33)	-0.090	-0.057	0.013
	(0.071)	(0.068)	(0.044)
NumCurrApply100	-0.860***	-1.239***	-1.119***
	(0.144)	(0.125)	(0.082)
Years member	-0.299***	-0.453***	0.076***
	(0.017)	(0.017)	(0.010)
Adj R2	0.051	0.055	0.014
N	897,696	939,620	939,620
C. Experience (Years LinkedIn Member)			
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)
NumCurrApply100	-0.722***	-1.233***	-1.071***
	(0.124)	(0.112)	(0.070)
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 and the number of applicants at time of viewing divided by 100 (NumCurrApply100). Legend: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 6: Heterogeneous Treatment Effects: Gender

	1	2	3
	Start Ext	Start Int	Finish Int
A. Gender			
Treatment (Female)	0.383*** (0.111)	0.302** (0.112)	0.212*** (0.058)
Treatment*Male	-0.033 (0.141)	-0.174 (0.141)	-0.170* (0.074)
Male	1.102*** (0.073)	1.426*** (0.074)	0.498*** (0.038)
NumCurrApply100	-0.704*** (0.131)	-1.171*** (0.119)	-1.114*** (0.074)
Adj R2	0.049	0.052	0.013
N	1,020,017	1,072,330	1,072,330
B. Gender + Number Seen			
Treatment (Female)	0.421*** (0.114)	0.242+ (0.125)	0.162* (0.066)
Treatment*NumCurrApply100 (Female)	-0.063 (0.068)	0.066 (0.084)	0.056 (0.051)
Treatment*Male	-0.068 (0.144)	-0.122 (0.155)	-0.138+ (0.083)
Treatment*Male*NumCurrApply100	0.060 (0.081)	-0.057 (0.101)	-0.034 (0.062)
Male	1.102*** (0.073)	1.425*** (0.074)	0.498*** (0.038)
NumCurrApply100	-0.696*** (0.132)	-1.178*** (0.119)	-1.122*** (0.075)
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 and the number of applicants at time of viewing divided by 100 (NumCurrApply100). Legend: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 7: Main Results Robustness Checks

	1	2	3
	Start Ext	Start Int	Finish Int
A. Without Controls or Fixed Effects			
Treatment	0.355*** (0.065)	0.225*** (0.065)	0.094** (0.034)
Adj R2	0.000	0.000	0.000
N	1,134,109	1,192,098	1,192,098
B. Preferred Main Results			
Treatment	0.349*** (0.065)	0.208** (0.065)	0.089** (0.034)
Adj R2	0.049	0.052	0.013
N	1,134,109	1,192,098	1,192,098
C. Logit with Fixed Effects			
Treatment	0.041*** (0.008)	0.023** (0.007)	0.035** (0.014)
Adj R2			
N	944,489	1,049,361	717,813
D. Clustered Errors at Job Level			
Treatment	0.355*** (0.064)	0.225*** (0.065)	0.094** (0.033)
Adj R2	0.000	0.000	0.000
N	1,134,109	1,192,098	1,192,098
E. US Only			
Treatment	0.325*** (0.096)	0.143 (0.106)	0.078 (0.049)
Adj R2	0.039	0.030	-0.003
N	527,193	446,921	446,921
F. Non-US Only			
Treatment	0.415*** (0.091)	0.231** (0.084)	0.097* (0.047)
Adj R2	0.061	0.064	0.017
N	606,916	745,177	745,177
G. Jobs in Control and Treatment			
Treatment	0.349*** (0.065)	0.208** (0.065)	0.089** (0.034)
Adj R2	0.047	0.052	0.013
N	1,130,546	1,190,953	1,190,953
H. Start Search During Experiment			
Treatment	0.365*** (0.069)	0.227** (0.069)	0.086* (0.036)
Adj R2	0.049	0.053	0.012
N	1,007,804	1,056,338	1,056,338

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. In Panel C I had to omit one of the job postings because the results would not converge while it was included. Legend: + $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

Table 8: Likelihood Apply to Masculine Job

	1	2	3
	Start Ext	Start Int	Finish Int
Treatment	0.938*** (0.210)	0.381* (0.193)	0.486* (0.197)
Treatment*MasculineJob _S	10.759*** (0.402)	8.412*** (0.377)	
Treatment*Male*MasculineJob _S	-12.375*** (0.490)	-9.587*** (0.454)	
Male*MasculineJob _S	14.182*** (0.253)	12.385*** (0.233)	
Treatment*Male	3.841*** (0.289)	2.814*** (0.263)	1.584*** (0.269)
MasculineJob _S	-11.357*** (0.209)	-10.844*** (0.194)	
Male	-3.425*** (0.146)	-2.440*** (0.133)	-1.248*** (0.136)
Treatment*MasculineJob _F			4.234*** (0.346)
Treatment*Male*MasculineJob _F			-4.778*** (0.425)
Male*MasculineJob _F			5.784*** (0.216)
MasculineJob _F			-4.781*** (0.176)
N	481,648	614,674	282,986

