UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
Search and Signaling on an Online Labor Market

Permalink
https://escholarship.org/uc/item/2bt1q6vw

Author
Lu, Sibo

Publication Date
2018

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Search and Signaling on an Online Labor Market

by

Sibo Lu

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor John Morgan, Chair
Professor Noam Yuchtman
Professor Ming Leung
Professor Joshua Blumenstock

Summer 2018

Search and Signaling on an Online Labor Market

Copyright 2018
by
Sibo Lu

Abstract

Search and Signaling on an Online Labor Market

by

Sibo Lu

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor John Morgan, Chair

This dissertation consists of three papers on search and signaling on a large online labor market. The abstract of each chapter is as follows:

Online resume and job post sites like LinkedIn and Monster.com has made it increasingly easy for employers to search for and invite workers who would not otherwise have applied to their job post. I leverage the A/B test of a new resume screening tool on a large online freelancing platform to identify which types of workers and employers are most likely to benefit from lower screening costs in the invitation channel. The screening tool increased the number of workers contacted by 12% but there was no increase in overall hiring. Treated employers substituted away from screening and hiring workers in their existing applicant pool to sending invitations. They hired workers with higher platform reputation at higher cost. Moreover, invitations are more concentrated among a minority of workers than applications. Thus the screening tool likely led to an increase in inequality for workers on the platform. Treated employers also had better job outcomes - at least half of this effect is due to employers avoiding hiring less well matched workers from their applicant pool. Benefits accrue primarily to employers looking for expert freelancers and willing to pay higher prices.

How do employers respond to hiring tools? We examine the hiring decision as a process comprised of a series of decisions. We claim hiring indicators act as a “minimal cue” by elevating a job applicant to being noticed early in the hiring process, but giving way to the information an employer privately gathers later in the hiring process. Regression discontinuity analyses of over 1.5 million job applications by freelancers for over 150,000 short-term jobs on an online market for contract labor demonstrate support for our contention. Dramatically, being algorithmically recommended increases a job applicant’s unconditional likelihood of being hired over applicants of observationally similar quality by approximately 50%. 70% of this effect can be attributed to the increase in likelihood of recommended job applicants being viewed. Furthermore, conditional on being interviewed, the recommendation has no influence on the hiring decision. Underscoring the idea of it being a ’minimal cue’, the effect of a recommendation is stronger

for low value jobs than high value ones.

The rapid growth of online information on workers, like LinkedIn's profiles or Monster.com's resume database, has dramatically lowered the cost for employers to directly reach out to (headhunt) workers. The promise is that by providing employers with an additional channel of sourcing matches, it will increase the probability of filling vacancies and the average quality of hires. I construct a search theoretic model of hiring that explicitly models both the headhunting channel and the application channel. I study equilibrium outcomes when workers can optimally respond by deciding whether or not to apply. I show that while lower headhunting costs improve the average quality of hires, it can actually decrease the probability of filling vacancies in equilibrium. This is due to workers optimally choosing to apply with a lower probability.

Acknowledgements

I will forever be grateful to my committee for their advice, patience, and support. A special thank you goes to John, my chair, who I am honored to call my dear friend. His strength in the face of life's challenges will always be an inspiration to me.

I would not have been able to finish this thesis without my friends at Haas. Moshe: my coauthor, sounding board, and "foot in the door" to the online platform. Oren: always there to motivate me and distract me from the doldrums. Abhay: my go-to for commiserating and yet so often a source of optimism too. Sam, Nan, Paulo, and so many more who made the day to day so much easier.

And of course my wife Alyssa, who had to put up with me during this formidable journey; not at all an easy task.

Chapter 1: Who benefits from lower screening costs? Evidence on employer attention, hiring, and job outcomes from the introduction of a new screening tool on an online labor market platform.

Abstract

Online resume and job post sites like LinkedIn and Monster.com has made it increasingly easy for employers to search for and invite workers who would not otherwise have applied to their job post. I leverage the A/B test of a new resume screening tool on a large online freelancing platform to identify which types of workers and employers are most likely to benefit from lower screening costs in the invitation channel. The screening tool increased the number of workers contacted by 12% but there was no increase in overall hiring. Treated employers substituted away from screening and hiring workers in their existing applicant pool to sending invitations. They hired workers with higher platform reputation at higher cost. Moreover, invitations are more concentrated among a minority of workers than applications. Thus the screening tool likely led to an increase in inequality for workers on the platform. Treated employers also had better job outcomes - at least half of this effect is due to employers avoiding hiring less well matched workers from their applicant pool. Benefits accrue primarily to employers looking for expert freelancers and willing to pay higher prices.

Introduction

Online job sites have become the predominant place to find work in the US, with 76.3% of unemployed job seekers looking online in 2011 (Faberman and Kudlyak 2016). Recently major job sites have increasingly promoted the ability for employers to invite workers from a vast pool of potential matches, without limiting themselves to a small set of traditional applicants. LinkedIn is the most well known and earliest example of this, but Monster.com, Zipcruiter, and Indeed now all have their own competing feature. This invitation channel relies on two components: a large database of available worker profiles and a quick and cost effective way to screen through them.

Despite the ubiquity of online job sites, there is limited evidence on their effects on match quality, overall hiring, or labor composition. As Kuhn (2014) puts it: "It remains to be seen whether making each "horse" in the race for jobs faster yields substantial aggregate benefits in terms of improved quality of matching and lower overall unemployment". What evidence there is has been negative; Kory and Pope (2014) exploit geographic variation in the sharp rise of Craiglist and show no effect on the unemployment rate. Follow up study by Brencic (2016) provide evidence that Craigslist had instead cannibalized existing online job sites. Beyond match efficiency, to the extent that job sites direct employer attention and there is heterogeneity in adoption, there may also be welfare concerns as they affect who is hired. For example, if invitations shift employer attention from active job seekers to passive job seekers, we may expect an increase in unemployment duration.

One reason that the effects of job sites are difficult to study is that most only observe employers' search activity, but not interview, hiring or eventual job outcomes. There is also extensive multi-homing - Ziprecruiter for example sends your job post to "100+ Job Boards"- so job sites often have only a partial view of the employer's activity. To overcome this, I use data from a large online freelancing platform. Job search and employer screening on this platform is analogous to the process on major job sites. In particular, this platform also allows for both traditional applications and enables employers to directly contact the entire pool of freelancers on the platform. Since work is completed on the platform, I observe the entire process from initial search, interviewing, hiring, and eventual job outcome. I exploit an experiment conducted by the platform that introduced a new screening tool to treated employers. This screening tool allows employers to quickly filter the pool of freelancers on the platform by characteristics like platform reputation, skills, and asking wages.

Looking at the invitation channel alone, a stylized way to view this intervention is through the lens of a standard fixed sample search model like Stigler (1961). Instead of a customer searching for a good in the presence of price dispersion, here it is an employer searching for a qualified employee among a pool of workers. The pool of workers contains a distribution of worker quality, and the employer pays a fixed cost to search each worker and obtain an informative signal over each searched worker's quality. By filtering out poorly matched workers, the screening tool can be modeled as both increasing the mean quality and decreasing the variance of the pool over which the employer will search. This model predicts that 1) employers will be more likely to hire and 2) workers hired will ex post be of higher quality i.e. be more likely to complete the job to the employer's satisfaction. The effect on the number of workers invited is ambiguous. The screening tool increases the benefit of each additional search, so we might expect an increase, but there is decreasing marginal returns. In the extreme case, if the screening tool reduced variance to zero, then the employer would optimally only search once if at all.

As expected, treated employers sent more invitations to freelancers, an increase of 1 additional application every 7.7 job posts. These invitations translated to an increase in the probability of hiring via the invitation channel. Treated employers screened for and hired freelancers with higher platform reputation and had better job outcomes. Surprisingly however, there was no effect on the fill rate. Instead, the increase in hiring via invites crowded out hiring via applications. This crowd out effect is also seen in employer attention, as measured by interviews. At least half of the improvement in job outcomes came from hires from the application channel, suggesting that the screening tool helped employers avoid least qualified applicants.

The extent of crowd-out was partially due to heterogeneous take-up: the experiment showed no change in the probability of invites being sent at the job level. Since employers who send invitations are already much more likely to fill their vacancies, this limited the upside of the treatment. I also show suggestive evidence that the screening tool improved fill rate for employers looking for higher expertise and higher cost freelancers, while bargain hunting employers may have been negatively affected. Thus the screening tool primarily benefited both employers and freelancers at the top end. Looking beyond the experiment, I show that freelancers bid higher when invited and invitations are concentrated among a minority of freelancers on the platform. Thus substituting towards invites likely increased inequality on the platform.

This rest of the paper is organized as follows. Section 1 lays out the setting and describes the experiment in detail. The main results of the experiment are reported in Section 2. In Section 3 I explore why the screening tool had no effect on fill rate. Finally Section 4 concludes the paper with a discussion on the implications on welfare distribution across workers.

Related Literature

There is an extensive literature in labor and personnel economics on how information affects the hiring process. Several recent papers have examined how the addition or removal of a specific piece of information on applicants have changed the composition of hires, motivated by concerns over minority welfare. Autor and Scarborough (2008) show that skill test hiring increased the productivity of hires with no measurable impact on minority hiring. Agan and Starr (2016) conduct a pre-post resume audit study examining the effect of "Ban the Box" policies. They find that removing information on previous criminal convictions caused employers to substitute to other signals and decreased call-back rates for Black applicants. Shoag and Clifford (2016) finds that the use of credit scores likely increased total employment but harmed groups with poorer credit scores. The technology reported relates to this literature in that the screening technology changes the cost of acquiring signals on match quality and has implications for the composition of hires. However the filters allow employers to screen on many different pieces of information, each with different implications.

This paper also relates to the literature on the effectiveness of job search assistance. Though the intervention reported here is on the employer side, it lowers the cost of employers acquiring information on workers, much like job search interventions such as resume editing or job faires. Recent papers highlight crowd out effects, where employers substitute from non-treated workers to treated workers, with little net effect on overall employment. Crepon et al (2013) finds this in a field experiment where unemployed youths in France were randomized into a reinforced counseling scheme. Unemployed youths assigned to the program were more likely to have found a stable job than those who were not but there were minimal net benefits in overall youth employment. Gautier et al (2018) finds a similar result for a Danish activation program for unemployed workers. I also find a 1 to 1 crowd out from the introduction of the filters, with no effect on the probability that a vacancy is filled. This is despite a low overall probability of fill relative to conventional labor markets. Access to detailed employer and job post data, as well as clickstream filter usage data, allows me to explore the role of heterogeneous adoption in limiting effects on fill rate.

A feature that differentiates this paper from earlier literature on information is the richness of the platform information. Since work is completed on platform, I can directly observe and link search activity, hiring, and outcomes. Growing. Most closely related papers using similarly rich data from online labor platforms are Horton (2017b) and Barach and Horton (2018). The first examines a field experiment where treated employers were presented with a set of freelancers to invite. These freelancers were chosen by the platform using a machine learning algorithm, trained on previous invitations made by employers on the platform. Horton report a substantial increase in fill rate (from a baseline about half of that reported here) with no detectable crowd out of other applicants. However the experiment was under powered to detect even large effects on job outcomes. Barach and Horton (2018) consider an experiment that hid past wages of applicants from treated employers. They show that employers increased screening via interviews and subsequently hired freelancers at lower wages, with no detectable decrease in fill rate.

1 Empirical Setting

This paper leverages an experiment that took place on a large online labor market platform between between November 2016 and February 2017. This platform specializes in connecting employers with gig-economy freelancers, similar to Toptal, Upwork, and Freelancer.com. As of the end of 2016, millions of freelancers have created profiles and hundreds of thousands of jobs have been posted. This platform is a particularly informative setting as work is completed on the platform, unlike recruitment sites like Monster.com or Indeed.com whose data ends at interview or hire. I am thus able to observe not only hires, but also the amount spent on each hire and the employer's feedback at the end of the job.

Employers on the platform are primarily individuals hiring for small firms or small teams within medium to large firms. Job postings are organized by job categories that represent a wide range of tasks that can be accomplished virtually. Based on dollars spent, the top skills are technical skills, such as web programming, mobile applications development (e.g., iPhone and Android), and web design. The top five countries for employers are: the United States, the United Kingdom, France, Germany, and Israel. Top top five countries for workers are: the United States, India, Philippines, Russia, and Ukraine.

There is a growing body of research in labor economics and personnel economics using data from online labor markets. Pallais (2014) showed the high value employers place on past on-platform work experience by randomly hiring new entrants. Horton (2017) showed platforms can improve match efficiency by recommending workers to job posts. Barach and Horton (2018) report the results of a field experiment where the platform randomly hid applicants' compensation history. They showed that employers increased screening via interviews and subsequently hired workers at lower wages.

Posting a job

Hiring on the platform is qualitatively similar to hiring in conventional labor markets. First, a would-be employer on the platform creates a job post, describing the nature of the work and specifying required skills. Additionally, the employer chooses a contractual form (hourly or fixed-price) and estimates how long the project is expected to last. The job post is then submitted for review by the platform and then posted to the marketplace. All the information provided by employers is viewable by potential applicants. Additionally, the platform also presents verified attributes of the employer, such as their number of past jobs and their average wage rate paid. Figure 1 shows how a typical job post appears to potential applicants on the platform.

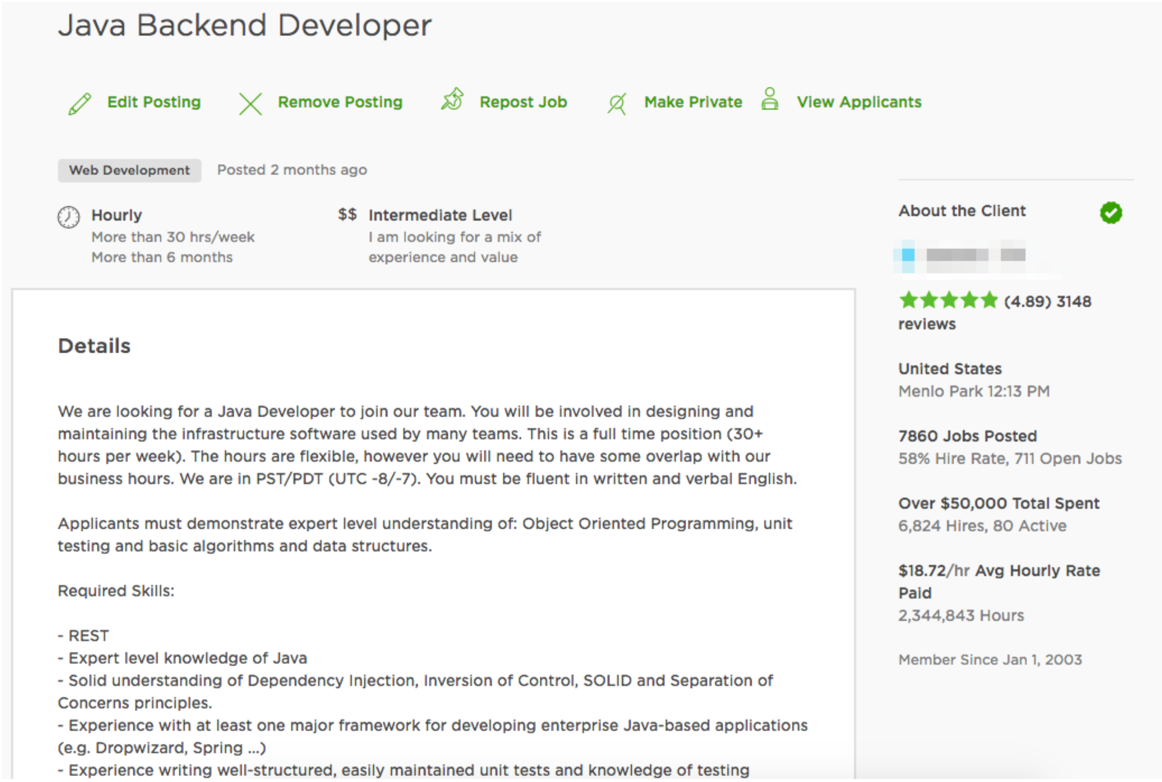

Figure 1: Job Post for Java Backend Developer

Applying to a job

Freelancers can apply to any public[1] job post on the platform. Their application includes a bid specifying an hourly rate or a fixed amount that they are willing to work for and a cover letter explaining why they should be hired. For the rest of the paper I refer to these applications as "organic applications".

After applying, the application immediately appears in the employer's "applicant tracking system" or ATS. This is the dashboard the employer sees upon logging into his account and clicking on an open job posting. Each application in the ATS shows the applicant's name, picture, bid, self-reported skills, country, and a few pieces of platform-verified information, such as total hours-worked and average feedback rating from previous projects (if any). Figure 2 shows the employer's view of the ATS.

[1]Some employers may choose to post private jobs, which are only visible to freelancers employers directly invite to apply. These job posts are often used to rehire a freelancer the employer has previously worked with. Private jobs are not included in this analysis.

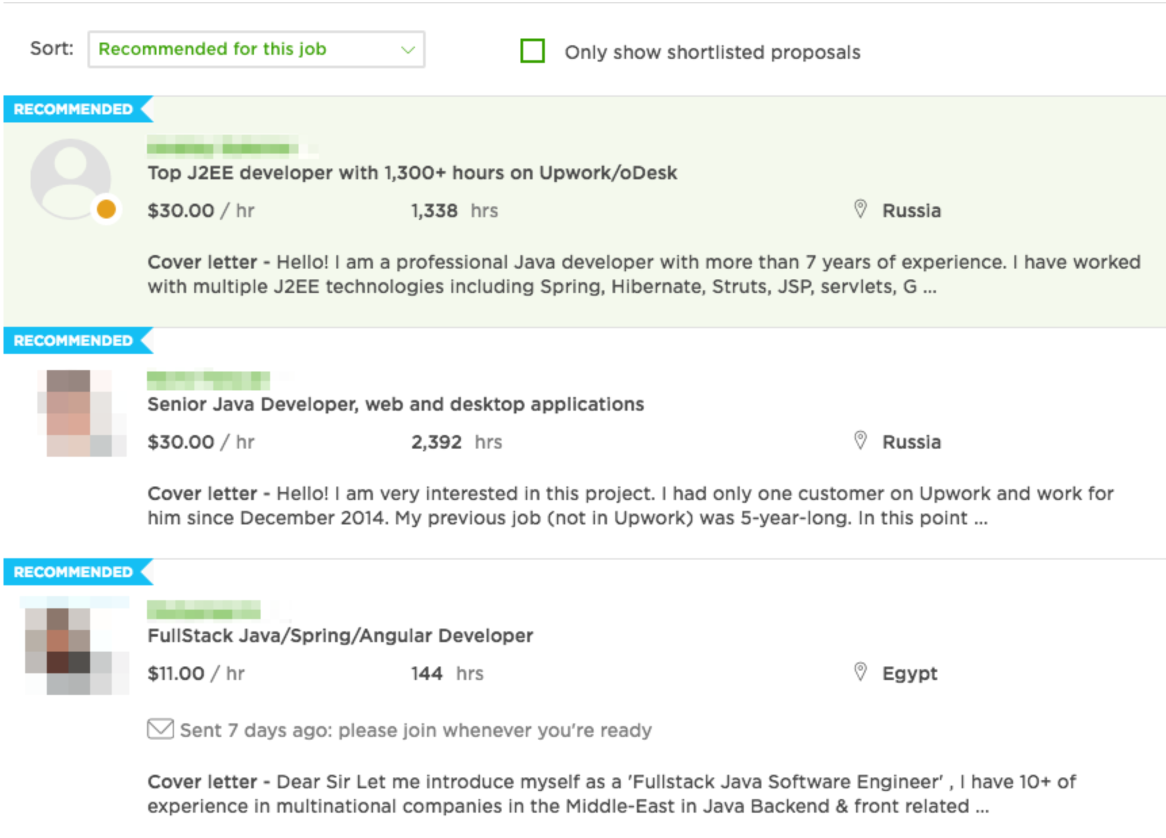

Figure 2: The Application Tracking System (ATS)

Sending Invitations

In addition to reviewing organic applicants, employers can directly reach out to freelancers on the platform by sending an "invitation to interview". After the job post is approved by the platform, the employer is presented with a list of freelancers that the platform believes might be good matches for the job. Employers can send invitations via this list, or use the search features on the same page to look for freelancers themselves. This page is shown in Figure 3.

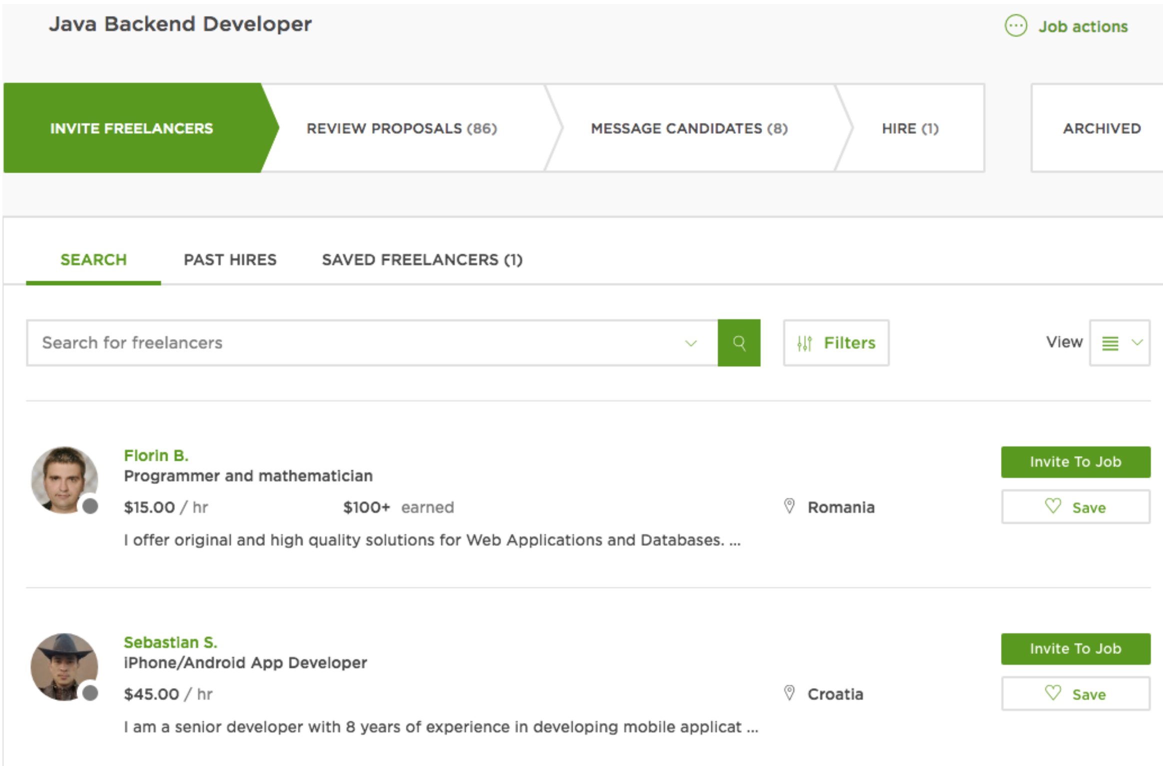

Figure 3: Invite Freelancer Page

If the employer chooses to invite a freelancer, the freelancer will see a brief message informing him of the employer's interest and a summary of the job post. An example of an invitation for an Administrative Assistant job post is shown in Figure 4.

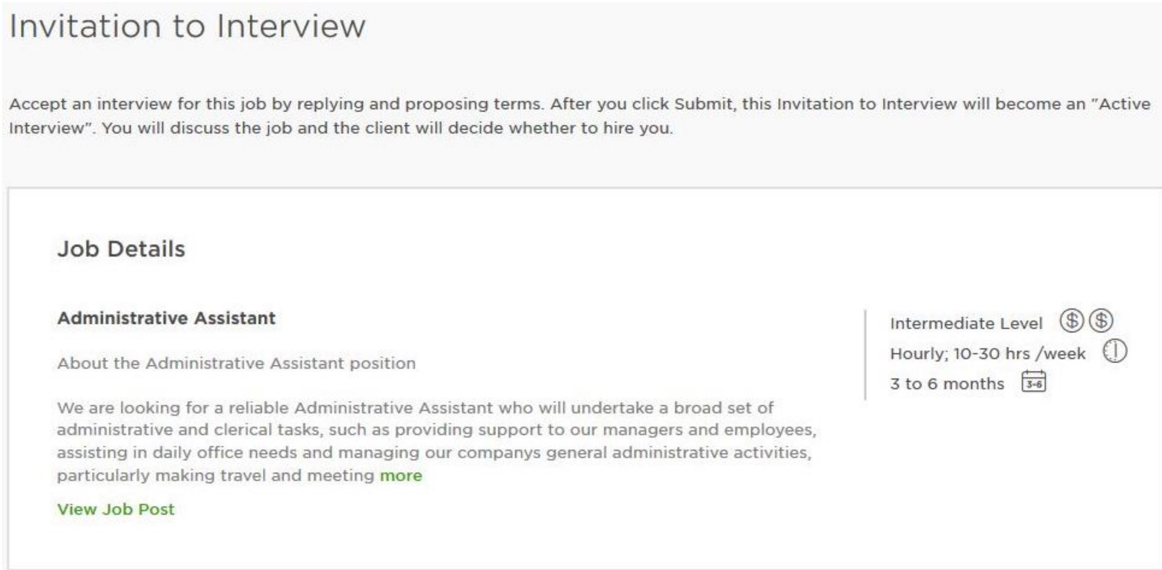

Figure 4: Invitation as received by the freelancer.

If the freelancer chooses to apply after receiving the invitation, the process is the same as that for organic applicants. For the rest of the paper I refer to these applications as "invited applications".

Viewing Applications & Interviewing Applicants

An employer is able to learn more about a specific applicant by clicking on any of the applications listed in the ATS. The detailed application consists of a cover letter and the applicant's profile, which is similar to a traditional resume. An employer viewing the freelancer's profile would see the freelancer's self-written overview, past work history with prior employer feedback, and any skill tests that they may have completed. Figure 5 shows how a freelancer's profile would look to an employer on the platform.

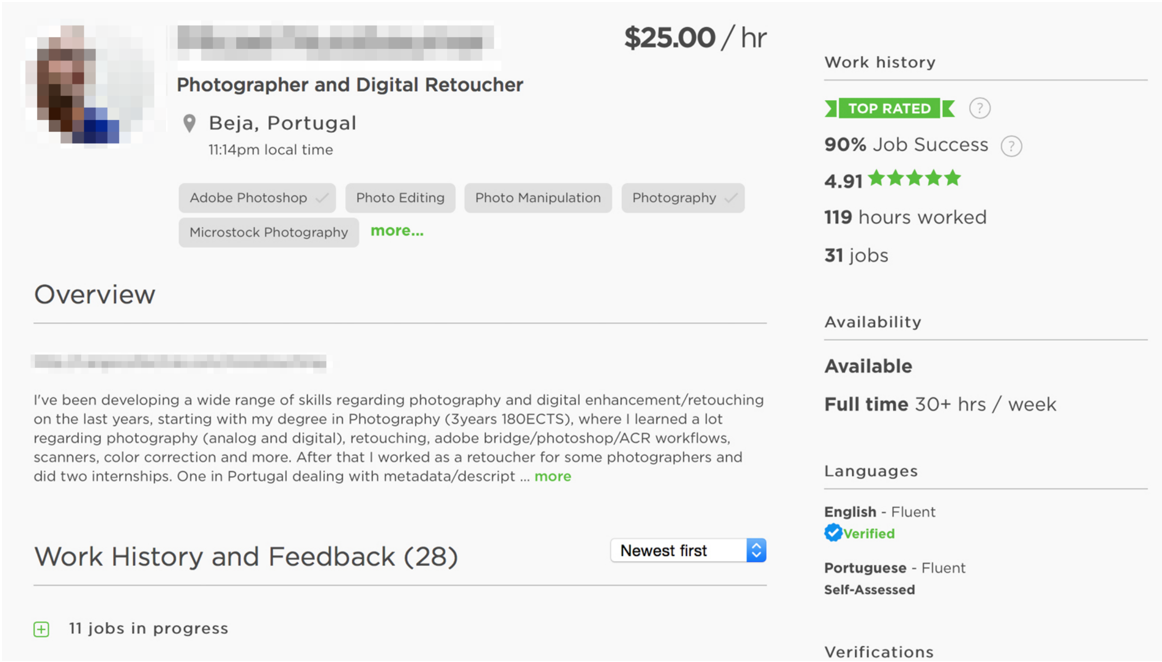

Figure 5: A freelancer's profile.

After viewing an application, the employer can either message the applicant to conduct an "interview" or directly hire the applicant. Employers are encouraged by the website to interview their applicants via their internal messaging system. The platform disallows contact information in both job posts and freelancer profiles, so while interviews may occur via an external service like Skype, employers must start the process via the platform.

Hiring Applicants

An employer is free to hire whomever they wish. The employer hires the worker on the terms proposed by the worker or make a counteroffer, which the worker can accept, reject, or negotiate. If an employer chooses to hire anyone, over 90% of the time they hire only one freelancer. Once a freelancer is hired, employer and employee exchange work details and the job is completed virtually. Payment is conducted through the website.

1.1 Experimental Design

In November 2017, the platform introduced a new search function on the Invite Freelancer page. Previously employers could only search using a text query, for example "python developer", by typing in the search box near the top of the page. The new search function provided a host of explicit options via Filters, allowing the employer to narrow results by freelancers' amount earned in previous jobs, hourly rate stated on their profile, previous job success, and many other characteristics. The complete set of filters are shown in Figure 6. Filters were provided in addition to the original search box.

As is common practice for the platform with any major new functionality, Filters were first introduced as an experiment. The experiment ran from November 2016 to February 2017 and the experiment was randomized at the employer level. Employers were assigned to treatment or control right after they post their first job during the experiment period, so that each employer either sees the Filters for every job they post or never see the Filters. The platform did not otherwise publicize any new search functionality during the experiment period.

For this paper the analysis sample is restricted to only public jobs, which any freelancer on the platform can apply to. When posting a job, employers have the option of keeping the post private so that only invited freelancers can apply. These are often used to rehire freelancers employers had previously employed or to hire referred freelancers. Though these hiring behaviors are very interesting, they are beyond the scope of this paper. Since the data has few but large outliers, I also drop job posts above the 99th percentile in number of organic applications or in number of invitations sent. Similarly, I drop all job posts by employers who posted 21 or more job posts during the experiment period.

Balance across employer characteristics between treatment and control is shown on Table 1. There is a slight imbalance in job characteristics. A back of the envelope calculation shows that this imbalance would at most account for at most 0.5% of main results; nevertheless I control for these job characteristics in all regressions.

The rest of the paper will primarily focus on the platform's largest job category: Web,

Table 1: Balance Table: All

	Control	Treatment	Difference
Job Post Characteristics			
Proportion Expert/Expensive	0.24 (.00)	0.24 (.00)	0.01 (.00)***
Proportion Hourly Jobs	0.43 (.00)	0.43 (.00)	-0.00 (.00)
Employer Characteristics			
Proportion New Employers	0.23 (.00)	0.23 (.00)	0.00 (.00)
Proportion from US	0.48 (.00)	0.47 (.00)	-0.00 (.00)

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$

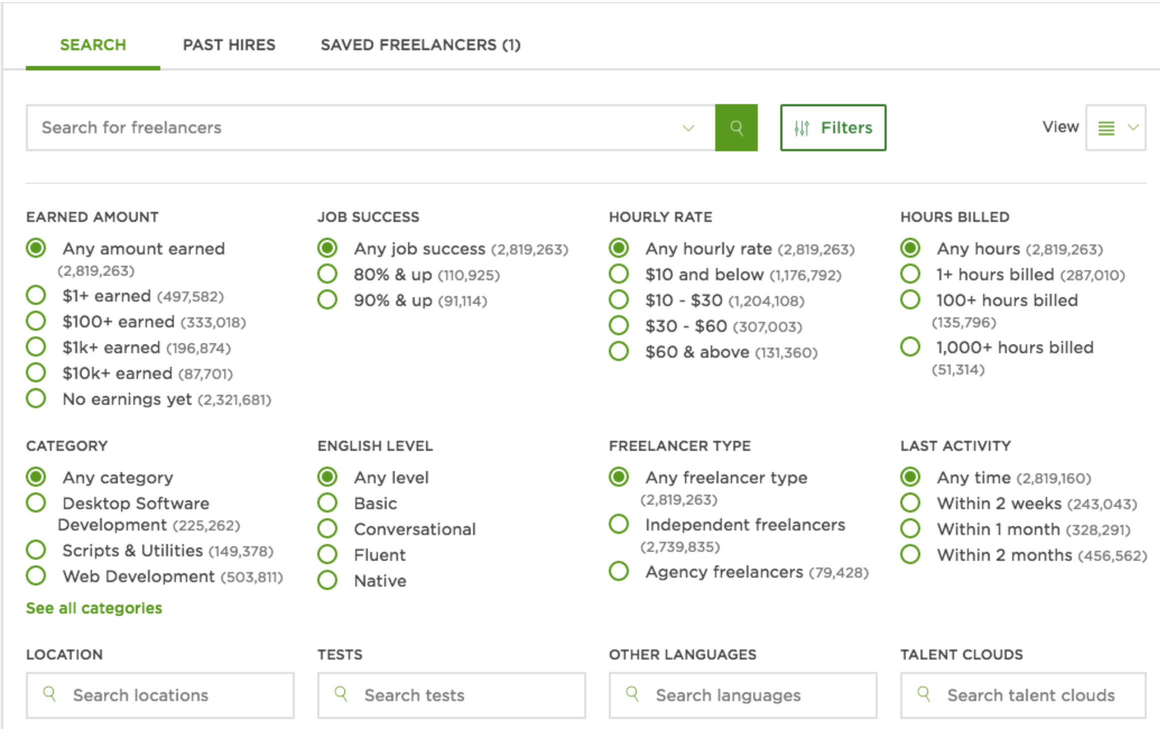

Figure 6: New Search Functionality: Filters

Mobile and Software Development. This category represents half of all earnings on the platform. Results across all categories are reported in the Appendix.

Table 2: Balance Table: Web, Mobile and Software Dev

	Control	Treatment	Difference
Job Post Characteristics			
Proportion Expert/Expensive	0.26 (.00)	0.27 (.00)	0.01 (.00)***
Proportion Hourly Jobs	0.42 (.00)	0.42 (.00)	-0.01 (.00)*
Employer Characteristics			
Proportion New Employers	0.25 (.00)	0.25 (.00)	-0.00 (.00)
Proportion from US	0.45 (.00)	0.44 (.00)	-0.00 (.00)

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$

2 Experimental Results

Main experimental results are presented below in chronological order with respect to the hiring process, from sending invitations to post job completion outcomes. Unless otherwise noted, results are from a linear regression at the job post level:

$$y_j = \alpha + \beta \text{Treated}_j + \mathbf{X}_j + \epsilon_j$$

where y_j is an outcome of interest, Treated_j is an indicator equal to 1 if the job was posted by a client in the treatment group, and $\mathbf{X}_j$ is a collection of pre-randomization job opening and employer characteristics.

2.1 Recruitment and Screening

Table 3: Effect on Invitations

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
Invites	0.33 (0.059)***	11.74
Extensive Margin	0.00 (0.005)	0.87
Intensive Margin	0.64 (0.103)***	9.70
Accepted Invites	0.13 (0.024)***	13.62
Probability Invite Accepted	0.00 (0.004)	0.36

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$ Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.

As expected, treated employers sent more invitations to freelancers. As shown in Table 3, these invitations resulted in a 14% increase in applications from the invitations channel, approximately 1 additional application every 7.7 job posts, with no overall change in the probability of acceptance by the freelancer. However there was no increase in the extensive margin - the percentage of job posts where the employer sent at least one invitation is effectively identical between treatment and control job posts.

Improving the employer's ability to screen the invitation pool might decrease the number of organic applicants that the employer screens, as the employer substitutes attention from one pool to another. On the other hand, if the new screening tool increases the employer's overall engagement on the platform, it may instead crowd in attention on organic applicants. This latter effect is most likely to be seen at the extensive margin, driven by jobs posted by employers with low engagement.

To test this, I look at the number of organic applicants the employer interviews. Here "interview" is defined as the employer sending a message on the platform to the applicant. As shown in Table 4, treated employers interviewed almost 4% fewer organic applicants, consistent with substitution of attention. There is a decrease in the extensive margin, i.e. the probability that the employer interviewed any organic applicant, so there is no evidence that the new screening tool led to any crowd-in of attention on organic applicants. It would be interesting to see how the total amount of employer attention was affected. Unfortunately when the employer sends an invite, the platform automatically generates a message to the invitee, so it is difficult to identify which invited applicant was interviewed.

Table 4: Interviews of Organic Applicants: Web, Mobile and Software Dev

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
Organic Interviews	-0.07 (0.026)***	-3.89
Extensive Margin	-0.01 (0.004)***	-1.88
Intensive Margin	-0.05 (0.035)	-1.91

Notes: Significance Indicators: $p \leq 0.1 : *, p \leq 0.05 : **, p \leq 0.01 : ***$. Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.

2.2 Effect on Hiring

The substitution effect holds true in employers' hiring decisions. There is a precisely measured zero effect on both mean hires per job and the probability that any hire was made (fill rate) between treatment and control. However, there is a one to one substitution from hiring via the organic applicant channel to hiring via the invitation channel. It is surprising that a new screening technology did not increase the fill rate, despite affecting how the employer searched for and screened applicants. Potential explanations are discussed below in section 3.

The new screening tool allowed employers to easily narrow down the invitation pool by worker characteristics. As is common with market platforms, a key proxy for quality is platform reputation. While the platform provides several signals of worker quality, including skill tests, contracts completed, and previous client feedback, the most prominently displayed signal is Job Success. This signal represents the percentage of jobs completed with positive feedback by the worker. It is only calculated for freelancers with at least 4 distinct employers in last 2 years. As shown in Table 6, treated employers were more likely to invite and hire workers with a high Job Success signal. Here the regression is at the worker level.

Table 5: Effect on Hires: Web, Mobile and Software Dev

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
Mean Hires	-0.00 (0.005)	-0.38
Mean Invite Hires	0.01 (0.003)**	4.73
Mean Org Hires	-0.01 (0.005)*	-1.92
Probability of Hire	0.00 (0.004)	0.07
Probability of Invite Hire	0.01 (0.003)**	5.05
Probability of Org Hire	-0.01 (0.004)*	-1.68

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$ Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.

Table 6: Percentage of contractors with $\geq 90\%$ Job Success: Web, Mobile and Software Dev

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
Invites	0.03 (0.005)***	4.47
Accepted Invites	0.02 (0.005)***	3.08
Hired Invites	0.03 (0.011)**	3.92
Org Applicants	0.00 (0.002)	0.64
Interviewed Org Applicants	0.00 (0.005)	0.68
Hired Org Applicants	0.01 (0.007)	1.81

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$. Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.Job Success represents the % of jobs completed with positive feedback. It is only calculated for freelancers with at least 4 distinct employers in last 2 years. Freelancers with no job success are included with the dependent variable set to 0, not dropped.

2.3 Effect on Outcomes

Since jobs are completed on the platform, I observe the amount the worker was paid as well as the eventual outcome of the job. In this subsection all regressions are at the contract level. A contract is created when a worker is hired to a job post, so it is unique at the job post by contractor level. All regressions in this subsection are at the contract level. At the end of contract, the employer is asked "How likely are you to recommend this freelancer to a friend or a colleague" on a scale from 0 to 10 with 0 being "Not at all likely" and 10 being "Extremely Likely". This feedback is private, i.e. not directly shown to the freelancer. Since the platform handles all billing between employer and worker, it also collects data on refund requests, disputes, and complaints by either side.

These proxies for employer satisfaction are aggregated by the platform into 3 categories: Good, Neutral, Bad. A contract has a Good outcome if the client provides a score of 9 or 10 and there are no disputes or complaints from either the employer or the worker. A contract has a Bad outcome if a client provides a score of 6 or below or if there are any refund requests, disputes, or complaints. All other contracts receive a Neutral outcome.

As shown in Tables 7 and 8, there is a 6% decrease in the probability that a contract ends in a Bad outcome. Though imprecisely measured, this improvement occurs with both workers hired via the organic application pool and the invitation pool. Since the experiment did not affect the composition or size of the applicant pool, this suggests that the new screening tool helped employers avoid hiring poorer matches from the organic application pool by substituting to better matches from the invitation pool. There is a small and noisily measured improvement in the probability that a contract ends in a Good outcome as well.

Table 7: Percentage of completed contracts with a Good Outcome: Web, Mobile and Software Dev

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
All Hires	0.01 (0.007)	1.09
Invite Hires	0.00 (0.012)	0.14
Org Hires	0.01 (0.007)	1.16

Notes: Significance Indicators: $p \leq 0.1 : *, p \leq 0.05 : **, p \leq 0.01 : ***$. Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.the platform categorizes a completed contract as having a Good outcome if 1) client provides a feedback score at least 9 out of 10, and 2) there are no complaints or disputes from either side. Regression is at the contract level.

Table 8: Percentage of completed contracts with a Bad Outcome: Web, Mobile and Software Dev

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
All Hires	-0.01 (0.005)**	-6.32
Invite Hires	-0.01 (0.008)	-5.50
Org Hires	-0.01 (0.005)*	-6.22

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$. Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.the platform categorizes a completed contract as having a Bad outcome if 1) client provides a feedback score of 6 or below out of 10, or 2) there are complaints or disputes from either side. Regression is at the contract level.

The question remains whether the improvement in match quality also came at a higher monetary cost to the employer. I observe the total amount the employer paid the worker for each contract. The platform allows for two types of payments - fixed price or hourly rate. Fixed price jobs are typically paid with a lump sum at the end of the contract. Hourly rate jobs pay based on the number of hours worked and the worker is paid periodically, e.g. biweekly. The payment type is specified when the job is posted and shown to the worker within the job post. Since payment amounts are highly skewed, regressions are run on logged values.

Total amount paid by treated employers is substantially higher in both fixed price and hourly rate jobs. For hourly rate jobs I can also look at the hourly rate agreed upon at the start of the contract. As shown in 11, the difference between treatment and control are directly the same though smaller and noisily measured.

Table 9: Log(Amount) paid in each fixed price employment contract: Web, Mobile and Software Dev

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
All Hires	0.13 (0.048)***	5.23
Invite Hires	0.13 (0.091)	4.43
Org Hires	0.13 (0.054)**	5.21

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$. Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.Contracts at 99th percentile or above are dropped. Regression is at the contract level. Percentage change displayed is calculated on non-logged values for amount paid.

Table 10: Log(Amount) paid in each hourly rate employment contract: Web, Mobile and Software Dev

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
All Hires	0.14 (0.062)**	3.61
Invite Hires	0.32 (0.115)***	8.28
Org Hires	0.08 (0.072)	2.13

Notes: Significance Indicators: $p \leq 0.1 : *, p \leq 0.05 : **, p \leq 0.01 : ***$. Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.Contracts at 98th percentile or above are dropped. Regression is at the contract level. Percentage change displayed is calculated on non-logged values for amount paid.

Table 11: Log(Hourly Rate) paid in each hourly rate employment contract: Web, Mobile and Software Dev

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
Hires	0.02 (0.012)	0.66
Invited Hires	0.02 (0.023)	0.81
Organic Hires	0.01 (0.014)	0.55
Interviewed	0.00 (0.010)	0.12
Invites	0.01 (0.013)	0.21
Org Applicants	-0.00 (0.006)	-0.16

Notes: Significance Indicators: $p \leq 0.1 : *, p \leq 0.05 : **, p \leq 0.01 : ***$. Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.Contracts at 98th percentile or above are dropped. Regression is at the contract level. Percentage change displayed is calculated on non-logged values for amount paid.

3 Why was there no effect on vacancies?

The screening tool increased the number of applications, shifted the distribution of applications towards higher quality freelancers, and improved eventual job outcomes. Thus it is surprising that there was no increase in the probability of hire. This section explores 3 potential reasons. First I show that vacant job posts are not substantially different in observables from filled job posts, thus it is unlikely that none of the vacant job posts would have remained vacant regardless of screening technology. Rather, as shown in the second subsection, employers that use invitations are substantially more likely to hire. Thus, since the intervention did not shift the extensive margin, its effect on fill rate is limited. Third I show suggestive evidence that employers willing to pay more for more experienced freelancers were more likely to fill their job in treatment vs. control. This benefit was balanced out in aggregate by a negative effect on employers looking for cheaper freelancers.

3.1 Vacant job posts are not drastically different from filled job posts

One concern is that these jobs would never have been filled, no matter the screening technology. The unfilled job posts may be posted by employers who never intended to hire, for example. In the analysis sample, 53.1% of unfilled jobs were posted by employers who had previously hired on the platform. This compares to 74.9% for filled jobs. Or these unfilled job posts may have been particularly unattractive to freelancers, for example they may be poorly specified or make unreasonable demands. The median unfilled job received 8 applications from freelancers who've worked on the platform, compared to 10 for filled jobs. Thus while unfilled jobs are perhaps on average of lower quality, it appears that the majority could have been filled if a reasonable match was found.

Another concern is that there simply were no qualified freelancers available on the platform for the unfilled job posts. However the unfilled job posts do not seem substantially more demanding in their requested freelancer experience, as shown by Table 12. They are also not concentrated in a few job categories with few hires, as shown by Table 13.