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 and the number of applicants at time of viewing divided by 100 (NumCurrApply100). MasculineJob_S takes the value 1 if at least 80% of those who started an application for this job posting in the control were male (note members in both the control and the treatment apply to “MasculineJob_S” positions). MasculineJob_F takes the value 1 if at least 80% of those who finished an application for this job in the control were male (note members in both the control and the treatment apply to “MasculineJob_F” positions). Legend: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$

6.3 Survey

In June 2014 I administered an online survey to obtain insight into how job applicants perceive the number of applicants. I used a snowball sampling technique and ended up with $N = 188$ respondents. Of those, 96 had a LinkedIn profile and would consider using LinkedIn to apply for a job. Of this group, 51% said that it takes them over an hour to apply for a job, 36% said it takes 31-60 minutes, and the remaining 12% said it takes 5- 30 minutes.

Survey respondents were shown two almost identical job postings as pictured in Figure 6. The “BLUE” posting has no information and is the same as the control in the field experiment. The “PURPLE” shows the number of applicants; this number was randomly assigned to be 2, 26, 72, 273, or 4,124 for each survey respondent. Survey respondents were asked, “If you were going to apply to either Posting BLUE or Posting PURPLE below, which posting would you prefer to apply to?” Excluding those who could not tell the difference between the BLUE and PURPLE posting, or who thought that the lack of information on the BLUE posting meant 0 applications ($N = 92$), 45% preferred the treatment (PURPLE) to the control (BLUE). For female respondents, 45% preferred the treatment compared to only 44% of the male respondents, but the difference is not statistically significant.

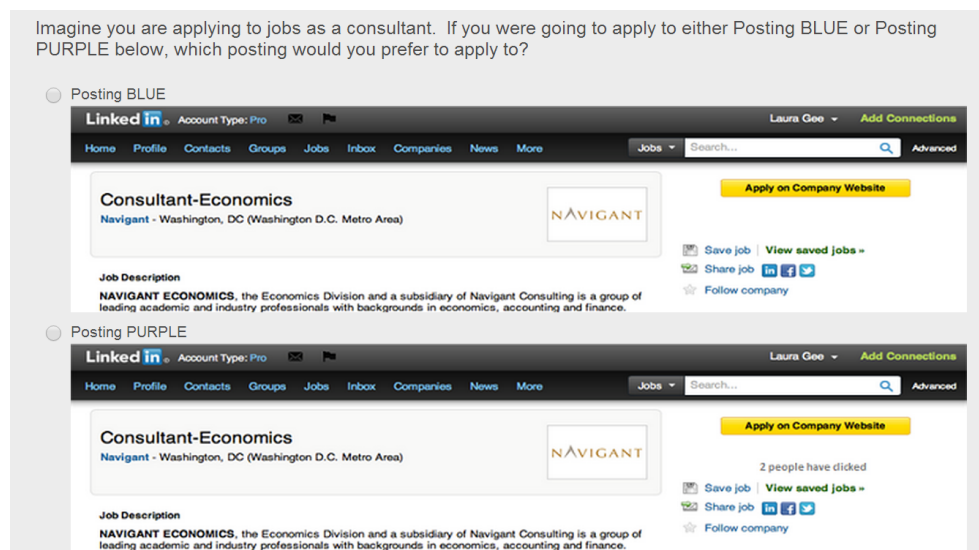


Figure 6: Type of Response by Number Seen

Note: This is the survey question that respondents answered. The number shown was randomly assigned to be either 2, 26, 72, 273 or 4124.

The main purpose of this survey was to determine how people’s beliefs

about applying to a job were affected by viewing the number of applicants. After making the choice between the BLUE and PURPLE posting, respondents were told, “In your own words please explain why you chose the BLUE or PURPLE posting.” The responses fell into four broad categories: (1) those who dislike ambiguity by a preference for more information, (2) those who prefer to avoid congestion/competition, (3) those who herd toward more popular job postings, and (4) other.⁵⁴ A research assistant was able to categorize 74 of the responses into one of the three non-other categories. Interestingly, the respondents seem to interpret the same number (e.g., 2, 26, etc.) differently. For example, some believe that seeing two applicants means there is low congestion/competition, while others think this is high. The fact that people view the same number many ways may explain why there is no pattern of herding/congestion in the field study. This difference in perception can be seen in Figure 7, which shows the proportion of respondents that interpreted the number shown as a sign of congestion/competition, signaling quality, or as extra useful information. Figure 7 shows that every number seen has a variety of interpretations, with the exception being 4,124, which the vast majority interpreted as a signal of congestion/competition.