Requested Freelancer Expertise	Filled Jobs	Unfilled Jobs
Cheap/Inexperienced	27.4%	24.4%
Expert/Expensive	21.0%	25.0%
Intermediate	48.9%	46.8%
Unspecified	2.7%	3.7%

Table 12: Percentage of Job Posts with Each Requested Freelancer Expertise

Job Category	Filled Jobs	Unfilled Jobs
Accounting & Consulting	1.5%	1.9%
Admin Support	9.0%	6.3%
Customer Service	0.5%	1.1%
Data Science & Analytics	2.6%	2.2%
Design & Creative	24.4%	17.0%
Engineering & Architecture	2.3%	2.5%
IT & Networking	2.9%	3.5%
Legal	0.8%	1.0%
Sales & Marketing	8.1%	14.0%
Translation	4.2%	2.1%
Web, Mobile & Software Dev	29.0%	36.8%
Writing	14.9%	11.5%

Table 13: Percentage of Job Posts in Each Job Category

3.2 Employers using invitations are already likely to hire

As shown in table 3, the screening tool did not change the probability that an employer sent any invitations for their job post. Looking at job posts in the control group, employers who sent at least one invitation are 9.4 percentage points more likely to hire. One hypothesis is that employers more familiar with the platform are more likely to use the invitation channel. However, as shown in Table 14, this is not the case here. Rather, by looking at only experienced employers and controlling for job post characteristics, employers who had used the invitation channel in previous job posts are substantially more likely to do so for their current job post. This suggests an employer specific unobservable - perhaps lower search cost or higher intent to hire - that is associated with their willingness to use the invitation channel. Since employers that use invitations are already likely to hire, the screening tool's benefit on fill rate will be limited if it only increases invitations on the intensive margin.

Table 14: Invitation Usage vs Previous Usage

	Dependent variable:			
	Sent Invite			
	All	Experienced	Experienced	Experienced
	(1)	(2)	(3)	(4)
First Job Post	−0.006			
	(0.005)			
Previously Sent Invite		0.119***		
		(0.005)		
Previously Had Invite Accepted			0.121***	
			(0.005)	
Previously Hired Invite				0.062***
				(0.006)
Observations	65,311	48,899	48,899	48,899
R^2	0.024	0.034	0.036	0.026
Adjusted R^2	0.024	0.034	0.035	0.026

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$ All regressions include the fixed effects for requested freelancer expertise tier, job category, and whether the job is fixed price or hourly wage. Standard errors are clustered at the employer level. Regressions 2 to 4 include only posts where the client had previously posted before on the platform.

3.3 Treatment Heterogeneity

The analysis above has looked at aggregate measures of how the screening tool affected employer hiring. There may be heterogeneities that balance each other in the aggregate. In particular, more attractive job posts or more selective employers might benefit more from a new screening tool. When posting a job, employers can specify on their post whether they are seeking a "Expert/Expensive", "Intermediate", or "Cheap/Inexperienced" freelancer. These expertise tiers are used by employers to signal both the difficult of their job and the wage their are willing to pay. As shown in Table 15 job posts specifying "Expert/Expensive" had a significantly higher fill rate in treatment vs. control, whereas there seems to be a negative effect on job posts specifying "Cheap/Inexperienced" though this is more noisily measured.

Table 15: Heterogenous Treatment Effect by Expertise Tier

	Dependent variable:
	filled
Treated	-0.012
	(0.009)
Expert/Expensive	-0.110^{***}
	(0.008)
Intermediate	-0.047^{***}
	(0.007)
Unspecified	-0.124^{***}
	(0.017)
Treated x Expert/Expensive	0.028^{**}
	(0.012)
Treated x Intermediate	0.010
	(0.010)
Treated x Unspecified	0.022
	(0.024)
Observations	65,311
R^2	0.048
Adjusted R^2	0.048

Notes: Significance Indicators: $p \leq 0.1 : *, p \leq 0.05 : **, p \leq 0.01 : ***$ Omitted level is "Cheap/Inexperienced". Regression includes fixed effects for job category, whether the employer had previously posted a job, whether the job was fixed price or hourly wage, and whether the job is fixed price or hourly wage. Standard errors are clustered at the employer level.

Click-stream data from the platform allows me to see the exact filters selected by each treated employer in each job. Table 16 shows the ten most commonly selected filters. In line treated employers inviting and hiring workers with higher reputation on the platform, Job Success Score is the second most used filter.

Table 16: Top 10 filters used

Filter	Percentage Selected
Category	74.3%
Job Success	69.4%
English Level	36.6%
Hourly Rate	30.3%
Hours Billed	28.3%
Earned Amount	24.7%
Location	21.6%
Freelancer Type	19.1%
Last Activity	15.7%
Talent Clouds	13.7%

Unfortunately there is no such data for employers in the treatment group, as they were unable to use the filters. Though I control for some job and employer characteristics, since the analysis below only includes treated employers, it is only correlational. There may be selection effects unobservable to the researcher or the platform, for example, that are correlated with both the choice of filter and hiring. Nevertheless, query data provides some suggestive evidence that employers who used the screening tool to search for higher quality freelancers had higher fill rate. Table 17 shows that employers who used the highest option in Job Success filter were substantially more likely to fill their job. Similarly, as shown in Table 18 employers who choose to filter out freelancers with lowest asking rates had more success hiring.

Table 17: NSS Facets vs Fill Rate, Used Facets Openings Only

	Dependent variable:
	Filled
Job Success 90 or higher	0.052^{***}
	(0.013)
Job Success 80 or higher	0.014
	(0.023)
Observations	10,369
R^2	0.045
Adjusted R^2	0.044
Residual Std. Error	0.481 (df = 10352)
Note:	*p<0.1; **p<0.05; ***p<0.01

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$ All regressions include the fixed effects for requested freelancer expertise tier, job category, and whether employer is new to the platform. Standard errors are clustered at the employer level. Omitted level is the default option "Any NSS". 435 job posts at least 1 query with ≥ 90 selected and 1 query with ≥ 80 selected; these are included in "90 or higher".

Table 18: Rate Facets vs Fill Rate, Used Facets Ops Only

	Dependent variable:
	Filled
10 or higher	0.030^{**}
	(0.013)
Observations	7,679
R^2	0.040
Adjusted R^2	0.038
Residual Std. Error	0.486 (df = 7663)

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$ All regressions include the fixed effects for requested freelancer expertise tier, job category, and whether employer is new to the platform. Standard errors are clustered at the employer level.

4 Discussion

The introduction of the new screening tool shifted employer attention and hiring from organic applicants to invited applicants. Within the invitation pool, treated employers invited and hired a higher percentage of freelancers with high platform reputation, but also at higher pay. This suggests that the new screening tool primarily benefited freelancers who are already successful on the platform. Since overall probability to hire was unaffected, this benefit came at a 1-1 cost to less successful freelancers.

This substitution may be a concern to welfare distribution, as invitations are highly concentrated among a minority of freelancers on the platform. Figure 7 shows the distribution of invitations received vs. organic applications sent among the freelancers in the experiment sample. Freelancers are included if they received at least one invitation or sent at least one organic application to a job post included in the experiment sample. Figure 8 shows the distribution of hires via the two channels for the same sample of freelancers. Both figures show a much higher concentration in the invitation channel as compared to organic applications.

There is also evidence that freelancers adjust their hourly rate bids upwards when they are invited. Given that an invitation is a strong signal of employer interest, this is consistent with optimal bidding in models of multiple attribute auctions. Regressions 1 and 2 in Table 19 explores how a freelancer i bids on job j when they are invited vs when they apply organically:

$$\text{Bid}_{ij} = \alpha_i + \beta \text{Invited}_{i,j} + \gamma' \mathbf{X}_j + \epsilon_{i,j}$$

Regression 1 shows that a freelance bids \$1 more per hour when invited - this is a sizable difference even in the traditional labor market. Regression 2 is the same analysis at the contract level and shows that the higher bid translates to a higher contract rate. Regressions 1 and 2 include freelancer fixed effects, but a freelancer may increase their rates over time as they accumulate platform experience. To control for this, regressions 3 and 4 has as dependent variable the difference between the freelancer's bid and their stated asking rate on their profile at the time of the bid. The estimates are similar across the two sets of regressions.

Despite the benefits shown by the experiment, substitution to invitations can have long run adverse effects for the platform. One such adverse effect is the potential for congestion among freelancers with the highest platform reputation. Another is the potential for unraveling in the organic application channel - as the probability of being hired via organic application decreases, fewer qualified freelancers would be willing to use this channel. This may negatively impact the clients that solely use the organic application channel. New freelancers may find it more difficult to start working, as without platform reputation, they primarily rely on the organic application channel.

Thus substituting to the invitation channel results in two clear patterns of redistribution. First, work is reallocated from freelancers with lower platform reputation to freelancers with higher platform reputation; a rich gets richer effect. Second, invited freelancers are able to extract higher pay from employers. The long-run optimal allocation of work on the platform is beyond the scope of this paper, but an interesting and important avenue to pursue. Turning to the traditional labor market, where it is increasingly easy for recruiters to directly reach out to workers via services like LinkedIn, we might expect a similar shift in employer attention and offers towards workers with the highest and most visible reputation. Not only might this limit the promise of online job sites to lower search frictions, but it may also lead to an increase in inequality of opportunity within the labor market.

Figure 7: Lorenze Curve for Invites Received vs Organic Applications Sent

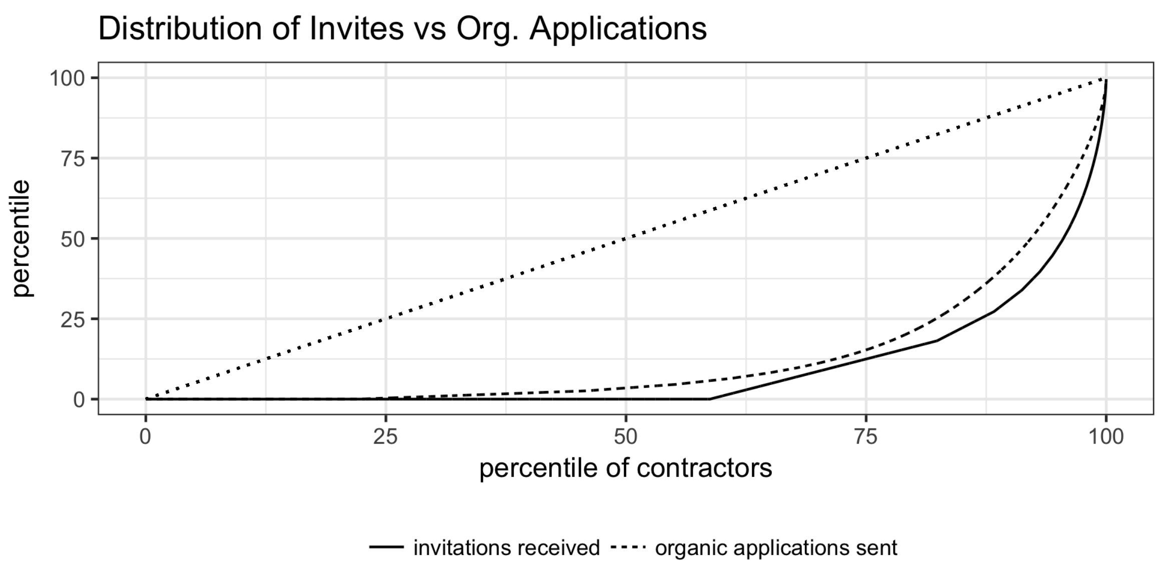

Table 19: Bids: Invite vs Organic Application

	Dependent variable:			
	Hourly Bid		(Hourly Bid - Profile Rate)	
	All	Hired	All	Hired
	(1)	(2)	(3)	(4)
Invited	1.072***	0.729	1.296***	1.549***
	(0.046)	(0.499)	(0.046)	(0.142)
Difference from Profile Rate	No	No	Yes	Yes
Median Hourly Rate	16	16	16	16
Freelancer FE	Yes	Yes	No	No
Observations	774,353	11,974	774,353	11,974
Adjusted R^2	0.920	0.921	0.009	0.020

Notes: Significance Indicators: $p \leq 0.1 : *, p \leq 0.05 : **, p \leq 0.01 : ***$
Hourly rate job posts only. All regressions include the fixed effects for requested freelancer expertise tier, job category, and whether employer is new to the platform. All standard errors are clustered at the employer level. To better account for time varying contractor specific confounding factors, the dependent variable in regressions 3 and 4 is the hourly rate of the bid minus the contractor's stated hourly rate on their profile at the time of the bid.

Figure 8: Lorenze Curve for Hires via Invitations vs Organic Applications

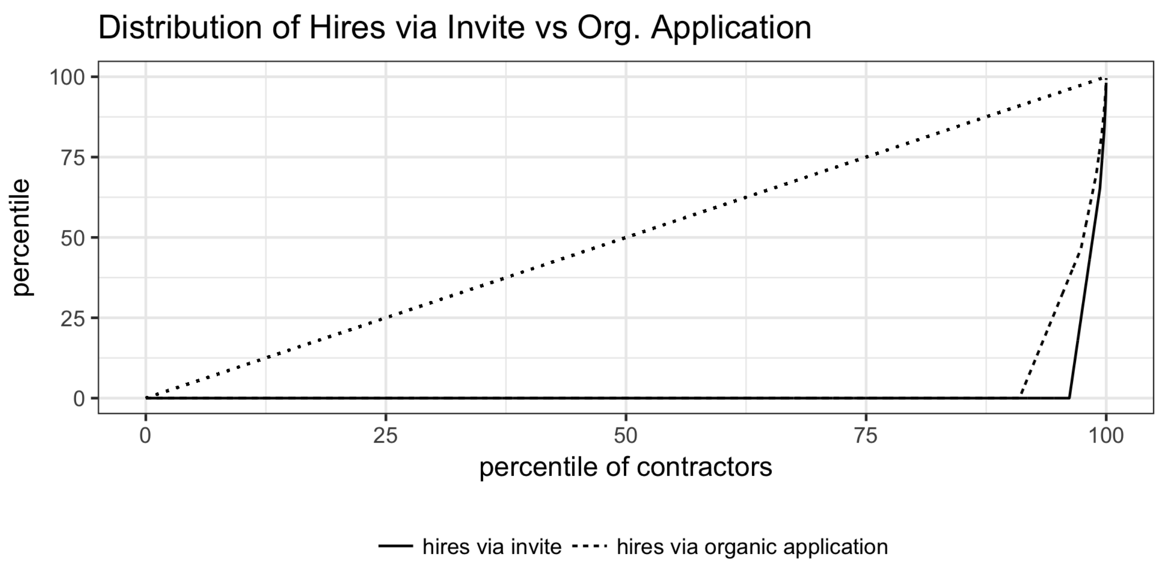

Appendix: Results for Full Sample

Table 20: Effect on Recruitment Intensity and Yield: All

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
Invites	0.28 (0.038)***	10.39
Extensive Margin	0.00 (0.003)	0.40
Intensive Margin	0.57 (0.064)***	8.63
Accepted Invites	0.11 (0.014)***	12.68
Probability Invite Accepted	0.00 (0.002)	0.93

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$ Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.

Table 21: Effect on Hires: All

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
Mean Hires	0.00 (0.004)	0.33
Mean Invite Hires	0.01 (0.002)***	5.28
Mean Org Hires	-0.01 (0.004)	-0.97
Probability of Hire	0.00 (0.003)	0.40
Probability of Invite Hire	0.01 (0.002)***	5.96
Probability of Org Hire	-0.01 (0.002)**	-1.17

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$ Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.

Table 22: Interviews of Organic Applicants: All

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
Organic Interviews	-0.03 (0.016)**	-1.82
Extensive Margin	-0.01 (0.002)***	-1.52
Intensive Margin	-0.01 (0.023)	-0.26

Notes: Significance Indicators: $p \leq 0.1 : *$, $p \leq 0.05 : **$, $p \leq 0.01 : ***$. Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.

Table 23: Percentage of contractors with $\geq 90\%$ Job Success: All

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
Invites	0.04 (0.004)***	5.48
Accepted Invites	0.03 (0.003)***	3.70
Hired Invites	0.02 (0.006)***	3.14
Org Applicants	0.00 (0.001)	0.03
Interviewed Org Applicants	-0.00 (0.003)	-0.59
Hired Org Applicants	-0.00 (0.004)	-0.34

Notes: Significance Indicators: $p \leq 0.1 : *, p \leq 0.05 : **, p \leq 0.01 : ***$. Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.Job Success represents the % of jobs completed with positive feedback. It is only calculated for freelancers with at least 4 distinct employers in last 2 years. Freelancers with no job success are included with the dependent variable set to 0, not dropped.

Table 24: Percentage of completed contracts with a Good Outcome: All

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
All Hires	0.00 (0.004)	0.23
Invite Hires	-0.01 (0.007)	-1.39
Org Hires	0.00 (0.004)	0.65

Notes: Significance Indicators: $p \leq 0.1 : *, p \leq 0.05 : **, p \leq 0.01 : ***$. Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.the platform categorizes a completed contract as having a Good outcome if 1) client provides a feedback score at least 9 out of 10, and 2) there are no complaints or disputes from either side.

Table 25: Percentage of completed contracts with a Bad Outcome: All

	Treatment Effect ($\hat{\beta}$)	$\hat{\beta}$ as % of Control Mean
All Hires	-0.00 (0.002)*	-3.11
Invite Hires	-0.00 (0.004)	-1.04
Org Hires	-0.01 (0.003)*	-3.50

Notes: Significance Indicators: $p \leq 0.1 : *, p \leq 0.05 : **, p \leq 0.01 : * * *$. Regressions include fixed effects for job type (hourly wage or fixed price), requested freelancer expertise tier, and whether employer is new to the platform. Errors are clustered at the employer level.the platform categorizes a completed contract as having a Bad outcome if 1) client provides a feedback score of 6 or below out of 10, and 2) there are complaints or disputes from either side.

Chapter 2: Foot in the Door or Head and Shoulders Above the Rest? How algorithmic recommendations affect employer search and hiring behavior

with Moshe A. Barach and Ming D. Leung

Abstract

How do employers respond to hiring tools? We examine the hiring decision as a process comprised of a series of decisions. We claim hiring indicators act as a "minimal cue" by elevating a job applicant to being noticed early in the hiring process, but giving way to the information an employer privately gathers later in the hiring process. Regression discontinuity analyses of over 1.5 million job applications by freelancers for over 150,000 short-term jobs on an online market for contract labor demonstrate support for our contention. Dramatically, being algorithmically recommended increases a job applicant's unconditional likelihood of being hired over applicants of observationally similar quality by approximately 50%. 70% of this effect can be attributed to the increase in likelihood of recommended job applicants being viewed. Furthermore, conditional on being interviewed, the recommendation has no influence on the hiring decision. Underscoring the idea of it being a 'minimal cue', the effect of a recommendation is stronger for low value jobs than high value ones.

Since at least Spence (1972) and Arrow (1973), we have recognized that hiring decisions are fraught with uncertainty because it is difficult for employers to judge the ability of a job applicant. To remedy this, screening tools have been developed to resolve this uncertainty by providing objective and quantifiable measures of underlying skill. For example, personality tests have been utilized to identify better employees (Autor & Scarborough 2008, Goodstein & Lanyon 1999), or credit reports can be used to determine the qualifications of an applicant (Bos et al 2016, Clifford & Shoag 2016). The evidence generally demonstrates that these objective screening devices have the ability to improve hiring decisions (Wigdor and Green 1991).

Despite the opportunity that screening indicators may present for employers, we lack understanding as to how and when these indicators are used in the 'real world'. The extant work on screening technologies (Autor & Scarborough 2008, Goodstein & Lanyon 1999) provides a limited perspective because, for the most part, they investigate the hiring outcomes when screening indicators are implemented, but do not provide any visibility into how employers respond to them. On one hand, algorithmic recommendations are likely to make superior hiring decisions than purely cognitive ones because computers are not swayed by extraneous cues or subject to biases that human employers have been accused of (Khaneman 2011). Extant work demonstrates that algorithmic recommendations perform "better" than purely cognitive, or human derived, decisions across a range of well-circumscribed tasks (Dawes et al 1989, Grove et al 2000). On the other hand, employers are often unwilling to incorporate such quantifiable indicators of job applicant ability into their hiring decision because they believe they are better judges of a job applicant's suitability for employment (Highhouse 2008). Employers believe they could learn more from"informal discussions" with job applicants (Rynes, Colbert, and Brown 2002) and often ignore quantifiable indicators of applicant quality because they value other, possibly unrelated, aspects of a job applicant's characteristics (Rivera 2012).

Disentangling these mixed perspectives is of particular importance given the increasing prevalence of computational analytics brought to bear on many aspects of labor market hiring. In particular, computational hiring algorithms have been broadly implemented by large, online job search sites, such as Monster, LinkedIn, or Career Builder to match potential job applicants with employer positions. The nascent work on algorithms in hiring has predicted extensive benefits of algorithmic screening tools to increase the likelihood of an employment match on online platforms for contract work (Horton 2015), identify employees who stay longer and perform better in the low-skilled service sector (Hoffman 2016), and pinpointing computer programmers with better leadership potential and cultural fit for a large technology company (Cowgill 2016). In sum, while algorithmic screening tools are likely to dramatically alter employer hiring decisions in the future, for better or worst (O'Neil 2016), we have little understanding as to whether or how people will interact with them.

In this paper, we ask whether and how algorithmic hiring recommendations affect employer hiring. We reconcile the potential conflicting predictions by examining hiring as a process, not merely a single decision, consisting of multiple decisions that employers make on a set of job applicants (Fernandez and Weinberg 1997). We suggest that the algorithmic recommendation exerts different effects depending on which stage of the hiring process one examines. In particular, we predict hiring algorithms provide a "minimal cue" to an employer, much in the spirit of how Arrow (1973) and Stiglitz (1975) conceived of education as a filter in the hiring process. We expect to see that employers will be more likely to hire algorithmically recommended job applicants than applicants who are not recommended. More specifically, we expect the algorithm to exert the most influence early in the hiring process, by altering the job applicants an employer decides to investigate in more detail. Beyond this step in the hiring process, we expect that the algorithmic information will be supplanted by private information employers glean directly from job applicants, as the employer chooses whom to hire. In sum, we believe being algorithmically recommended will help an applicant get their "foot in the door".

Analyses of over 1.5 million job applications from over 200,000 freelancers for over 130,000 jobs posted by more than 81,000 employers on one of the largest online market for contract labor demonstrate support for our contentions. Utilizing a Regression Discontinuity design (Imbens and Kalyanaraman 2012), we reveal that algorithmic hiring recommendations do indeed exert a causal effect by increasing a job applicant's likelihood of being hired by 50% over non-recommended, but observationally similar, job applicants. In support of our contention, we demonstrate that 70% of this effect is isolated to early in the hiring decision process, by increasing the likelihood that an applicant will be investigated further by the employer among an initial pool of job applicants. The effect of the algorithmic recommendation diminishes dramatically and has no effect on being hired conditional on a job applicant being interviewed by an employer. Substantively, this surprisingly large hiring advantage equates to a non-recommended freelancer bidding $1/hr being hired at rates similar to a recommended freelancer bidding $50/hr. A recommended freelancer with 1 past completed job is similarly likely to be hired as a non-recommended freelancer with 100 past completed jobs. Furthermore, we find that the effect of being recommended is greater when employers are hiring for low skilled jobs than high skilled ones, thereby corroborating our belief that being recommended is merely a "minimal cue".

This paper makes at least three notable contributes. First, it makes both a theoretical and substantively practical contribution to the labor market literature. While previous theories relied on merely observing hiring outcomes to infer mechanisms that affect employer decision making, they can only provide a partial explanation if, as we suspect, employers respond to cues of an applicant's ability in one part of the hiring decision but not another. For example, this knowledge may alter how we view the audit and correspondence study findings of discrimination scholars who demonstrate bias in the callback rates of minority job applicants (Pager 2007, Tilcsik 2011). A potential criticism of this work is that the outcomes these scholars observe are only an initial step in the hiring process and that discrimination may not be as severe as is inferred from these callback rates. However, our findings suggest that these very early decisions employers make to either include or exclude a job applicant from further consideration have the potential to magnify the disparities we see in employment because any slight preference only need to manifest very early in the hiring process to have a very strong impact on eventual outcomes (Bendick et al 1999). Related, if, as we suspect, employers are utilizing different information to make decisions in different stages of the hiring process, then it becomes theoretically vital for researchers to identify how the individual steps in a hiring process affect the eventual hiring outcome (Fernandez-Mateo and Fernandez 2016). Clearly, one potential reason remediation strategies often do not accomplish their intended consequences (Kalev et al 2006) is due to the fact that scholars have not been able to isolate how employers make hiring decisions in the "real world". Knowledge as to what information is used in which step of the hiring process moves us towards a clearer understanding as to how to correct potential biases as well.

Second, empirically, our regression discontinuity strategy allows us to disentangle the impact of a screening indicator versus a human capital advantage (Spence 1973, Stiglitz 1975). Simply put, job applicants who are recommended may merely be "head and shoulders above the rest". Certainly, other types of external cues, such as an applicant's level of education, are also correlated with actual differences in their abilities. The simple answer to our question could be that the applicants should be more likely to be hired, because they are simply better. Empirical investigations of how an indicators may act independently of actual ability, are difficult to come by. Our empirical strategy effectively allows us to control for actual ability, while isolating the effect of the algorithmic recommendation on actual hiring outcomes.

Finally, our understanding as to how the temporary "gig-economy" market, in general, and skilled freelancing in particular, operates is lacking, despite its rapid recent growth as an employment relationship (Barley and Kunda 2004). A recent investigation revealed that a full 15.8% of the US Labor force consider themselves employed in non-standard work in 2015, up from 10.1% in 2005, which includes skilled temporary contract labor and part-time employment (Katz and Krueger 2016). Other estimates that include individuals who dabble part-time in counts of self-employed freelancers suggest up to 53 Million individuals have attempted to freelance at least part-time (Edelman 2014). Scholars recognize this trend as replacing traditional work relationships with a market-based, on-demand workforce (Cappelli 1999, Leung 2014), and industry analysts and consultants have all identified them as being a critical sector of the job market (Gartside et al 2013), leading the popular press to suggest we are in "the age of the virtual worker".

Algorithms and Hiring Decisions

The general consensus is that algorithmic decision making, which we define as the use of computer machine learning and statistical predictions of historic decisions and outcomes to identify more versus less successful choices, generally perform "better" than purely cognitive, or human derived, decisions across a range of well-circumscribed tasks. The intrusion by computers into the realm of decisions traditionally considered highly nuanced and subjective is exemplified by the use of algorithms to assist employers in making hiring decisions, an increasingly prevalent occurrence. For example, all the major job sites, such as LinkedIn, Monster, and CareerBuilder, use algorithmic recommendations in funneling job candidates to appropriate job openings. These sites also use algorithms to advertise job openings to potential candidates. Social network sites like LinkedIn and Facebook algorithmically recommend recruiters and companies to users looking for a job. This is particularly notable given that a full 84% of Americans who have actively looked for a job in 2014 and 2015, applied for a job online.

Extant investigations have demonstrated that algorithms can successfully identify better versus worst job candidates. A meta-analysis of twenty-five samples across seventeen studies demonstrated that mechanical methods of predicting advancement at work, supervisor ratings, and training performance outperformed human expert judgment (Kuncel 2013). While mechanistic predictions are expected to work well for straightforward tasks, Cowgill (2016) examined how predictive algorithms may fare in a job context that required advanced social skills, such as the needs to work in teams of computer programmers. He finds that algorithmic recommendations continue to outperform human judgments in predicting likelihood passing a face-to-face interview, of accepting a job offer, and on job performance metrics.

Algorithmic prediction are likely to elicit superior results in the hiring domain, for several reasons. First, it is well-known that employers exhibit bias in identifying which applicants are best suited because they are often swayed by cues and signals that have little predictive power (Kahneman et al 1982). For example, Rivera (2012) explained how recruiters for elite jobs, such as investment banking and consulting, were swayed to hire job applicants who were similar to them in terms of the schools they graduated from and even the types of sports or activities they pursued in college even at the risk of hiring less qualified applicants. Even more alarmingly, employers are subject to implicit biases that they may not be necessarily aware of. For example, employers often draw on stereotypes that are associated with the job applicant, in determined how well-suited they may be for a job. Experiments show that people often associate the gender of a job applicant with how appropriate they may be as an employee, with female applicants seen as being more suited to jobs that require nurturing and men applicants being more suited to competitive jobs (Cejka and Eagly 1999).

Second, algorithms have the ability to learn from a breadth of experiences, by incorporating observations from multiple individuals. In effect, algorithms incorporate the wisdom of the crowds by taking advantage of data collected across individuals and over time to identify the best solutions. Conversely, individuals often have to rely on only their limited samples to deduce what factors lead to success or failure. Finally, while algorithms are designed around optimizing an observable outcome; in the real world, decision makers rarely have the opportunity to realize the consequences of their decisions (Einhorn and Hogarth 1978). In a sense, human decision makers cannot avail themselves of counter-factual learning opportunities that algorithms can.

The fact that algorithms take into account information across a broad range of observations as well as possibly improve on the biases that human decision makers make in deciding whom to hire suggests that algorithmic recommendations will provide employers with a hiring cue that they should heed. Broadly, algorithms in hiring can be seen as an innovation in screening technology that provides a source of additional information for the employer. In this sense, an algorithm's recommendation can be considered as providing additional information as to a job candidate's underlying ability, much as a signal or cue has been construed of in the past (Arrow 1973, Steglitz 1995). In doing so, those job applicants who are recommended are likely to be preferred by employers just by dint of the fact they are recommended.

However, the process of hiring is considered a highly subjective decision by most employers. As such, employers are likely to dismiss algorithmic recommendations, instead trusting their own hiring judgments, to their detriment. As Hoffman and his colleagues (2016) demonstrate, when given discretion, employers often ignored the recommendations of a job assessment even though this assessment is effective at identifying employees with lower turnover rates. The fact that employers tend to rely more on their private information gathering, while ignoring algorithmic cues, is particularly worrisome in light of the fact that information gleaned from face-to-face interviews are often extraneous and likely add no predictive power. For example, as DeVaul and colleagues (1987) demonstrated, medical students at the University of Texas at Houston, who were initially rejected largely due to their interview performance, performed as well in measures of attrition and in both pre-clinical and clinical performance through medical school and one year of postgraduate training, as those students who were originally accepted in the program due to their interview performance.

Algorithmic recommendations and the hiring process: A "minimal cue"

Employers are likely to downplay the advice of algorithmic recommendations in light of information they gather themselves. Employers often prefer to rely on their gut instincts in deciding whom to hire and ignore quantifiable indicators of ability, such as test scores (Highhouse 2008). They do so because they believe they will make a better decision regarding who will be a better employee. This so-called "myth of expertise", stems from two, interrelated beliefs that employers hold regarding their ability to seek out the ideal job applicant. First, employers believe that as experts, they are able to spot idiosyncrasies in a job applicant's profile (Jeanneret and Silzer, 1998) which an algorithm would not be able to identify. The canonical example is the applicant with a "broken leg", who would be otherwise disadvantaged from the perspective of an algorithm, that cannot take into account such a rare event, whereby an employer who would be able to glean such details from an applicant through an interview, and therefore take that information into account.

Second, algorithms are accused of being unable to interpret configurations of traits that an expert employer would be able to piece together into a holistic impression (Prien et al 2003). Because employers believe that each candidate is unique, they would rather be able to evaluate the information that a job applicant supplies in light of the other pieces of information. In doing so, individuals often are so focused on developing detailed narratives that they often incorporate extraneous information, preferring to predict less likely scenarios that incorporate more detailed information than simpler, but less detailed explanations (Tversky and Kahneman 1982). This is why employers prefer that job applicants who are able to provide a coherent narrative in describing their past work experiences and how it may apply to the job at hand (Bruner 1991).

Given this, we expect that employers will not be willing to allow algorithmic hiring mechanisms to trump information they may glean for themselves directly from the applicant. Instead, we believe a hiring algorithm acts as a "minimal cue" to an employer. By minimal cue, we mean that employers would be reasonably swayed by algorithms but continue to retain discretion in the hiring process. One way this could manifest is that algorithmic assessments assist an employer in identifying better versus worst job candidates initially to screen in, but, when the decision becomes one of selecting which job applicant to hire, employers will likely rely on their personal judgment, gleaned from direct contact with the applicant, over and above any external indicators such as a recommendation. This is because, "subjective impressions of candidates that employers develop through interviews are strong drivers of hiring decisions, often carrying more weight than candidates' resume qualifications (Rivera 2012: 1000).

We identify how algorithmic hiring recommendations affect an employer's hiring decision by conceptualizing the hiring decision into a multi-step process and identifying where in the process a recommendation is expected to have more versus less influence. Hiring a job applicant is not merely a single decision, but rather a series of decisions an employer engages in that successively winnows down the full set of job applicants until a person, or people are, is offered a job (Fernandez and Weinberg 1997). This process is best illustrated as a funnel and can be characterized as progressing along a continuum along two interrelated trajectories. First, as the hiring process proceeds along the funnel, there is an increasingly smaller number of applicants being considered by the employer. This occurs as the employer reduces the number of job applicants whom they feel are worthy of additional consideration. Second, as the hiring process continues along the funnel, there is also a markedly increasing amount of detailed information an employer gathers regarding those job applicants. As the pool of potential applicants is winnowed, employers gather more detailed information regarding the suitability of the applicants. This information is used by the employer to eventually decides whom to hire from the reduced the pool of applicants. See Figure 1 for an example of an employer hiring decision process. Note, employers may vary in the number and distinctiveness of the highlighted stages, however, this conception is generally corroborated by the scant work that has able to observe such details (Fernandez and Weinberg 1997, Fernandez-Mateo and Fernandez 2016).

Figure 1: Illustrative Applicant Pool to Job Offer Hiring Process

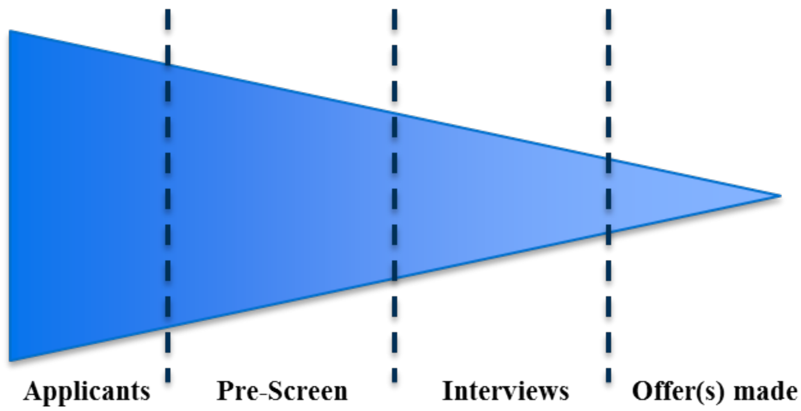

Beginning with a set of job applicants, an employer's first decision is to choose which applicants, to review. This initial screening step may be accomplished by simple heuristics, such as the level of education or some minimal qualification that the applicants are expected to possess (Bills 1990). The employer then reviews this subset by looking closely at an applicant's qualifications or specific relevant job experiences. Among this further narrowed set of applicants, the employer may choose whom to invest more time by interviewing them or requesting additional information. Finally, armed with this privately acquired information, an employer will select whom to hire.

Our expectation is that algorithmic hiring recommendations will likely exert the most influence early in the hiring decision, by increasing the likelihood that a job applicant will be more seriously considered. Concurrently, as an employer gathers first-hand information regarding the job applicants, algorithmic recommendations will exert a declining influence while individuating information from the job applicant and an employer's discretion begins to exert greater influence.

Algorithms are therefore useful in assisting an employer in quickly narrowing down the choice set that they will then invest time in choosing among. This helps because search is inherently costly, and when search is made more complex, it often works to demotivate us by leading us to make poor choices or deciding not to choose at all (Lepper and Iyengar and Lepper 2009, Leung 2017). For an employer to exhaustively investigate a complete pool of job applicants would not only require time and effort, but may also render the decision too complex as the greater the number of options, the greater the number of characteristics one will have to consider (for detailed review, see Chernev et al 2015). A simpler strategy to quickly narrow down the set of available options is for an employer to identify a very small number of dimensions, such as whether one is recommended or not, to compare the job applicants. These metrics will be used to quickly include or exclude job applicants for further consideration.

There are at least three scope conditions worth highlighting. First, the employer is not subject to any organizational specific hiring rules, such as policies that require the use of algorithms. This would necessarily lead to employers having limited discretion in how they hire and therefore not necessarily reflect how algorithms may be viewed "in the real world". Second, the full set of job applicants consist of both algorithmically recommended applicants as well as non-recommended applicants and are all visible to the employer. If a hiring algorithm limits an employer's ability to even view applicants, perhaps by automatically removing them from a job applicant pool, then an employer would not even have the opportunity to exercise their discretion. Third, the employer has a reasonable number of job applicants from which to hire. Because we believe the algorithm is most useful in assisting an employer winnow down their choices, it would not be as useful for jobs with very few applicants as the employer would expend very little cost in examining all of them.

1 Empirical Setting

We investigate how employers use algorithmic recommendations in a large online labor market platform for the hiring of gig-economy freelancers, similar to oDesk, Upwork, and Freelancer.com. As of the end of 2015, millions of employers and freelancers have created profiles. In 2015, employers spent over $1 billion on wages through the platform. Employers on the platform are primarily individuals hiring for small firms or small teams within medium to large firms. In our data 97% of employer accounts have only one associated User ID, so only a single individual is involved in the hiring process. These single user employer accounts are associated with 80% of all job posts in our data.

Job postings are organized by job categories that represent the spectrum of business tasks that can be accomplished virtually. In general, these are office tasks that are completed individually. The median job lasts about 10 hours and costs about $75. Based on dollars spent, the top skills in the marketplace are technical skills, such as web programming, mobile applications development (e.g., iPhone and Android), and web design. Based on hours worked, the top skills are web programming, data entry, search engine optimization, and web research. The top five countries for employers are: the United States, the United Kingdom, France, Germany, and Israel. Top top five countries for workers are: the United States, India, Philippines, Russia, and Ukraine. Transacting on the platform

Posting a job

The process for filling a job opening is qualitatively similar to the process in conventional labor markets. First, a would-be employer on the platform creates a public job posting. Public jobs can be seen by all workers on the platform. Employers choose a job title and describe the nature of the work. Additionally, employers choose a contractual form (hourly or fixed-price), specify what skills the project requires (both by listing skills and by choosing a category from a mutually exclusive list), and estimate how long the project is likely to last. Once the job posting is written, it is reviewed by the platform and then posted to the marketplace. All the information provided by employers is viewable by potential applicants. Additionally, the platform also presents verified attributes of the employer, such as their number of past jobs and their average wage rate paid. Figure 2 shows how a typical job posting might appear once it is posted to the platform.

Figure 2: Job Post for Java Backend Developer

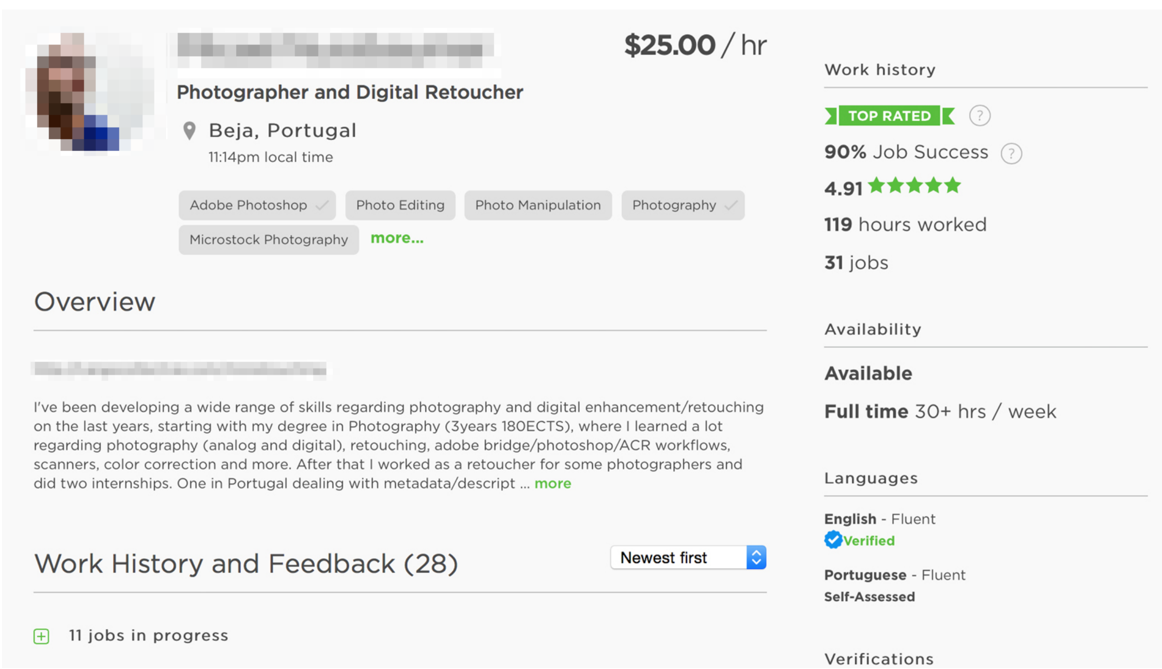

Applying to a job

Applicants can apply to any public job posting on the platform. When they apply, they include a bid (the amount they are willing to work for), and a cover letter, which consists of a paragraph of text meant to convince the employer that the applicant is "right" for the job. After applying, the applicant immediately appears in the employer's applicant tracking system or ATS, which is the dashboard the employer sees upon logging into his account and clicking on an open job posting. Each application in the ATS shows the applicant's' name, picture, bid, self-reported skills, country, and a few pieces of platform-verified information, such as total hours-worked and average feedback rating from previous projects (if any). Figure 3 shows the employers' view of the applicant tracking system (ATS).

Figure 3: The Application Tracking System (ATS)

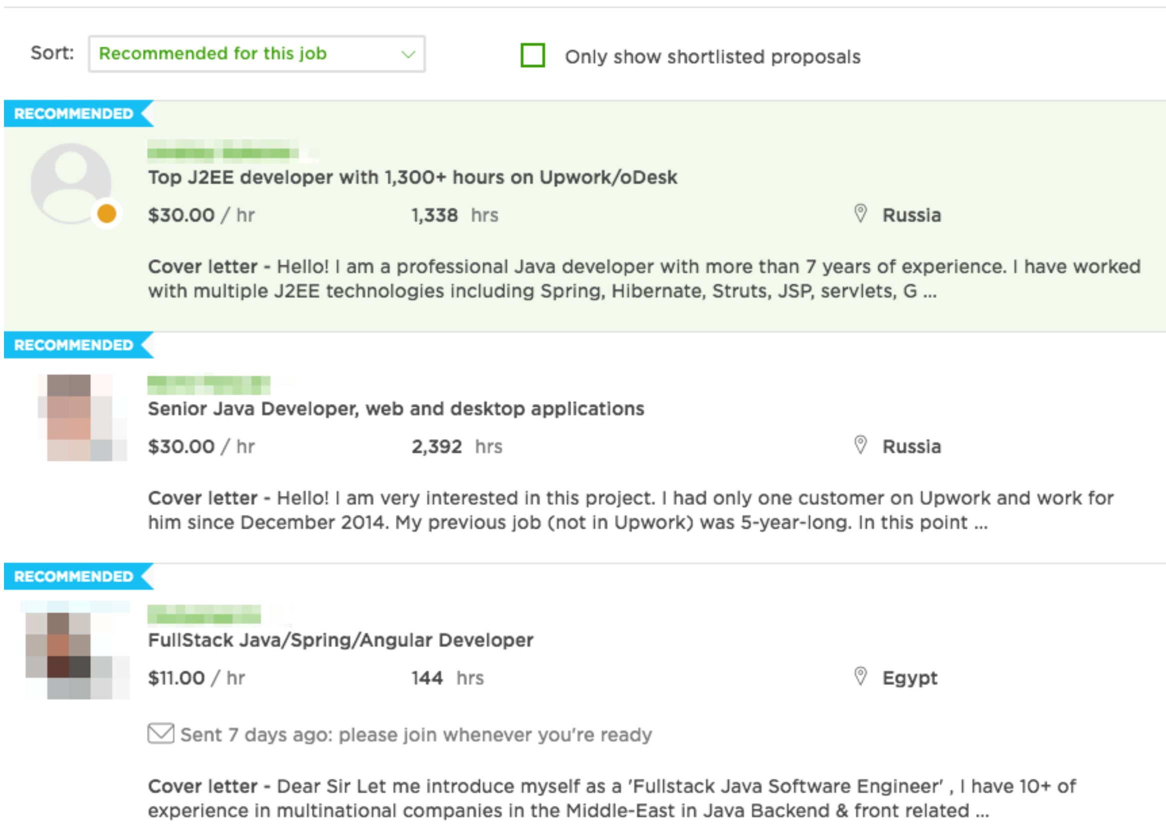

Algorithmic Recommendation

For each applicant to a job, the platform uses a proprietary algorithm to identify those who are recommended to the employer. This machine learned algorithm predicts whether an applicant will be hired and if so, whether the job will be completed to the client's satisfaction. The algorithm uses a variety of indicators, such as the applicant's past employment, current profile characteristics, and the job's stated specifications.