Figure 7: Jobs Landing Page

Note: This figure shows the proportion of respondents who interpreted the number of applicants as either (1) giving information about competition or congestion, (2) giving information allowing them to herd toward more popular postings, or (3) having more information in general. The proportion is shown for each number of applicants, either 2, 26, 72, 273, or 4,124. For example, for those who saw the number 26, 43% felt this signaled competition, 21% felt it signalled popularity, and 36% liked the additional information.

Here are a few examples of each type of response:

⁵⁴ “Other” includes responses that comment on aesthetic appearance, or are vague.

1. Like Information

- I'd rather have the information to guide both how much time I invest in customizing my resume/ linkedin profile / cover letter and to set my expectations (Female / Shown 4,124)
- I figure more information is better. Given that I know they CAN post the number of clicks, it feels deceptive to hide that information. (Male / Shown 72)

2. Avoid Congestion/Competition

- If over 4000 people have applied to a job posting, I would be unlikely to get the job. Therefore, it isn't worth the time to apply. (Male / Shown 4124)
- When I saw that two people had already clicked on the posting of the purple it made me feel very anxious. I guess that I like to think that I am the only person who is applying and therefore I have a high probability of getting the position. (Female / Shown 2)

3. Herd Toward Popular

- That additional piece of information helps validate my interest by showing me how desirable that position is to other job seekers. (Male / Shown 273)
- The information on the people who have clicked on the job tells me it is a desirable job with a reputable company (Female / Shown 273)

Another goal of the survey was to determine if people felt that competition was declining as the number seen dropped. Survey respondents were asked the following two questions:

- If a job posting that you applied to said 10 people had already begun that application how likely do you believe you would be to get the to the next step in the interview process and eventually get a job offer?
 - Very Unlikely (0-20%)
 - Unlikely (21-40%)
 - Undecided (41-60%)

- Likely (61-80%)
 - Very Likely (81-100%)
- If a job posting that you applied to said 100 people had already begun that application how likely do you believe you would be to get the to the next step in the interview process and eventually get a job offer?
 - Very Unlikely (0-20%)
 - Unlikely (21-40%)
 - Undecided (41-60%)
 - Likely (61-80%)
 - Very Likely (81-100%)

The results from the 137 respondents who answered both questions are represented in Figure 8. The distribution is concentrated around “Very Likely” and “Likely” when only 10 applicants are seen, but shifts toward “Unlikely” and “Very Unlikely” when 100 applicants are seen. This result implies that, as subjects see higher relative numbers, they believe they face greater competition. This supports the use of the relative difference in number seen to test for competition aversion. The shift in the distribution is similar for female and male respondents.




Figure 8: Likelihood of Job Offer

Note: This figure shows the proportion of respondents who said they believed they were likely to go on to the next step of the interview process and eventually get a job offer if they saw 10 vs. 100 applicants.

A final goal of the survey was to determine if people felt that the quality of the position was improving as the number seen increased. Survey respondents were asked the following two questions:

- If a job posting that you applied to said 10 people had already begun that application how likely do you believe you would like that job?
 - Very Unlikely to like job (0-20%)
 - Unlikely to like job (21-40%)
 - Undecided on if will like job (41-60%)
 - Likely to like job (61-80%)
 - Very likely to like job (81-100%)
- If a job posting that you applied to said 100 people had already begun that application how likely do you believe you would like that job?
 - Very Unlikely to like job (0-20%)
 - Unlikely to like job (21-40%)
 - Undecided on if will like job (41-60%)
 - Likely to like job (61-80%)
 - Very likely to like job (81-100%)

The results from the 137 respondents who answered both questions are represented in Figure 9. The proportion reporting they are “Very Likely” or “Likely” to enjoy the job is larger when 100 applicants are seen rather than 10. This shift in the distribution is not very large, but it implies that individuals do believe there is a positive quality signal as the number of applicants shown rises. The shift in the distribution is similar for female and male respondents.




Figure 9: Likelihood of Liking Job

Note: This Figure shows the proportion of respondents who said they believed they were likely to "like" a job if they saw 10 versus 100 applicants.




Figure 10: 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.




Figure 11: 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). The left panel is for male job seekers and the right panel is for female job seekers. A correlation closer to -1 is evidence of competition aversion. A correlation of 1 is evidence of herding.



Figure 12: Average Correlation Between Number Seen and Start Application By Number Seen

Note: This figure shows the average correlation for all people who saw this specific number on a job posting. First, for each person in the treatment a correlation between number seen and starting application was computed. Then, for each number shown I find the average correlation for people who saw that number. So each dot on this graph represents an average of multiple user's correlations. For example, for the 35,251 users in the treatment who viewed a job posting with a 1 shown the average correlation those users had was 0.05 (min: -1, max: 1). While for the single user who viewed a job posting with a 4,000 shown they had a -0.25 correlation. A correlation closer to -1 is evidence of competition aversion. A correlation of 1 is evidence of herding.