The algorithm assesses each applicant and scores him or her on a zero to one scale. While the exact score of the algorithm is never made public, All job applicants who score a 0.50 or above on this algorithm at the time of their application are automatically recommended by the platform, resulting in a "recommended" flag being placed above their picture. Additionally, applicants are default sorted by algorithm score, placing that applicant at the top of the list an employer views. Figure 3 shows how applicants both those marked as recommended as well as those not marked as recommended appear to employers in the ATS. Applicants are never informed by the platform of their ranking or whether or not they were recommended. From frequent discussions on the platform's message board, it is clear that freelancers are aware of the recommendation algorithm but do not know enough detail to game the algorithm. A situation that the platform actively works to maintain by not revealing how the algorithm is calculated.

Viewing Applications & Interviewing Applicants

An employer is able to learn more about a specific applicant and view that applicant's detailed application by clicking on any of the job proposals that which are listed in the ATS. Viewing an applicant's detailed application is similar to reading an applicant resume. While it is possible for an applicant to be hired without the employer reviewing the detailed application, this is extremely rare. In our data only 1.8% of hired applicants are hired without being viewed by the employer. An employer viewing a detailed application would see more information about that worker, including the workers self-written overview as well as the worker's past work history. Figure 4 shows how a detailed application would look to an employer on the platform. For example, we can see in figure 4 that applicant self-categorizes himself as a Photographer and Digital Retoucher, and lists that he he is skilled in: Adobe Photoshop, Photo Editing, as well as a few extra skills. Employers can also review applicants work history and feedback on this page. This work history contains an item for each job the applicant has completed on the platform and contains both the review and feedback left by the applicant's previous employer on that job.

Figure 4: The Detailed Application

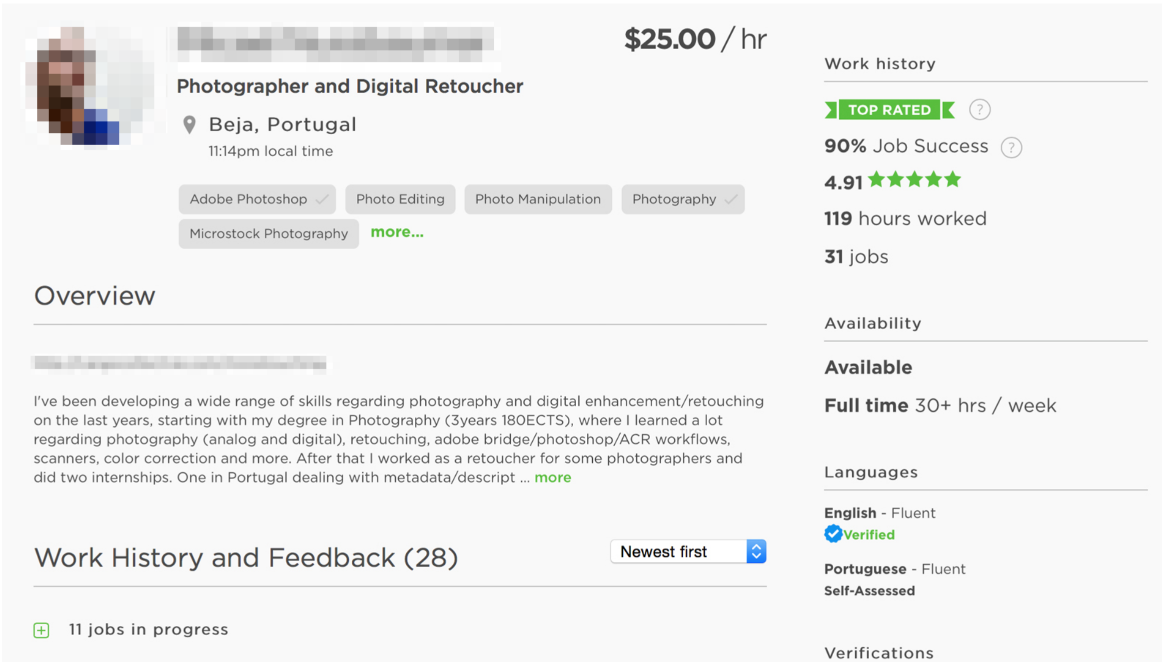

After viewing an application, the employer can either message the applicant to conduct an “interview” or directly hire the applicant. Employers are encouraged by the website to interview their applicants and about 70% of hired applicants are interviewed. Figure 5 shows the prompt through which employers can initiate interviewing of a candidate via messages.

Figure 5: Internal Messaging System

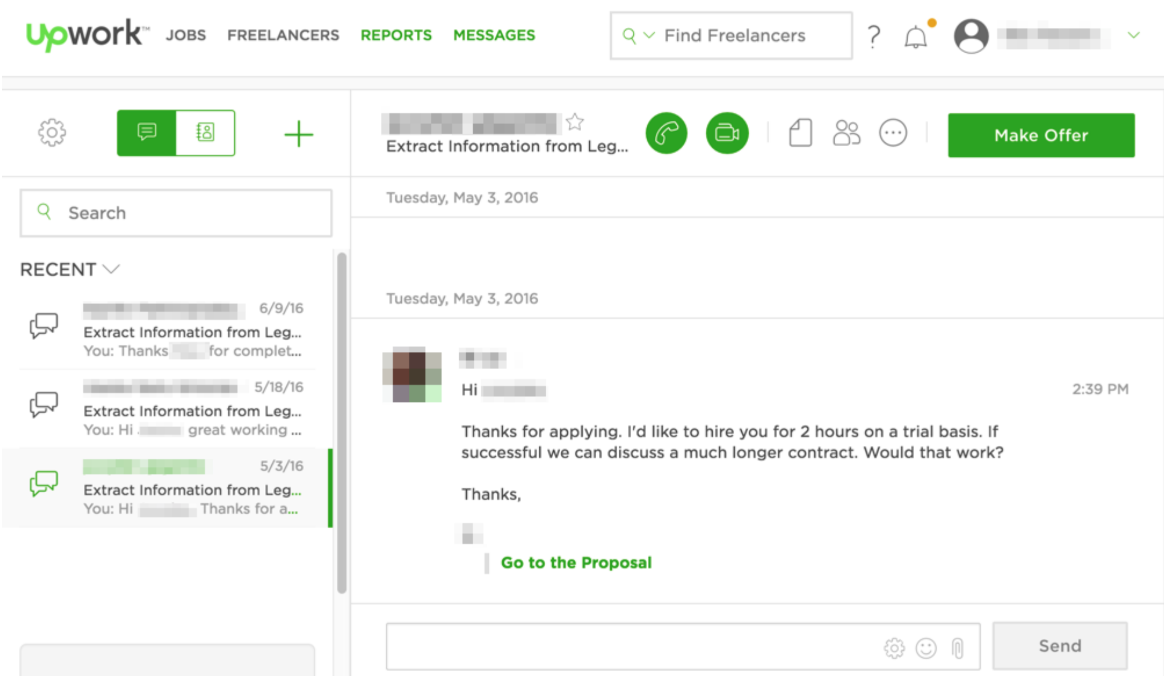

Hiring Applicants

An employer is free to hire whomever they wish. The employer hires the worker on the terms proposed by the worker or make a counteroffer, which the worker can accept, reject, or negotiate. On average only 43% of job posting are filled. If an employer chooses to hire anyone, 90% of the time they hire only one freelancer, though employers are able to hire more. Once a freelancer is hired, employer and employee exchange job details and the job is completed virtually. Payment is conducted through the website.

Data and Variables

We study a random sample of job postings submitted to the platform between March 3, 2016 and May 30, 2016. We limit our sample to only non-invited applicants, as applicants who are invited to the job might be previously known to the employer. We additionally winsorize our sample by dropping job postings which are above the 99th percentile with respect to either number of applications or number of applicants hired per job. Our remaining sample consists of 1,574,204 applications from 211,620 unique freelancers, submitted to 136,548 unique job postings posted by 81,314 unique employers. Table 1 presents summary statistics for all applicants to jobs in our sample.

Table 1

Opening Summary Stats, Organic Applications Only (n = 136,548)

Statistic	Mean	St. Dev.	Pctl(25)	Median	Pctl(75)	Max
Applicants	20.75	21.24	6	14	27	137
Recommended	6.44	7.54	2	4	9	86
Viewed	5.81	8.85	1	3	7	128
Interviewed	1.67	3.15	0	1	2	112
Hired	0.44	0.77	0	0	1	35

Notes: This table reports summary statistics on the non-invited applicant pool. All reports are on a per-opening basis. An application is "recommended" if the applicant is marked by the ML algorithum as a recommended applicant. An application is "viewed" if the employer clicked on a worker's application to learn more about the applicant. An applicant is "interviewed" if the employer sent a message to the applicant.

2 Empirical Strategy

We seek to understand how an applicant being recommended by the algorithm affects employer's hiring decisions on the platform. Unfortunately, we cannot simply compare selection rates of recommended and non-recommended applicants, as these applicants differ in quality in ways which might not be observable leading to omitted variable bias. To eliminate such problems, we take advantage of the fact that job applicants are unable to precisely manipulate their algorithm score. As such, the variation in treatment - being recommended by the hiring algorithm - near the recommendation threshold can be thought of as good as randomized. See Appendix C for tests which show that applicant's are unable to manipulate whether or not they are recommended. Taking advantage of this institutional detail allow for us to use a regression discontinuity design to control for heterogeneous quality of applicants. To get a better idea of the amount of variation we have close to the recommendation discontinuity, Table 2 presents summary statistics for what we call our discontinuity sample, which is limited to applications which are close to the algorithmic threshold (between .45 and .55).

Table 2

Opening Summary Stats, Organic Applications Only with Algorithm Score in [0.45, 0.55] (n = 117,247)

Statistic	Mean	St. Dev.	Pctl(25)	Median	Pctl(75)	Max
Applicants	6.05	6.54	2	4	8	73
Recommended	2.36	2.89	1	1	3	36
Viewed	1.77	2.84	0	1	2	56
Interviewed	0.50	1.10	0	0	1	38
Hired	0.13	0.38	0	0	0	14

Notes: This table reports summary statistics on the non-invited applicant pool. Applicantions with algorithum scores greater than 0.55 and less than 0.45 are removed. All reports are on a per-opening basis. An application is "recommended" if the applicant is marked by the ML algorithum as a recommended applicant. An application is "viewed" if the employer clicked on a worker's application to learn more about the applicant. An applicant is "interviewed" if the employer sent a message to the applicant.

Graphical Strategy

We begin with a graphical analysis of the effects of an applicant being recommended by the hiring algorithm on the employer's likelihood of selecting that applicant at various stages in the hiring funnel (Lee and Lemieux 2010). These results help to visualize the relationship between employer choices along the hiring funnel and the rating of applicants by the hiring algorithm as well as indicating the magnitude of the effect of being recommended. Figure 6 plots the average probability of an application being viewed, interviewed or hired by the application's algorithm score. The average value of each outcome was calculated for bins of width of .002 on either side of the .5 recommendation discontinuity for algorithm scores ranging between .45 and .55. The binwidth was chosen to clearly demonstrate the effect of the recommendation discontinuity without over-smoothing the data. Alternative bin widths do not materially alter the observed relationships.

Figure 6: Regression Discontinuity Plots

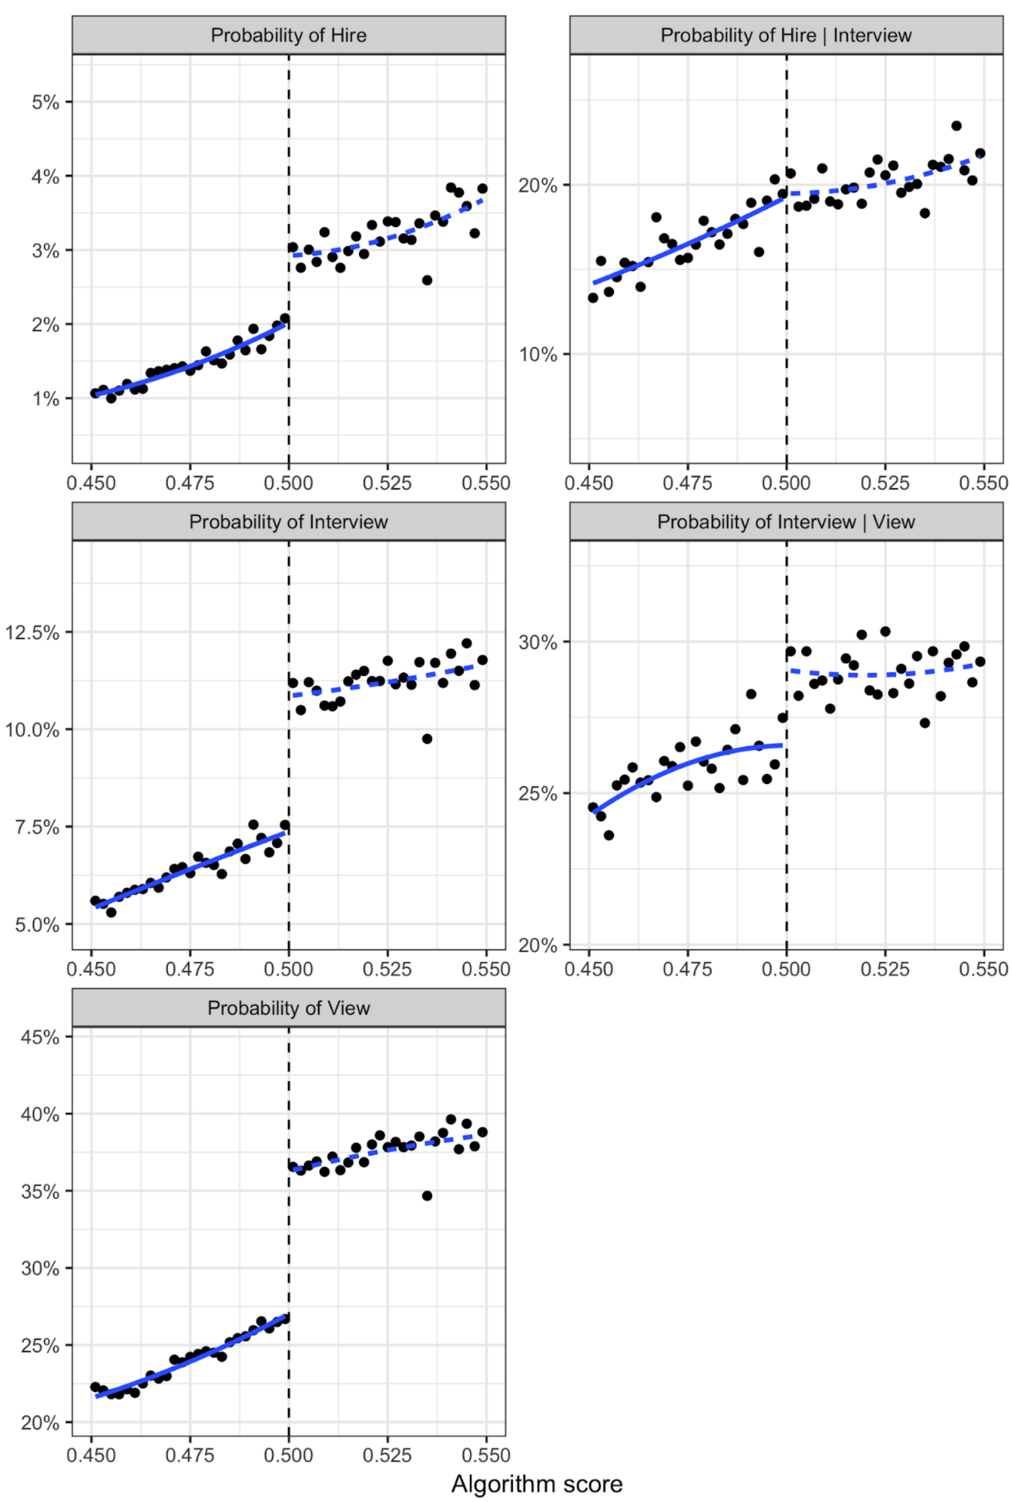

Scatter plot is means for bins of size 0.02. Fitted line is from separate quadratic fits on either side of cutoff.

Regression Strategy

We additionally report regression results which estimate the average treatment effect of the recommendation on employer decisions for applicants with hiring algorithm scores close to the recommendation discontinuity. These results allow us to precisely measure the impact of the algorithmic recommendation on various employer decisions, As we study an online marketplace, where data is plentiful, we report results from a local linear regression. We minimized concerns that the functional form of the relationship between the algorithm score of each application and the outcome of interest is correctly specified by using data in a very narrow bandwidth around the discontinuity.

3 Empirical Results

We examined each step along the hiring funnel as we view the hiring decision as a process consisting of a series of choices. We begin by analyzing the employer's first choice, how to narrow the pool of applicants after the job is posted. Employers must choose which applicant are worth their time and then view those applicant's detailed applications. We then proceed to the decision the employer faces over which applicants the employer chooses to interview using the the platform's internal messaging system. We then investigate the employer's final decision, which applicant(s) should they hire.

Effects on the decisions along the hiring funnel

Table 3 displays results of our preferred regression model, which is a local linear regression allowing for slopes to vary on either side of the cutoff (Hahn, Todd and van der Klaauw 2001). We calculate the bandwidth used in each model using the Imbens-Kalyanaraman Optimal Bandwidth Calculation procedure (Imbens and Kalyanaraman, 2009). All regression results in Table 3 are derived from a model of form:

$$Y_{ij} = b_1 \mathrm{Rec}_{ij} + b_2\, 1[Rec_{ij} = 1]\, \mathrm{score}_{ij} + b_3 1[Rec_{ij} = 0]\, \mathrm{score}_{ij} + X_{ij} + \gamma_j + e_{ij}$$

Table 3

Local Linear RD estimates of the effect of being Recommended

	Dependent variable:				
	Applicant Viewed	Applicant Interviewed		Applicant Hired	
	(1)	(2)	(3)	(4)	(5)
Recommended	0.092^{***}	0.029^{***}	0.013^{**}	0.010^{***}	-0.011
	(0.002)	(0.002)	(0.006)	(0.001)	(0.013)
Norm. Score (< 0)	0.971^{***}	0.486^{***}	0.848^{***}	0.109^{***}	0.393
	(0.046)	(0.052)	(0.188)	(0.011)	(0.308)
Norm. Score (<= 0)	0.649^{***}	0.378^{***}	0.364^{*}	0.211^{***}	0.324
	(0.062)	(0.076)	(0.198)	(0.018)	(0.319)
Sample	All Applicants	All Applicants	Viewed Applicants	All Applicants	Interviewed Applicants
Job-level FE	Yes	Yes	Yes	Yes	Yes
I-K Optimal Bandwidth	0.048	0.034	0.037	0.061	0.048
Dep. Var. Mean	0.292	0.084	0.275	0.021	0.185
Baseline	0.267	0.075	0.275	0.021	0.195
Observations	689,329	483,507	157,329	842,672	56,805

Notes: The bandwidth used in each model was calculated using the Imbens-Kalyanaraman Optimal Bandwidth Calculation. The estimations are from a local-linear model which allows for slope differences on either side of the cutoff. Covariates include applicant's bid on the job posting and within job algorithm score rank fixed effects. Baseline is dependent variable mean for applications within [-0.02, 0) in normalized algorithm score. Model (1) the outcome is an indicator if the applicant is viewed. Model (2) the outcome is an indicator if the applicant is interviewed. Model (3) the outcome is an indicator if the applicant is interviewed conditional on the applicant being viewed. Model (4) the outcome is an indicator if the applicant is hired. Model (5) the outcome is an indicator if the applicant is hired conditional on the applicant being interviewed. Heteroskedasticity-robust standard errors clustered at the job posting level are reported. Significance indicators: $p \le 0.05 : *, .p \le 0.01 : **$ and $p \le .001 : ***$.

These regressions use only data within a small bandwidth around the cut-point. In Column (1), the outcome of interest, Y_{ij} is an indicator that applicant i to job opening j was viewed by the employer. Recommended(Rec) is an indicator which is equal to 1 if the applicant was recommended, score is the applicant's private algorithm score, and X_{ij} is a vector of controls which include: the rank order of the application on the page, the log of the applicant's hourly wage. γ_j is a job-opening fixed effect. This job-opening fixed effect ensures that we only make use of variation between applicants within a job posting, and that our results are not driven by heterogeneous employers or heterogeneous job postings.

The outcome of interest in Model (2) is the probability that the applicant is interviewed. The outcome of interest in Model (3) is the probability that a viewed applicant is interviewed. The outcome of interest in Model (4) is the probability that an applicant is hired. The outcome of interest in Model (5) is the probability that an interviewed applicant is hired.

Probability of Viewing Applications

We begin at the top of the hiring funnel and first seek to understand to what extent employers rely on the recommendation in deciding whether to view an applicant's detailed application.The bottom left panel of Figure 6 shows the probability of viewing an applicant is increasing in an applicant's algorithm score; better applications, as rated by the algorithm, are more likely to be viewed by employers. Additionally, there appears to be a large discontinuity in the probability of viewing an applicant across the recommendation threshold. This indicates that an employer is much more likely to view the detailed application of an applicant who is marked as recommended by the algorithm than that of a similarly skilled applicant who is not not marked as recommended. The regression results in Table 3 confirm these findings and show that for applicants close to the threshold, being marked as recommended by the hiring algorithm makes an applicant 9.2 percentage points more likely to be viewed by an employer or about 34% more likely to be viewed from a baseline of 26%. Employers seem to rely heavily on the recommendation to save time in determining which applicants are worth their further investigation.

Probability of Interviewing Applicants

Once an employer has deemed an applicant worth investigating, and has viewed the applicant's application, the employer must decide which applicants are worth further consideration and should be interviewed via the online messaging system. Column (2) of Table 3, shows that employers are 2.9 percentage points more likely to interview a candidate who is recommended by the machine learning algorithm from a baseline probability of being interviewed of 8.4% for applicants near the recommendation threshold. This is roughly a 34% increase in the probability of being interviewed for a job.

This increase in likelihood of being interviewed could be driven by the algorithm affecting only the employer's choice of which applicants to view, and through that influencing who is interviewed. Additionally, employers could continue to rely on the recommendation even after viewing an applicant's detailed application. In column (3) of Table 3 we limit the sample to only applicant's who were viewed by the employer. After conditioning on applicants who have been viewed by the employer, the effect of the recommendation is substantially smaller, but still positive and significant. The recommendation increases employer's likelihood of interviewing an applicant they have already viewed by 1.3 percentage points from a baseline probability of about 27.5% at the threshold. This translates to an increase in the likelihood of interviewing a viewed applicant of about 4%. Comparing this 4% increase in probably of interviewing to the unconditional increase in probably of interviewing of 34% it is clear that while employers are still affected by the hiring algorithm even after viewing an application, the advice of the algorithm is downplayed in light of the information the employer obtains through viewing an applicant's detailed application.

Probability of Hiring Applicants

From the top left panel of Figure 6 we can see that there is a significant positive relationship between hiring algorithm score and an applicant's likelihood of being hired. Moreover, there is a large discontinuity in the likelihood of hiring an applicant right at recommendation threshold indicating that employers are more likely to hire similar quality recommended applicants than non-recommended applicants. However, the top right panel of Figure 6 displays no discontinuity indicating that recommended interviewed applicants are no more likely to be hired than non-recommended interviewed applicants.

These results are confirmed by Model (4) in Table 3, which indicates that applicants who are marked as recommended by the hiring algorithm are about 1 percentage point more likely to be hired from a baseline probability of being hired of about 2.1% at the discontinuity. This translates to algorithmically recommended applicants being about 50% more likely to be hired than a similar quality applicant just below the recommendation threshold. Model (5) of Table 3, subsets the population to only applicants who are interviewed by the employer, the effect is not statistically different from zero. This indicates that after exerting effort and gathering their own information about an applicant, employers find no value in the algorithmic recommendation. Employers are not willing to allow the hiring algorithm to trump information they glean for themselves.

By providing a "minimal cue" to an employer early in the hiring funnel, the hiring algorithm still has a substantial yet veiled effect on which applicant is hired. Being marked as recommended by the algorithm increases each applicant's raw probability of being hired from 2.1% to 3.1% an increase of 47%. For the 70% of hired applicants, in our data, who were interviewed prior to being hired the probability of being hired is equal to:

$$\Pr[\text{Hired}] = \Pr[\text{Viewed}]\,\Pr[\text{Interviewed}|\text{Viewed}]\Pr[\text{Hired}|\text{Interviewed}]$$

Since this relationship is multiplicative we can attribute the proportion of the overall effect of the recommendation on hiring which comes from the algorithm increasing the probability of an application being viewed early in the funnel as the ratio of the increase of Pr[Viewed] to Pr[Hired]. The probability an applicant is viewed increases from 26.7% to 35.9% an increase of about 35%. Thus, we can attribute 72% (.35/.47) of the algorithm to helping applicants get their foot in the door by increasing their probability of their application being viewed by the employer.

Relative Signal Strength We seek to understand how powerful a signal the recommendation is compared to other commonly used signals which affect applicant selection such as applicant's bid and applicant's past on-platform experience. Figure 7 plots the probability of an employer hiring an applicant by both the applicant's bid and the applicant prior experience for applicants just above and just below the recommendation threshold. To help isolate the signaling value of the recommendation, the sample is limited to only applicants with algorithm scores within the Imbens-Kalyanaraman Optimal Bandwidth of .063 of the recommendation threshold. The effect of the recommendation is so large, that a recommended applicant who bids $50/hour is about just as likely to be hired as a very similar non-recommended applicant who bids only $1/hour. Additionally, a recommended applicant with only 1 prior on-platform job is about equally likely to be hired as a non-recommended applicant with 100 previous on platform jobs completed. The influence of the recommendation on an employer's' final hiring decision is so large that it completely subsumes other available and highly important signals of applicant quality such as applicant bid or number of previous jobs an applicant has completed on the platform.

Figure 7: Effect of Signals of Applicant Value by Recommendation

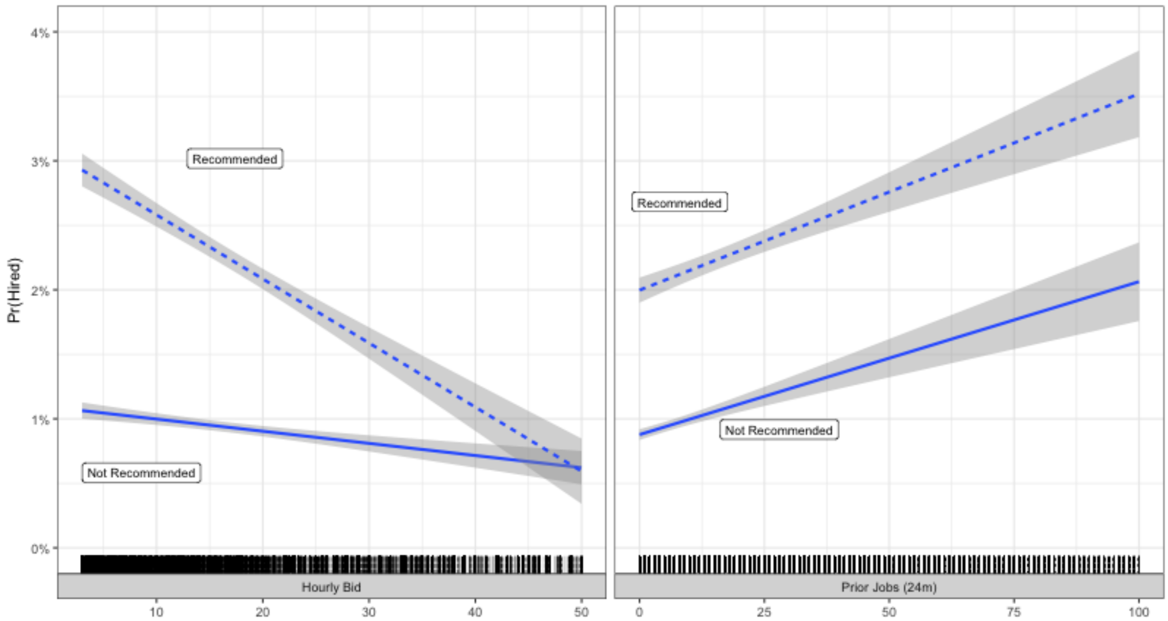

Job Outcomes

As the algorithm has such large effects on the hiring behavior of employers, it is logical to question if the machine learning algorithm leads to differential job outcomes. As we are comparing algorithm recommended applicants to equivalent quality non algorithm recommended applicants we do not expect there to be any measurable effect on the quality of the work product. However, if expectations are lower for non-recommended applicants than recommended applicants it is possible that we might detect differences in job satisfaction between employers who hired recommended applicants who are just above the threshold, and employers who hired non-recommended applicants who are just below the threshold. Figure 8 shows that while there is a large positive relationship between algorithm score and job satisfaction overall, we find no discontinuity across the recommendation threshold in job outcomes, as measured from employer feedback ratings.

Heterogeneous Effects

To further test that the recommendation acts as a 'minimal signal' we examine the possibility of heterogeneous effects of the recommendation based on the value of the job posted by the employer. Because employers are likely to care more about high-value jobs versus low value jobs, we expect them to rely on their personal instincts more for high-value jobs than low-value jobs. We expect employers to invest more of their personal effort in both screening and selecting job applicants for jobs that are more

Figure 8: Effect on Feedback

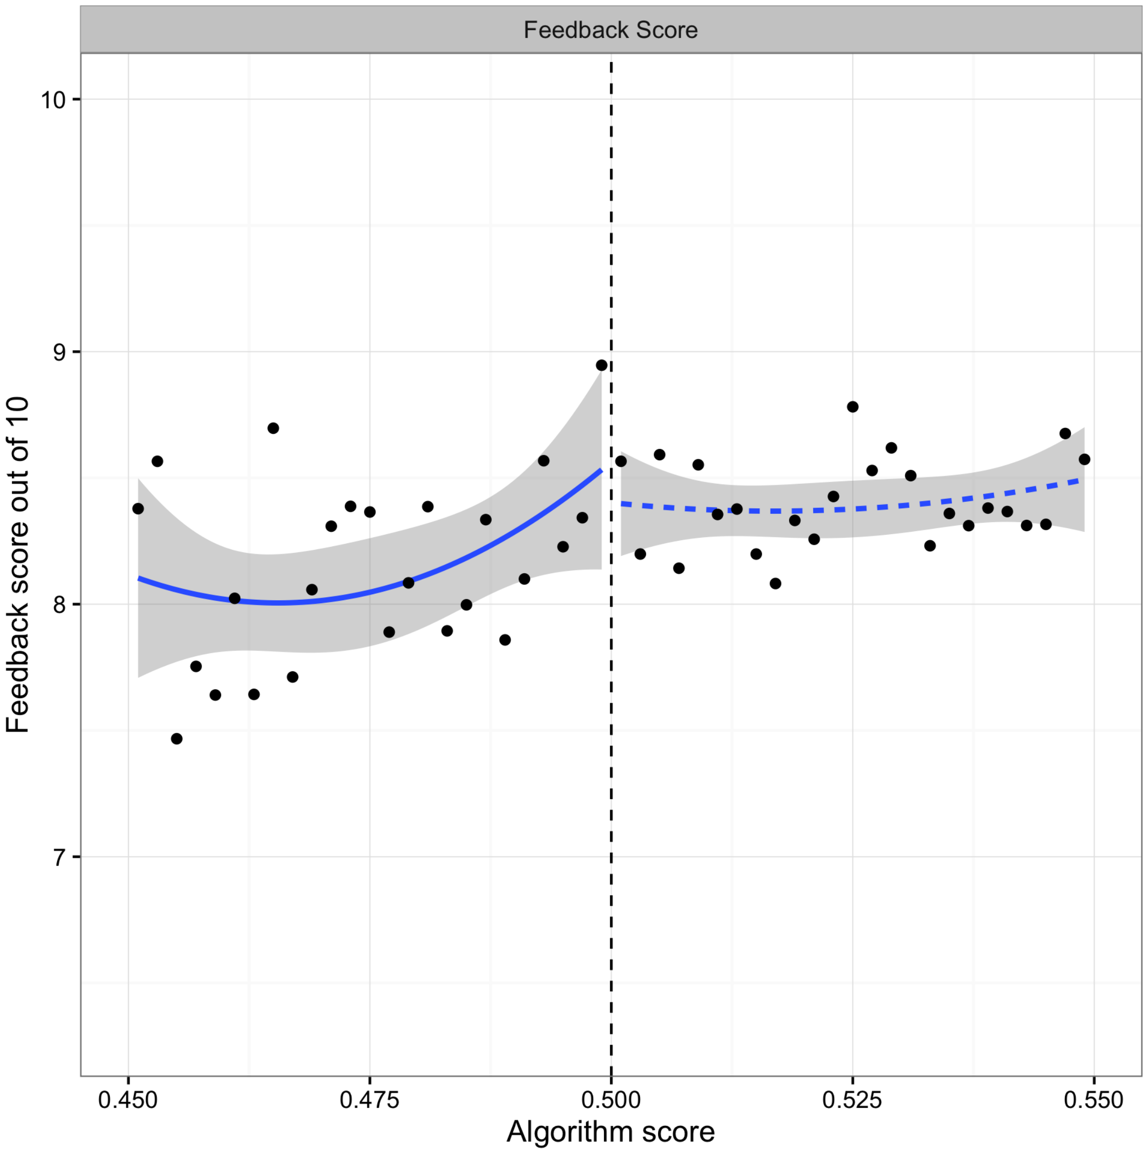

important. For high-valued jobs, employers will rely more heavily on personal, detailed, and private information than on a "minimal signal". Conversely, as cheap information, the recommendation will likely be relied upon more for low value jobs because it is less likely an employer will expend the cognitive effort required to investigate each applicant.

To help route internal resources, the platform assigns all jobs posted on the site a value category of "very low value", "low value", "medium value", or "high value". These value classifications are based on the predicted total earnings of the job, based on the employer's prior behavior on the platform and job characteristics. To get a better sense of the differences in value between lower and higher value jobs, Table 4 presents summary statistics by job value. To facilitate easy analysis, we separate the jobs in our sample into two groups: the first group includes "very low value" and "low value" jobs and the second group includes "medium value" and "high value" jobs. We choose to use these two groups to make interpretation easier and to increase the sample size of the comparison groups. Using all four categories does not alter the results.

Table 4

Table 8: Opening Summary Stats by Value Group, Organic Applications

	VLV	LV	MV	HV
Median Non-Zero Billing	\$80	\$210	\$500	\$700
Mean Number of Applications	17.42	25.38	27.35	30.34
Mean Number of Recommended	5.46	7.45	8.00	11.29
Mean Number of Hired	0.50	0.42	0.31	0.30
Total Number of Jobs	89,881	15,741	23,554	7,372

Table 5 reports results from an interaction model based on our preferred specification, a non-parametric local linear model allowing for slope differences on each side of the discontinuity. We interact the main effect of the recommendation with the the value category of the job to test for differential effects of the recommendation by job value. As job values is clearly fixed within a job, we are unable to use a job-fixed effects. Instead use an employer-fixed effects and add additional controls for job heterogeneity.

Turing to Table 5, we can see that the interaction term (recommended x LV) is positive and significant for all models except for the model (2) where the outcome of interest is the probability of hiring an applicant who has already been interviewed. This indicates that the effect of the recommendation is differentially larger for low-value jobs than high-value jobs in influencing who an employer chooses to view, chooses to interview, and chooses to hire, but that there is no differential effect of the recommendation on which interviewed applicants an employer hires. Recall, that our main finding showed that the recommendation has no effect on the employer's decision over which interviewed applicants he/she should hire. Regardless of the value of the job posted, once employers obtains private information from interviewing, they do not find any additional use for the recommendation. However, the unconditional effect of the recommendation on who is hired is nearly 5 times as larger for low value jobs compared to high value jobs. This confirms that employers are much more likely to be influenced by the recommendation on low value jobs than high value jobs and lends support to our interpretation of the recommendation as providing value only as a "minimal signal".

Table 5

Local Linear RD estimates of the effect of being Recommended by Job Value

	Dependent variable:				
	Applicant Viewed	Applicant Interviewed		Applicant Hired	
	(1)	(2)	(3)	(4)	(5)
Recommended (rec)	0.062***	0.019***	0.014	0.002*	−0.002
	(0.004)	(0.003)	(0.011)	(0.001)	(0.016)
Lower Value Jobs (LV)	−0.059***	−0.013***	0.015	0.004***	0.088***
	(0.005)	(0.003)	(0.012)	(0.001)	(0.017)
rec X LV	0.020***	0.010***	0.002	0.007***	−0.007
	(0.005)	(0.003)	(0.012)	(0.001)	(0.019)
Sample	All Applicants	All Applicants	Viewed Applicants	All Applicants	Interviewed Applicants
Job-level FE	No	No	No	No	No
Employer-level FE	Yes	Yes	Yes	Yes	Yes
I-K Optimal Bandwidth	0.059	0.047	0.036	0.06	0.049
Dep. Var. Mean	0.289	0.081	0.274	0.021	0.185
Observations	844,138	694,218	153,686	856,319	58,549

Notes: The bandwidth used in each model was calculated using the Imbens-Kalyanaraman Optimal Bandwidth Calculation. The estimations are from a local-linear model which allows for slope differences on either side of the cutoff. The omited group are jobs which are categorized as medium value or high value. Covariates include the applicant's bid on the job posting and the page order of the applicant, job type indicators, and first job posting indicators. Model (1) the outcome is an indicator if the applicant is viewed. Model (2) the outcome is an indicator if the applicant is interviewed. Model (3) the outcome is an indicator if the applicant is interviewed conditional on the applicant being viewed. Model (4) the outcome is an indicator if the applicant is hired. Model (5) the outcome is an indicator if the applicant is hired conditional on the applicant being interviewed. Heteroskedasticity-robust standard errors clustered at the job posting level are reported. Significance indicators: $p \leq 0.05 : *, .p \leq 0.01 : **$ and $p \leq .001 : ***$.

4 Discussion and Conclusions

We demonstrated that a recommendation made by a computational algorithm dramatically increases the likelihood of a job applicant being hired by an employer over similarly qualified others. In particular, this algorithmic recommendation increases the likelihood an applicant will be hired by 50% over similarly qualified applicants. More specifically, this effect is attributable to having a "foot in the door" advantage because 70% of this effect is attributable to recommended applicants being more likely to be noticed at the initial stages of the hiring process. When additional information regarding the job applicants are gathered by the employer, such as private, match specific information, the recommendation no longer has any effect. Conditional on being interviewed, recommended and non-recommended job applicants have an identical likelihood of being hired. Our regression discontinuity design allows us to account for the alternative explanation, that these applicants are "head and shoulders above the rest" due to their superior ability. In short, algorithmic recommendations act as a "minimal cue" - elevating job applicants in the initial stages, where cheap and simple information is used to sort applicants into groups that are worthy or unworthy of additional investigation.

With regard to labor market hiring, this paper neatly address two issues that the literature has grappled with to date. First, this paper empirically separated the actual quality of a job application from an external cue, neatly demonstrating the value of a cue in getting hired, over and above the human capital differences (Spence 1973, Steglitz 1975). This paper reveals the value of "getting a foot in the door", by revealing insight into the heretofore unobservable hiring decision process by identifying separate decisions that employers make in deciding whom to hire. Merely moving beyond an initial screening process dramatically alters one's employment outcomes. Having this visibility provides an explanation to labor market scholars who demonstrate that hiring outcomes are dramatically improved for those who are referred to an employer. One reason could simply be that those applicants are privileged early in the search process.

This paper may also shed light on our understanding of the discrimination and bias minority job seekers face. There is little doubt that members of certain social categories, such as race or gender, are disadvantaged in the labor market. While most of the work has worked on uncovering this bias, much less work has been able to isolate the particular mechanisms contributing to these outcomes. Our findings suggest that the disparate hiring outcomes may occur very early in the hiring process. For example, our theory suggests that any small differences in perceptions of a job applicant abilities that occur very early in the hiring process will have a large impact on eventual hiring outcomes (Bartos et al 2016). If so, then this finding provides support for the idea that labor market discrimination could be the result of very small perceived differences that dramatically alter the likelihood of moving beyond the very early stages of a job search - such as the callbacks on may receive from a correspondence study. This means these differences are used to sort applicants very early in the hiring process (Bothello and Abraham 2017) even if there are practically no differences in their actual abilities.

Finally, the gig-economy is burgeoning and additional visibility into how matches are made in this novel environment is valuable. Early work in this vein touted the potential benefits to how matching can be dramatically improved for the two-sided markets for labor (Horton 2015). We contribute to this stream of work by demonstrating not only that computational matching algorithms "work" by increasing the likelihood of a job applicant being hired, but also to demonstrate how this comes about. Our findings have implications for future work on how best to design market platforms, as ultimately, computationally derived algorithms are useless unless humans heed their recommendations.

Appendix A

Types of Jobs on the Platform

Web, Mobile & Software Dev
Desktop Software Development
Ecommerce Development
Game Development
Mobile Development
Product Management
QA & Testing
Scripts & Utilities
Web Development
Web & Mobile Design
Other - Software Development
IT & Networking
Database Administration
ERP / CRM Software
Information Security
Network & System Administration
Other - IT & Networking
Data Science & Analytics
A/B Testing
Data Visualization
Data Extraction / ETL
Data Mining & Management
Machine Learning
Quantitative Analysis
Other - Data Science & Analytics
Engineering & Architecture
3D Modeling & CAD
Architecture
Chemical Engineering
Civil & Structural Engineering
Contract Manufacturing
Electrical Engineering
Interior Design
Mechanical Engineering
Product Design
Other - Engineering & Architecture
Design & Creative
Animation
Audio Production
Graphic Design
Illustration
Logo Design & Branding
Photography
Presentations
Video Production
Voice Talent
Other - Design & Creative
Writing
Academic Writing & Research
Article & Blog Writing
Copywriting
Creative Writing
Editing & Proofreading
Grant Writing
Resumes & Cover Letters
Technical Writing
Web Content
Other - Writing
Translation
General Translation
Legal Translation
Medical Translation
Technical Translation
Legal
Contract Law
Corporate Law
Criminal Law
Family Law
Intellectual Property Law
Paralegal Services
Other - Legal
Admin Support
Data Entry
Personal / Virtual Assistant
Project Management
Transcription
Web Research
Other - Admin Support
Customer Service
Customer Service
Technical Support
Other - Customer Service
Sales & Marketing
Display Advertising
Email & Marketing Automation
Lead Generation
Market & Customer Research
Marketing Strategy
Public Relations
SEM - Search Engine Marketing
SEO - Search Engine Optimization
SMM - Social Media Marketing
Telemarketing & Telesales
Other - Sales & Marketing
Accounting & Consulting
Accounting
Financial Planning
Human Resources
Management Consulting
Other - Accounting & Consulting

5 Appendix B

Table B1 shows that the results of Table 3 are robust to both bandwidth choices and model specifications. The results do not substantively differ from those presented in table 3 indicating that our main specification is robust to different choices of both bandwidth or functional form.

Table B1: Regression Results Sensitivity Analysis

Bandwidth	Hired	Hired|Interview	Interview	Interview|View	View
Linear Model					
0.02	0.006 (0.002)	-0.016 (0.032)	0.026 (0.003)	0.009 (0.01)	0.086 (0.004)
0.03	0.008 (0.001)	-0.017 (0.021)	0.029 (0.002)	0.012 (0.007)	0.09 (0.003)
0.04	0.009 (0.001)	-0.011 (0.015)	0.031 (0.002)	0.014 (0.006)	0.092 (0.003)
0.05	0.009 (0.001)	-0.011 (0.012)	0.032 (0.001)	0.012 (0.005)	0.093 (0.002)
0.06	0.01 (0.001)	-0.008 (0.01)	0.032 (0.001)	0.013 (0.004)	0.093 (0.002)
0.07	0.011 (0.001)	-0.002 (0.009)	0.033 (0.001)	0.015 (0.004)	0.095 (0.002)
0.08	0.011 (0.001)	0 (0.009)	0.034 (0.001)	0.018 (0.004)	0.095 (0.002)
Quadratic Model					
0.02	0.007 (0.002)	0.019 (0.046)	0.026 (0.004)	0.007 (0.016)	0.085 (0.006)
0.03	0.007 (0.002)	-0.024 (0.03)	0.026 (0.003)	0.006 (0.011)	0.087 (0.004)
0.04	0.007 (0.001)	-0.029 (0.023)	0.028 (0.003)	0.011 (0.009)	0.089 (0.004)
0.05	0.008 (0.001)	-0.015 (0.018)	0.032 (0.002)	0.019 (0.007)	0.092 (0.003)
0.06	0.008 (0.001)	-0.017 (0.016)	0.032 (0.002)	0.012 (0.006)	0.093 (0.003)
0.07	0.008 (0.001)	-0.019 (0.014)	0.03 (0.002)	0.009 (0.006)	0.091 (0.003)

Notes: The results are from local linear and local quadratic regressions allowing for slope differences across the threshold. The bandwith was varied between .02 and .08 indicating that the results of table 3 are robust to bandwith and model specifications. Covariates used include: job opening fixed effects, the order of the application on the page, and the applicant's bid. Heteroskedasticity-robust standard errors clustered at the job posting level are reported.

6 Appendix C

Figure C1: Density of Algorithm Scores

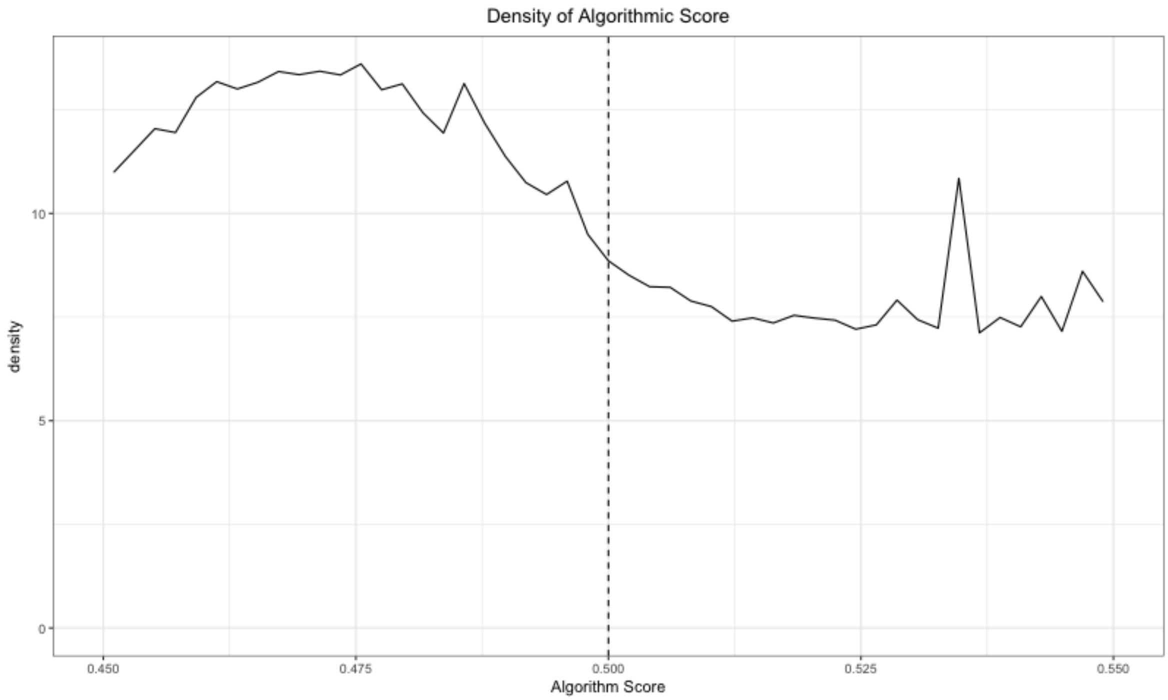

Notes: Figure C1 shows the density of best match scores across applications in our sample. The lack of bunching[1] above 0.5 indicates that applicants at the margin are unable to manipulate whether or not they are recommended. Since the algorithm is a black box to its users, this is not surprising. Moreover the algorithm relies on signals that are hard for users to manipulate, such as verified work history and past feedback.

[1]McCrary (2008). "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test". Journal of Econometrics. 142 (2): 698714

Chapter 3: A Model of Competing Recruiting Channels

Abstract

The rapid growth of online information on workers, like LinkedIn's profiles or Monster.com's resume database, has dramatically lowered the cost for employers to directly reach out to (headhunt) workers. The promise is that by providing employers with an additional channel of sourcing matches, it will increase the probability of filling vacancies and the average quality of hires. I construct a search theoretic model of hiring that explicitly models both the headhunting channel and the application channel. I study equilibrium outcomes when workers can optimally respond by deciding whether or not to apply. I show that while lower headhunting costs improve the average quality of hires, it can actually decrease the probability of filling vacancies in equilibrium. This is due to workers optimally choosing to apply with a lower probability.

In this paper I present a search theoretic model of hiring when the employer has the option of screening applicants as well as headhunting. Following Morgan and Vardy (2009), a firm is looking to fill one vacancy. Workers are either a quality match or not, but match quality is imperfectly observed by both the applicant and the firm prior to employment. A worker can choose to apply to the vacancy after observing a signal of their match quality, but doing so incurs a fixed cost. The firm can screen their applicant, if they receive one, for free. To capture the possibility of headhunting, the firm can also additionally exert a fixed cost to obtain a worker from the general pool.

The equilibrium of the model is pinned down by a pair of applicant and employer strategies. Both strategies are cutoffs; the worker will apply if and only if they observe a signal above a threshold and the employer will headhunt if and only if they observe a signal from their applicant below a threshold. The main comparative static in question is with respect to the headhunting cost. I show that while expected quality of hires increase as headhunting cost decreases, the probability of filling the vacancy decreases for a large set of parameters. This is because the worker will require a higher signal to apply, decreasing the probability that the employer has two workers to choose from.

There is an established literature using search theoretic models to study the employer's hiring decision. One area examines discrimination; Morgan and Vardy (2009) show that disparity can arise in a sequential search model if minority applicants provide a noisier signal of quality to majority employers. Cornell and Welch (1996) study a similar context but with fixed sample search. Search theory has also been applied to workers' job search intensity, strategies, and outcomes; Cyronek (2016) provides a detailed review of directed job search models while Benhabib and Bull consider optimal job search intensity when workers have the option of searching while employed. Equilibrium search-theoretic models of the labor market have examined market level outcomes of many firms searching for many workers, studying effects on unemployment rates, unemployment durations, and wage distributions; Rogerson et al (2005) provide an extensive review. Recently Devaro and Gürtler (2018) model the matching process as a strategic coordination game, wherein employers can advertise vacancies and workers can advertise their availability. To my knowledge there is not yet a model incorporating the possibility of headhunting with endogenous job search intensity.

1 Model

1.1 Setup

A firm has 1 vacancy which if filled by worker at wage w, gets the firm:

$$\pi(m_i, w) = m_i - w$$

Worker i is either able to do job or not so $m_i \in \{0, 1\}$.

Match value m_i is equal to 1 with probability $\rho \in (0, 1)$.

Consider the case where a worker (a) can apply and the employer j has the option of inviting another worker (b). The two workers have match values drawn independently from the same distribution.

Invites

Employers can directly reach out to the worker at cost $c_h > 0$. If employer pays c_h, employer receives a worker (b) and a noisy signal $z_b = m_b + \nu_b$ with $\nu_b \sim \mathcal{N}(0, \sigma_\nu^2)$.

Applications

If applicant a applies, employer sees a noisy signal $z_a = m_a + \nu_a$ with $\nu_a \sim \mathcal{N}(0, \sigma_\nu^2)$.

When worker a encounters the vacancy, they receive a signal $s = m_a + \epsilon$ with $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. Based on this they decide whether or not to apply at cost $c_a > 0$.

If worker a applies, their expected net payoff is $\Pr[\text{Hire}|s]w - c_a$.

Equilibrium is determined by the following decision rules:

1. Worker a applies IFF they receive a signal $s > s^*$.
2. If worker a applies and employer receives signal $z_a > z^*$, employer hires a immediately.
3. If worker a did not apply or if $z_a \leq z^*$, then employer reaches out to worker b and receives signal z_b.
4. Employer then hires the worker with highest expected m, if $E[m] - w > 0$.

Equilibrium is pinned down by the pair of thresholds s^*, z^*. Following Morgan and Várdy (2009), first define employer's posterior beliefs over the candidates after all signals are accounted for: $q_a(z_a, s^*) = E[m_a|z_a, s > s^*]$ and $q_b(z_b) = E[m_b|z_b]$. By Bayes' Rule and since s, z are independent conditional on m:

$$\begin{aligned} q_a(z_a, s^*) &= \frac{\Pr[z_a|m=1]\Pr[m=1|s>s^*]}{\Pr[z_a|m=1]\Pr[m=1|s>s^*] + \Pr[z_a|m=0](1-\Pr[m=1|s>s^*])} \\ &= \frac{\Pr[z_a|m=1]\Pr[s>s^*|m=1]\rho}{\Pr[z_a|m=1]\Pr[s>s^*|m=1]\rho + \Pr[z_a|m=0]\Pr[s>s^*|m=0](1-\rho)} \\ &= \frac{\phi[(z_a-1)/\sigma]\Phi[(s^*-1)/\sigma]\rho}{\phi[(z_a-1)/\sigma]\Phi[(s^*-1)/\sigma]\rho + \phi[z_a/\sigma]\Phi[s^*/\sigma](1-\rho)} \end{aligned} \tag{1}$$

and

$$q_b(z_b) = \frac{\phi[(z_a-1)/\sigma]\rho}{\phi[(z_a-1)/\sigma]\rho + \phi[z_a/\sigma](1-\rho)}$$

where $\phi(.)$ and $\Phi(.)$ respectively denote the probability density and cumulative distribution of a standard Normal random variable. It can be shown that $q_a(z_a, s^*)$ is strictly and continuously increasing in s^* and z_a and $q_b(z_b)$ is strictly and continuously increasing in z_b. Moreover, $q_a(z_a, s^*)$ and $q_b(z_b)$ have full support over $(0, 1)$.

To rule out trivial cases, I assume that:

1. Applicants apply if they will be hired for sure i.e. $w - c_a > 0$, and
2. Employers with no applicant will headhunt, i.e. $\int_w^1 (q_b - w)\, dG_b(q_b) - c_h > 0$

Employer's Strategy.

Given correct beliefs of applicant strategy s^* and after seeing signal z_a the employer infers q_a and will headhunt IFF the expected benefit of headhunting $\pi_h(q_a)$ is above the headhunting cost c_h. An applicant with $q_a < w$ would not be hired regardless of headhunting, so this case is equivalent to having no applicant. The interesting case is $q_a \geq w$, in this case $\pi_h(q_a)$ is the expected improvement above q_a and the employer will headhunt IFF:

$$\pi_h(q_a) = \int_w^1 (q_b - q_a)\, dG_b(q_b) - c_h > 0. \tag{2}$$

Since $\pi_h(q_a)$ is strictly and continuously decreasing in q_a for all $q_a \geq w$ and $\int_w^1 (q_b - w)\, dG_b(q_b) - c_h > 0$, there exists a unique q_a^* such that $\pi_h(q_a^*) = 0$. This is the minimum applicant quality at which the employer forgoes headhunting. For any applicant strategy s^*, there is a corresponding z^* such that $E[m|s > s^*, z = z^*] = q_a^*$. This z^* is the minimum signal the employer needs to see from the applicant to forgo headhunting and is thus the employer's best response to applicant strategy s^*.

From equation 1, $q_a(z, s^*)$ is strictly and continuously increasing in s^* for every $z \in \mathbb{R}$, so the employer's best response function $z^*(s^*)$ is also strictly and continuously decreasing in s^*. Moreover, as $s^* \to \infty, z^*(s^*) \to -\infty$. On the other extreme, the limit of $s^* \to \infty$ is the case where all applicants apply as in the Babbling Equilibrium described below. The employer's best response in this extreme ($\bar{z}$) is such that $E[m_a|z_a = \bar{z}] = q_a^*$.

Applicant's Strategy.

Applicants take employer strategy z^* and employer belief over applicant strategy s' as given. Recall that employer belief over s' is used to update their posterior belief $q_a(s', z_a)$ and thus affect the probability that an applicant is hired. Upon receiving signal s about themselves, the applicant will apply if the expected benefit of doing is strictly greater than the application cost. Thus the applicant applies IFF $\pi_a(s, z^*, s') = \Pr[\text{hired}|s, z^*, s'] - c_a > 0$.

To see that $\pi_a(s, z^*, s')$ is strictly increasing in s, recall that $z_a = s + \eta_a - \epsilon$. Thus for any fixed $y \in \mathbb{R}$, $\Pr[z_a > y|s] = 1 - \Phi((y - s)/(\sigma_\eta^2 + \sigma_\epsilon^2)^{\frac{1}{2}})$ is strictly increasing in s. Since employer's posterior beliefs over applicant quality is only affected by z_a when s' is held constant, and $q_a(s', z_a)$ is strictly increasing in z_a, $\Pr[\text{hired}|s, z^*, s']$ must be strictly increasing in s.

Therefore for any interior solution the applicant's best response is a threshold strategy: apply IFF $s > s^*$ for an unique $s^* \in \overline{\mathbb{R}}$[1]. For a given pair z^*, s', there are 3 possible cases for $\pi_a(s, z^*, s')$, and applicants' corresponding best response function $s^*(z^*, s')$ are as follows:

1. $\pi_a(s, z^*, s') < 0\, \forall s \in \mathbb{R}$: Never apply so $s^* = \infty$
2. $\exists s^* \in \mathbb{R}$ such that $\pi_a(s^*, z^*, s') = 0$: Apply IFF $s > s^*$ for the unique $s^* \in \mathbb{R}$
3. $\pi_a(s, z^*, s') > 0\, \forall s \in \mathbb{R}$: Always apply so $s^* = -\infty$

Existence of Babbling Equilibrium

It is possible to have an equilibrium where all applicants apply, i.e. $s^* = -\infty$. This equilibrium holds whenever the probability that an applicant with $m = 0$ is mistakenly hired is high enough to

[1] $\overline{\mathbb{R}}$ denotes the extended reals, i.e. $\overline{\mathbb{R}} = \mathbb{R} \cup (-\infty, \infty)$

overcome the applicant cost. More precisely $E[\text{hired}|m_a = 0]w - c_a > 0$ given employer believes that all applicant apply. Threshold q_a^* is as above, since employer gains no information from the act of applying, z^* is simply such that $E[m_a|z_a = z^*] = q_a^*$.

$$\begin{aligned} E[\text{hired}|m_a = 0] &= \Pr[z_a > z^*|m_a = 0] + \Pr[\max\{z_b, z(w)\} < z_a \leq z^*|m_a = 0] \\ &= (1 - \Phi[z^*/\sigma]) + \int_{z(w)}^{z^*} \Pr[z_b < z_a]\, dF(z_a) \\ &= (1 - \Phi[z^*/\sigma]) + \int_{z(w)}^{z^*} \rho\Phi[(z_a - 1)/\sigma] + (1-\rho)\Phi[z_a/\sigma]\, dF(z_a) \end{aligned}$$

where $dF(z_a)$ is CDF of ν_a and $z(w)$ is the signal that makes the employer indifferent between hiring a and not, if a was the only choice, i.e. $q_a(z(w)) = w$.

Existence of Interior Equilibrium

As described above, the search for an equilibrium is technically over a 3 dimensional space (z^*, s^*, s'), where s' is the employer's belief over s^*. I use the fact that, in equilibrium, the employer must have correct belief over s^* to collapse this space down to a 2 dimensional space (z^*, s^*). This is done by limiting applicant's best response strategies to cases where the employer has the correct belief over s^*. To clarify, the applicant takes z^* as given and then chooses an optimal s^* as if the employer also holds belief s^*.[2]

An interior equilibrium, if it exists, is at the intersection of the employer's best response function $z^*(s^*)$ and the set of applicant's best response functions corresponding to correct employer beliefs, i.e. $s^*(z^*) = s^*(z^*, s' = s^*)$. These best response functions are shown graphically in Figure 1.1. Since both best response functions, if finite, are strictly monotone, there can be at most 1 interior equilibrium. Thus if the interior equilibrium exists, it is also unique.

Since $s^*(z^*)$ is strictly increasing and $z^*(s^*)$ is strictly decreasing, a single crossing occurs if and only if there exists $z > z'$ such that $z^*(s^*(z)) > z$ and $z^*(s^*(z')) < z'$. Since both functions are strictly monotone, we just need to compare the following limits:

1. From above, $\lim_{s^* \to \infty} z^*(s^*) = -\infty$ and,
2. $\lim_{s^* \to -\infty} z^*(s^*) = \bar{z} \in \mathbb{R}$.
3. As long as $w > c_a$, there exists $\underline{z}$ such that all applicants would apply. This unique $\underline{z}$ is defined by $E[\text{hired}|m_a = 0, z^* = \underline{z}]w - c_a = 0$. Therefore $\lim_{z^* \to \underline{z}} = -\infty$.
4. Recall that $s^*(z^*)$ is defined as the applicant's best response to z^* if the employer holds belief $s^*(z^*)$. Thus for any z^*, there exists an unique and finite $s^*(z^*)$. This is because as s^* increases, the employer's belief over the applicant's q tends towards 1, so will hire the applicant with increasing probability after headhunting.

The inequality $z^*(s^*(z')) < z'$ is guaranteed by limits 1 and 4. The inequality $z^*(s^*(z)) > z$ holds if and only if $\bar{z} > \underline{z}$. Since $\bar{z}$ is the employer's best response to all applicants applying $(s^* = \infty)$ and $\underline{z}$ is the z^* such that all applicants' best response is to apply. Therefore:

[2] This means that at off-equilibrium points, the applicant is responding to a set $z^*, s' = s^*$ that is inconsistent with the employer's best response. Alternatively, the dimensional reduction can be achieved by having the applicant respond to a set z^*, s' where employer belief s' is consistent with z^* being their best response, i.e. $q_a(s', z_a) = q^*$. These two approaches lead to the same set of equilibria, since in equilibrium $s' = s^*$.

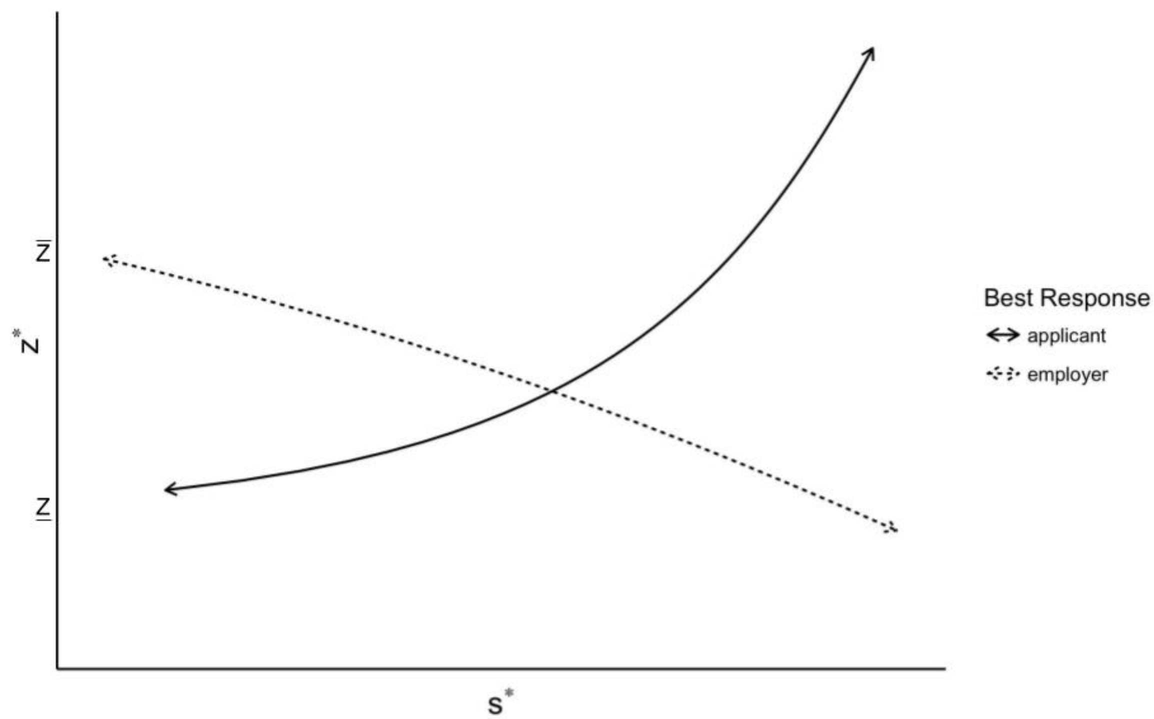

Figure 1: Unique equilibrium is achieved at the intersection of the two best response functions.

Theorem: The necessary and sufficient condition for an unique interior equilibrium to exist is that a Babbling Equilibrium does not exist.

1.2 Comparative Statics: Partial Equilibrium

I first consider the partial equilibrium effects of decreasing the cost of headhunting c_h, holding applicant strategy s^* constant. This is analogous, for example, to an experiment run by a large job site or online labor market to test a new headhunting tool.

Lemma: Assume that the change in c_h is within the set of parameters allowing an interior equilibrium. Then, holding s^* constant, a decrease in c_h causes:

1. A strict increase in z^*, i.e. employer requires a strictly higher signal from the applicant to forgo headhunting.
2. A strict increase in the probability that the employer hires from headhunting pool.
3. A strict increase in the expected quality of worker hired.
4. No change in the probability that a job is filled.
5. A strict increase in employer welfare.

Effects 1 and 2 follow from equation 2. Since s^* stays fixed, the distribution of applicants seen by employers also remains fixed. Thus changes in the quality of hire can only occur from the employer headhunting when they did not previously, and obtaining a worker of higher expected quality than their applicant. This gives effect 3.

Effect 4 follows from the assumption that employers headhunt when $q_a \leq w$. To see this, the job goes unfilled only when either:

a) $q_a, q_b \leq w$,

b) $q_a \leq w$ and the employer does not headhunt,

c) the applicant does not apply and $q_b \leq w$, or

d) the applicant does not apply and employer does not headhunt.

In interior equilibria, employers headhunt when $q_a \leq w$, so cases b) and d) do not occur. Since s^* is fixed and q_b is unaffected by applicant decisions, the probability of cases a) and c) are unaffected by the change in c_h.

1.3 Comparative Statics: General Equilibrium

As c_h is lowered from H to L in general equilibrium, the following changes occur to equilibrium strategies:

- The minimum applicant quality required for employer to forgo headhunting (q^*) increases. This follows directly from Equation 2.
- The minimum signal required for employer to forgo headhunting (z^*) increases.
- The minimum signal required for applicants to apply (s^*) increases as the higher z^* decreases the probability of hire for any given s.

Correspondingly, the probability of hiring an applicant decreases, both since the employer is more likely to headhunt and the applicant is less likely to apply. Similarly, the probability of hiring a headhunted applicant increases. Applicant welfare is thus strictly lower in expectation.

Most interestingly, the effect on the probability of fill is ambiguous. Returning to the two cases above where a job goes unfilled in an interior equilibrium, the effects of lowering c_h is as follows:

1. Applicant would have applied, but now do not, and headhunted worker is not hired. I.e. $s^*_H < s \leq s^*_L$ and $q_b \leq w$. This effect is strictly negative on fill rate.
2. Conditional on applying and a fixed signal z_a, the employer infers a strictly higher q_a since their belief over s^* has updated upwards. Thus for any realization such that the applicant applies and $q_b \leq w$, the employer is strictly more likely to hire. Note that as above, changing c_h does not affect fill rate when $q_a \leq w$.

The overall effect thus depends on the relative magnitude of the two effects, i.e. the relative frequency of these two sets of realizations. Figure 2 shows computationally simulated frequencies for 1 set of parameters.

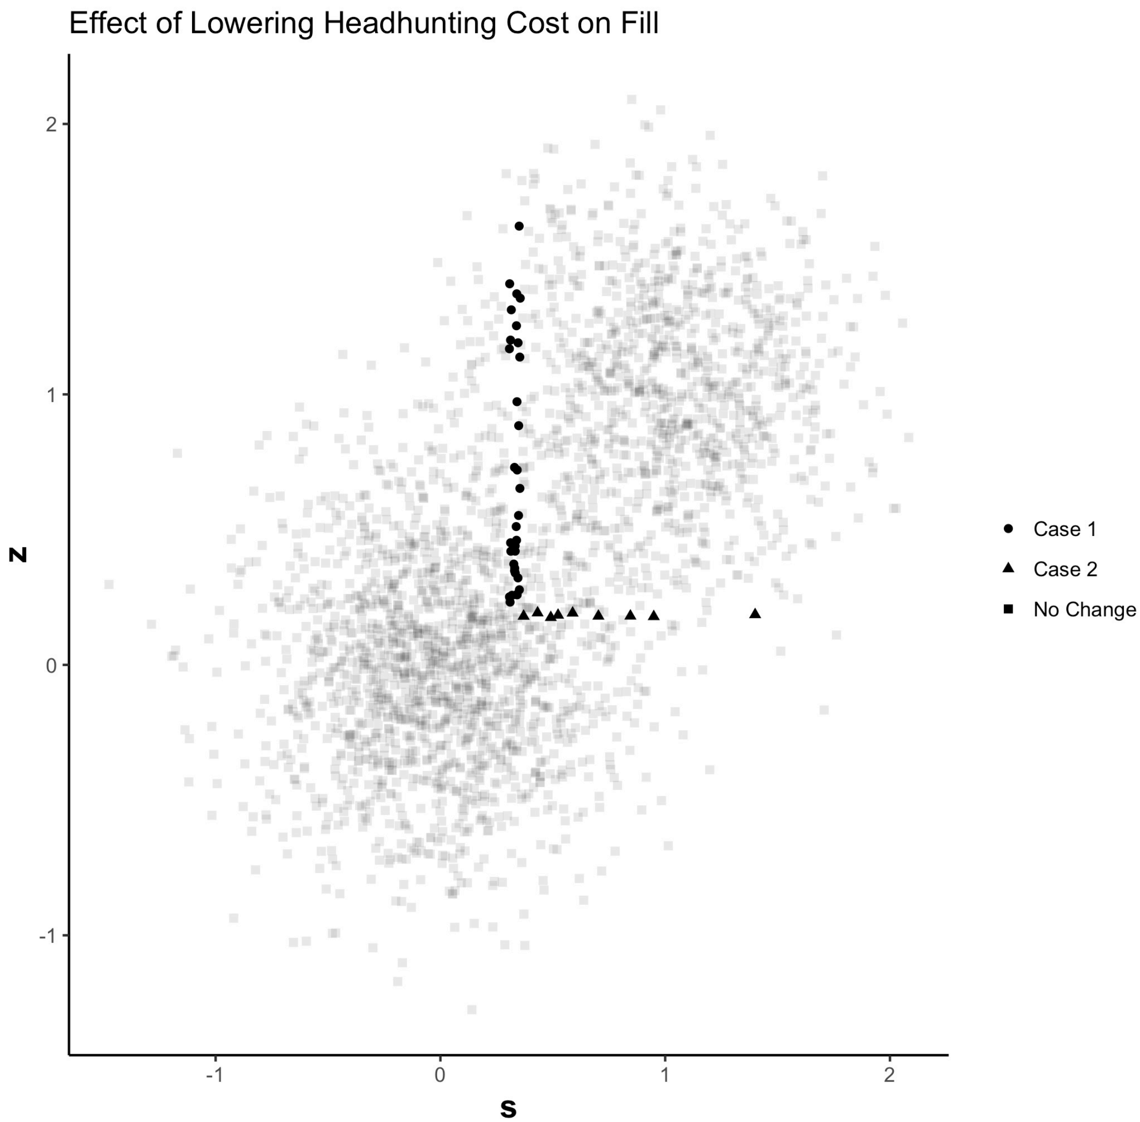

Figure 2: Simulation with c_h decreasing from 0.18 to 0.08. Other parameters fixed at $\rho = 0.4$, $w = 0.3$, $c_a = 0.12$, and $\sigma_\epsilon, \sigma_e ta = 0.4$. Each point is a realization where $q_b \leq w$.

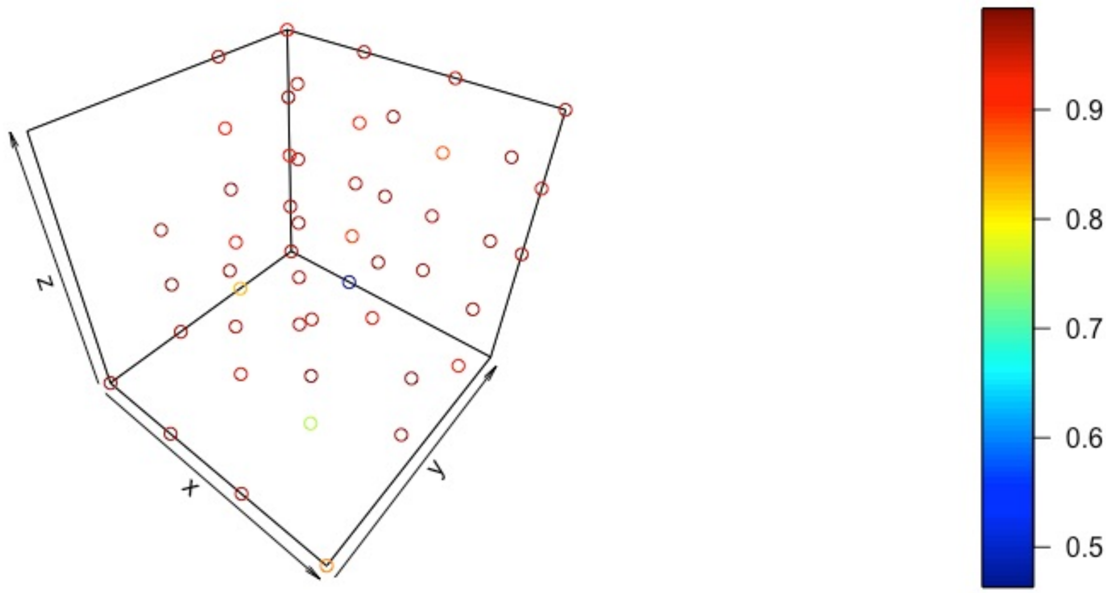

Figure 3: Simulated correlation between probability of fill and c_h across parameter sets with interior equilibria. Axes are: $x = \rho,\ y = w,\ z = c_a$

For any given equilibrium, let z^w the lowest signal that the employer will hire if the applicant was their only viable choice, i.e. $q(s^*, z^w) = w$. Then the expected change in fill rate for an arbitrarily small decrease in c_h is proportional to:

$$-\Pr[s = s^*, z > z^w] + \Pr[z = z^w, s > s^*]|\frac{\delta z^w}{\delta s^*}|$$

where the first term is the effect from Case 1 and the second the effect from Case 2. Figure 3 shows results from simulations across all parameters that allow for an interior equilibrium. This shows that effect 1 dominates and thus fill rate decreases with c_h for the vast majority of parameters.

Extension to Multiple Candidates

1 Applicant, 2 Potential Headhunted Workers. For the applicant, the decision to apply is still based on $\Pr[\text{hired}|s] - c_a > 0$. With two workers in the headhunting pool, the employer has a more valuable outside option to hiring the applicant. Thus the cutoff s^* is strictly higher. The employer now has 2 cutoffs: z_a^*, z_b^*, where employer headhunts once if $z_a < z_a^*$ and twice if $z_b < z_b^*$. These thresholds can be determined working backwards from z_b^* just as above.

2 Applicants, 1 Potential Headhunted Workers. Assume applicants move simultaneously, without knowledge of whether or not the other applicant will apply. We'll also look for a symmetric pure strategy equilibrium, where each applicant will apply IFF they receive $s > s^*$. Since probability of being hired is strictly lower, the threshold $s > s^*$ is strictly higher. On the employer side, they compare any applicants they have and again apply a cutoff z^* to decide whether or not to headhunt.

2 Discussion

The model shows that while decreasing headhunting cost on average improves the quality of hires, it can potentially unravel the application channel, decreasing the probability of filling vacancies and increasing the unemployment rate. So far the comparative statics has been restricted to comparison between interior equilibria. A broader examination should also consider the extensive margin, where decreasing headhunting cost may induce an employer to headhunt in the absence of an applicant when they previously would not have. One way to capture this is to consider a distribution of employers with heterogeneous headhunting costs relative to output. A further extension is to endogenous wages, perhaps allowing wages to differ between hires from the application channel vs. the invitation channel.

3 Appendix: Simulation

Simulation for the general equilibrium solution for a set of parameters ρ, σ, c_h, c_a is as follows. Brute force components used 5,000 draws of m_a, m_b, s, ζ.

1. Define function c_hh_benefit_calc(q) that calculates the expected benefit of headhunting if employer already has applicant with quality $E[m = 1|\zeta, s > s*] = q$. Note if $q < w$ then q is set to w. This is calculated by brute force.

$$\text{Expected benefit} = \int_q^1 E[m = 1|\zeta] - q \, d\zeta$$

2. Set q^* as the minimum applicant quality at which employer does not headhunt. Define function c_sjh_star_calc(q^*, s^*) to calculate the corresponding signal threshold ζ^*.

3. Determine applicant's decision to apply based on function c_p_hired_si_calc(s, ζ^*, s^*) that calculates the expected probability of hire if applicant with signal s applies, assuming employer believes applicant is using threshold s^* and employer uses ζ^*. This is calculated by brute force simulation of many headhunted candidates.

4. Solve for equilibrium s^*, ζ^* by using a root finding algorithm to find s^* that sets c_p_hired_si_calc(s^*, c_sjh_star_calc(q^*, s), s^*)$w - c_a$ to 0.

Calculation for the general equilibrium comparative static with respective to a parameter is by repeating the above simulation for varying values of the parameter in question.

References

Agan, Amanda Y. and Sonja B. Starr, "Ban the box, criminal records, and statistical discrimination: A field experiment". Working Paper, 2016.

Agrawal, Ajay K., John Horton, Nico Lacetera, and Elizabeth Lyons, "Digitization and the contract labor market: A research agenda". *Economics of Digitization*, 2013.

Arrow, Kenneth J. , "Higher education as a filter". *Journal of Public Economics*, 1973, 2(3), pp. 193–216.

Autor, David H and David Scarborough,"Does job testing harm minority workers? Evidence from retail establishments". *The Quarterly Journal of Economics*, 2008, pp. 219–277.

Barach, Moshe and John Horton, "How do Employers Use Compensation History: Evidence From a Field Experiment". Working Paper, 2018.

Bartoš, Vojtěch, Michal Bauer, Julie Chytilová and Filip Matějka, "Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition". *American Economic Review*, 2016, 106(6), pp. 1437-75.

Bendick, Marc Jr., Lauren E. Brown, and Kennington Wall, "No Foot in the Door". *Journal of Aging & Social Policy*, 1999, 10(4), pp. 5-23.

Benhabib, Jess and Clive Bull, "Job Search: The Choice of Intensity". *Journal of Political Economy*, 1983, Vol. 91, No. 5, pp 747-764.

Bills, David B, "Employers' Use of Job History Data for Making Hiring Decisions: A Fuller Specification of Job Assignment and Status Attainment". *The Sociological Quarterly*, 1990, 31 (1), pp. 23-35.

Brencic, Vera. "The impact of Craigslist's entry on competing employment websites". *IZA Journal of Labor Economics*, 2016, 5: 7.

Bruner, Jerome., "The Narrative Construction of Reality". *Critical Inquiry*, 1991, 18 (1), pp. 1–21.

Cappelli, P. 1999. The new deal at work: Managing the market-driven workforce. *Boston: Harvard Business School Press.*

Cejka M A and AH Eagly 1999. Gender-stereotypic images of occupations correspond to the sex segregation of employment. *Personality and Social Psychology Bulletin* 25(4):413-423.

Chandler, Dana, Steven D. Levitt, and John A. List 2011. Predicting and Preventing Shootings among At-Risk Youth *American Economic Review: Papers & Proceedings.* 2011, 101:3, 288–292.

Chan, Jaons and Jing Wang. 2016. Hiring Preferences in Online Labor Markets: Evidence of a Female Hiring Bias, forthcoming, *Management Science.*

Chernev, A. 2003. "When More Is Less and Less Is More: The Role of Ideal Point Availability and Assortment in Consumer Choice." *Journal of Consumer Research*, Vol. 30, September: 170-183.

Cornell, Bradford, and Ivo Welch. "Culture, Information, and Screening Discrimination", *Journal of Political Economy*, 1996, 104(3), pp. 542-71.

Cowgill, Bo. 2016. Productivity and Automation in Subjective Judgements: Theory and Evidence from Resume Screening. Working paper, Columbia University.

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., and Zamora, P., "Do labor market policies have displacement effects? Evidence from a clustered randomized experiment". *The Quarterly Journal of Economics*, 2013, 128(2), 531–580.

Cyronek, Travis, "Directed Search and Worker Application Strategies". 2016. Working Paper.

Dawes, R. M. 1971. A case study of graduate admissions: Application of three principles of human decision making. *American Psychologist*, 26, 180-188.

Dawes, R. M., Faust, D., and Meehl, P. E. (1989). Clinical versus actuarial judgment. *Science*, 243, 1668 1674. http://dx.doi.org/10.1126/science .2648573

Deming, David J., Noam Yuchtman, Amira Abulafi, Claudia Goldin and Lawrence F. Katz. 2016. "The Value of Postsecondary Credentials in the Labor Market: An Experimental Study." *American Economic Review*, 106(3): 778-806.

DeVaro, Jed and Oliver Gürtler, "Advertising and Labor Market Matching: A Tour through the Times". *Journal of Labor Economics* 36, no. 1, 2018, pp. 253-307.

DeVaul RA, Jervey F, Chappell JA, Caver P, Short B, O'Keefe S. 1987. Medical school performance of initially rejected students. *JAMA*. 1987 Jan 2;257(1):47-51.

Dunnette, M. D., Kirchner, W. K., and Erickson, J. R., "Predicting turnover of female office employees". *Personnel Administration*, 1960, 23, pp. 45-50.

Einhorn, J. H., and Hogarth, R. M., Confidence in judgment: Persistence of the illusion of validity. *Psychological Review*, 1978, 85, pp. 395-416.

Fernandez, R.M. and N. Weinberg, "Sifting and Sorting: Personal Contacts and Hiring in a Retail Bank". *American Sociological Review*, 1997, 62, December, pp. 883-902.

Edelman, D., "Freelancing in America Study: A National Survey of the New Workforce". Elance, Inc. 2014.

Gautier Pieter, Paul Muller, Bas van der Klaauw, Michael Rosholm, and Michael Svarer, "Estimating Equilibrium Effects of Job Search Assistance". *Journal of Labor Economics*, 2018, forthcoming.

Gomez-Uribe, Carlos A. and Neil Hunt, "The Netflix Recommender System: Algorithms, Business Value, and Innovation". *ACM Transactions on Management Information Systems*, 2015, 6 (4), Article 13.

Grove, William M., David H. Zald, Boyd S. Lebow, Beth E. Snitz, and Chad Nelson, "Clinical Versus Mechanical Prediction: A Meta-Analysis". *Psychological Assessment*, 2000, 12 (1), pp. 19-30.

Highhouse, Scott, "Stubborn reliance on intuition and subjectivity in employee selection". *Industrial and Organizational Psychology: Perspectives on Science and Practice*, 2008, 1, 333–342.

Hoffman, Mitch, Lisa B Kahn, and Danielle Li, "Discretion in Hiring," 2016. Working paper, Harvard University.

Horton, John J., "Supply constraints as a market friction: Evidence from an online labor market".

2016 Working Paper.

Horton, John J., “The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment”. *Journal of Labor Economics*, 2017, vol. 35, no. 2.

Imbens, Guido and Kalyanaraman, Karthik, “Optimal bandwidth choice for the regression discontinuity estimator”. *Review of Economic Studies*, 2012, 79 (3), pp. 933–959.

Iyengar, S. Sheena. and Mark Lepper. “When Choice is Demotivating: Can One Desire Too Much of a Good Thing?” *Journal of Personality and Social Psychology.* 2000, 76, pp. 995-1006.

Jeanneret, R., and Silzer, R. “An overview of psychological assessment”. In R. Jeanneret & R. Silzer (Eds.), *Individual psychological assessment: Predicting behavior in organizational settings*, 1998, pp. 3–26, San Francisco: Jossey-Bass.

Kahneman, Daniel, “Thinking Fast and Slow”. 2011 New York: Farrar, Strauss, and Giroux.

Kalev, A., F. Dobbin, and E. Kelly, “Best practices or best guesses? Assessing the efficacy of corporate affirmative action and diversity policies”. *American Sociological Review*, 2006, 71, pp. 589-617.

Katz, L.F. and A.B. Krueger, “The Rise and Nature of Alternative Work Arrangements in the United States, 1995-2015”. NBER 2016 Working paper.

Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. “Prediction Policy Problems”. *American Economic Review: Papers & Proceedings*, 2015, 105 (5), pp. 491–495.

Kroft Kory and Devin G. Pope, “Does online search crowd out traditional search and improve matching efficiency? Evidence from Craigslist”. *Journal of Labor Economics*, 2014, no. 32 pp. 259–303.

Kuhn, Peter and Hani Mansour, “Is Internet Job Search Still Ineffective?”. *The Economic Journal*, 2014, 124 (581), pp. 1213–1233.

Kuncel, Nathan, David Klieger, Brian Connelly, and Deniz Ones. “Mechanical Versus Clinical Data Combination in Selection and Admissions Decisions: A Meta-Analysis”. *Journal of Applied Psychology*, 2013, Vol. 98 (6), pp. 1060-1072.

Kuhn, Peter and Hani Mansour, “Is Internet Job Search Still Ineffective?”. *Economic Journal*, 2014, 124(581), pp. 1213–1233.

Leung, Ming D. “Dilettante or Renaissance Person? How the Order of Job Experiences Affects Hiring in an External Labor Market.” *American Sociological Review*, 2014, 79 (1), pp. 136-58.

Lewis, Kevin, “The limits of racial prejudice”. *Proceedings of the National Academy of Sciences of the United States of America*, 2013, 110 (47), pp. 18814-18819.

Lewis, E. C. and MacKinney, A. C. “Counselor vs. statistical predictions of job satisfaction in engineering”. *Journal of Counseling Psychology*, 1961, 8, pp. 224-230.

Lohr, Steve, “Creditworthy? Let’s Consider Capitalization”. *New York Times.* 2015, Jan 8, accessed Dec 8, 2016.

Meehl, P. E. 1954. “Clinical vs. statistical prediction: A theoretical analysis and a review of the evidence”. Minneapolis: University of Minnesota Press.

Morgan, John and Felix Várdy, "Diversity in the Workplace". *American Economic Review*, 2009, Vol. 99, No. 1, pp. 472-485.

Oyer, Paul and Scott Schaefer. " Personnel Economics: Hiring and Incentives, in the Handbook of Labor Economics", 4B, eds. David Card and Orley Ashenfelter, 2011, pp. 1769-1823.

O'Neil, Cathy. "Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy". 2016, Broadway Books.

Pager, Devah, "The use of field experiments for studies of employment discrimination: Contributions, critiques, and directions for the future." *The Annals of the American Academy of Political and Social Science*, 2007, 609 (1), pp. 104-33.

Pallais, Amanda, "Inefficient hiring in entry-level labor markets". *The American Economic Review*, 2014, 104 (11), pp. 3565–99.

Pellizzari, Michele, "Employers' Search and the Efficiency of Matching". *British Journal of Industrial Relations*, 2011, 49(1), pp. 25–53.

Prien, E. P., Schippmann, J. S., and Prien, K. O., "Individual assessment: As practiced in industry and consulting". Mahwah, NJ: Lawrence Erlbaum, 2003.

Rivera, L.A., "Hiring as Cultural Matching: The Case of Elite Professional Service Firms," *American Sociological Review*, 2012, 77 (6), pp. 999-1022.

Rogerson, Richard, Robert Shimer, and Randall Wright, "Search-Theoretic Models of the Labor Market: A Survey". *Journal of Economic Literature*, 2005, 43 (4), pp. 959–988.

Royce, W. S. and Weiss, R. L. "Behavioral dues in the judgment of marital satisfaction: A linear regression analysis". *Journal of Consulting and Clinical Psychology*, 1975, 43, pp. 816-824.

Rynes, S. L., Colbert, A. E. and Brown, K. G., "HR professionals' beliefs about effective human resource practices: Correspondence between research and practice". *Human Resources Management*, 2002, 41, 149–174.

Shoag, Daniel and Robert Clifford, ""No more credit score" employer credit check ban and signal substitution," Working Paper, 2016.

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

Stiglitz, J.E., "The theory of "screening", education, and the distribution of income", *American Economic Review*, 1975, 65 (3), pp. 283-300.

Thompson, R. E. A validation of the Glueck Social Prediction Scale for proneness to delinquency. *Journal of Criminal Law: Criminology and Police Science*, 1952, 43, pp. 451-470.

Tilcsik, András, "Pride and Prejudice: Employment Discrimination against Openly Gay Men in the United States." *American Journal of Sociology*, 2011, 117 (2), pp. 586-626.

Tversky, A. and Kahneman, D., "Judgments of and by representativeness". In D. Kahneman, P. Slovic & A. Tversky (Eds.), Judgment under uncertainty: Heuristics and biases, 1982, Cambridge, UK: Cambridge University Press.

Webb, S. C., Hultgen, D. D., and Craddick, R. A., "Predicting occupational choice by clinical and statistical methods". *Journal of Counseling Psychology*, 1975, 24, 98-110.

Wigdor, Alexandra, and Bert F. Green, Jr., eds., "Performance Assessment for the Workplace". Volume I (Washington, DC: National Academy Press, 1991.

Wormith, J. S., and Goldstone, C. S. "The clinical and statistical prediction of recidivism". *Criminal Justice and Behavior*, 1984, 11, pp. 3-34.

Varian, Hal R., "Computer mediated transactions", *American Economic Review*, 2010, 100 (2), pp. 1-10.

Yeomans, Mike, Anuj K. Shah, Sendhil Mullainathan, and Jon Kleinberg. "Making Sense of Recommendations", Harvard University Working Paper 2016.

Zheng, Alvin and Hong, Yili and Pavlou, Paul A., "Matching in Two-Sided Platforms for IT Services: Evidence from Online Labor Markets", 2016, Fox School of Business Research Paper No. 16-026.