I Z A Institute
of Labor Economics
Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 14622

Preference Signaling and Worker-Firm Matching: Evidence from Interview Auctions

Ron A. Laschever
Russell Weinstein

JULY 2021

I Z A Institute
of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 14622

Preference Signaling and Worker-Firm Matching: Evidence from Interview Auctions

Ron A. Laschever
Compass Lexecon

Russell Weinstein
University of Illinois and IZA

JULY 2021

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Preference Signaling and Worker-Firm Matching: Evidence from Interview Auctions*

We study whether there are improvements in worker-firm matching when employers and applicants can credibly signal their interest in a match. Using a detailed résumé dataset of more than 400 applicants from one university over five years, we analyze a matching process in which firms fill some of their inter- view slots by invitation and the remainder are filled by an auction. Consistent with the predictions of a signaling model, we find the auction is valuable for less desirable firms trying to hire high desirability applicants. Second, we find evidence that is consistent with the auction benefiting overlooked applicants. Candidates who are less likely to be invited for an interview (e.g., non-U.S. citizens) are hired after having the opportunity to interview through the auction. Among hires, these candidates are more represented among auction winners than invited interviewees, and this difference is more pronounced at more desirable firms. Finally, counterfactual analysis shows the auction increases the number and quality of hires for less desirable firms, and total hires in the market.

JEL Classification: J20, M51, D83

Keywords: labor markets, signaling, hiring, interview, matching

Corresponding author:
Russell Weinstein
University of Illinois at Urbana-Champaign
School of Labor and Employment Relations
Department of Economics
504 E. Armory Avenue
Champaign, IL 61820
USA

E-mail: weinst@illinois.edu

* We are grateful for helpful comments from Alex Bartik, Bo Cowgill, Peter Ganong, Andy Garin, Mitch Hoffman, Lisa Kahn, Kevin Lang, David McKenzie, Paul Oyer, Daniele Paserman, Johannes Schmieder, and seminar participants at Georgia State University, Loyola University Quinlan School of Business, University of Illinois at Urbana-Champaign, the LERA/ASSA meetings, the Midwest Economic Association/SOLE meetings, the Society of Labor Economists Annual Meeting, and the Wharton People and Organizations Conference. The opinions expressed in this article are the authors' own and do not reflect the views of Compass Lexecon.

1 Introduction

The way in which employers match with job applicants is central to understanding labor markets. Yet, this process largely remains a black box. Improving worker and firm outcomes through a more efficient matching process requires understanding the underlying frictions, as well as mechanisms that might reduce these frictions. We study worker-firm matching in a market with one such potential mechanism: allowing job applicants to credibly signal their interest in an employer through an auction for interviews.

Auctions for interview slots may address two important frictions in the matching process: uncertainty over applicant quality, and uncertainty over the likelihood that an applicant accepts an offer. Even if employers can successfully identify desirable applicants, there remains the challenge of identifying which candidates are truly interested in the job and would accept an offer with high probability. In recent years the cost of job applications has fallen as more postings and applications are online. This further raises the potential that applicants will have a low likelihood of accepting an offer. For example, the average Economics Ph.D. applicant in 2006-2008 applied to an average of 80 employers (Coles et al., 2010).[1] While applicants may try to signal preferences during the recruiting process, these actions are generally not costly and may be interpreted as cheap talk.[2]

Uncertainty over offer acceptance likelihood may lead to inefficiencies in matching, affecting the number and quality of matches in the market. The cost of identifying interested applicants may imply firms interview, and make offers to, candidates who are unlikely to accept. Alternatively, some firms may avoid making offers to desirable candidates, who are on average less likely to accept. As a result, these firms forego the opportunity to hire a desirable candidate who would have actually accepted an offer. Firms may also rely on other methods to identify desirable candidates likely to accept an offer, such as referrals or networks more broadly. However, use of these

[1] Unemployed job seekers on CareerBuilder sent an average of 13 job applications over a three month period (Marinescu and Rathelot, 2018). Faberman and Kudlyak (2019) find that people searching for hourly jobs on the SnagAJob website for more than a week, on average sent roughly 10 applications over an average five-and-a-half-week search spell. The average corporate job opening attracted 250 résumés (Glassdoor, 2015).

[2] Delong and Vijayaraghavan (2006) note that the investment banking firm SG Cowen offered informational interviews at their office prior to the formal hiring process, partly to determine which applicants were interested enough to take this opportunity.

methods may perpetuate labor market inequality, as some workers lack networks with high-quality jobs.

Though not common, there are a few markets in which all applicants have an equal opportunity to credibly signal their preferences for an employer. One example is the American Economic Association (AEA) job signaling mechanism, which allows candidates to send a signal of interest to two departments. Importantly, there is no requirement that employers interview the applicants sending the signal. In contrast, in our setting an employer is compelled to meet with some signaling job seekers.

A second example, and the focus of this paper, is the auction system used in the market for professional master's degree students, most commonly MBA students, at many top-ranked programs. These programs allow employers to choose some percentage of the applicants they interview, but require the remainder of the interview slots are allocated through an auction. Typically, firms first invite applicants for interviews, before applicants have had the opportunity to signal. Next, there is an auction for the remaining interview slots, and thus auction participants are students who were not invited for an interview by the firm. Each student is provided with an equal allotment of "bid points," and the auction winners are guaranteed interviews with the firm.

To the best of our knowledge, this is the first paper to study the impact of these auctions on initial hiring decisions.[3] Our setting provides a unique opportunity to learn about the underlying frictions in worker-firm matching. Additionally, it allows us to highlight potential benefits from broader implementation of these interview auctions, for example in online job application sites or local job centers.

To understand the expected impact of these auctions on applicants and firms, we build on the intuition of a basic signaling model. We introduce a stylized model in Section 2 to illustrate this intuition. Consider the case in which firm and worker quality are observable, but conditional on firm and worker quality, an applicant's acceptance likelihood is unobservable.[4] In this setting, the opportunity cost of signaling to a given firm is negatively correlated with the applicant's interest in the firm (as in the signaling model of Spence, 1973). Thus, on average, applicants who signal are

[3]Budish et al. (2017) study an MBA course allocation system that solves for the approximate competitive equilibrium for course allocation – finding a price for each course – and assigning students to schedules based on their reported preferences and an endowment of "fake money."

[4]More precisely, we assume firm and worker quality are observable up to an idiosyncratic noise over which neither firm nor employee can credibly signal.

more likely to accept an offer, relative to observationally-equivalent applicants who were invited for interviews, and did not have the opportunity to signal.

The main intuition is that these signals will be most valuable for less desirable firms trying to hire high desirability applicants.[5] Less desirable firms infer high desirability applicants are unlikely to accept an offer, but this rejection risk is lower if the applicant signaled. Under plausible conditions, the signal is less valuable for less desirable candidates and for high desirability firms, since acceptance probabilities are higher for these groups.[6]

We have three main predictions. First, the differential likelihood that a given firm makes an offer to, and hires, an *auction winner* relative to an *invitee* increases in applicant desirability, but less so for higher desirability firms. Second, this effect for less desirable firms should be increasing in the auction winner's bid value, as these are the auction winners most likely to accept an offer. Third, the difference in applicant desirability between hired invitees and hired auction winners increases with firm desirability. While this does not necessarily imply the auction yields a greater increase in underrepresented hires at high desirability firms, we predict that certain model extensions would yield this result, and we empirically test this.

Though the above discussion focused on the offer-acceptance signal, firms may hire auction winners for other reasons. Winning the auction allows candidates to reveal during the interview further information about their quality. For example, candidates may be able to discuss their international work experience and how it makes them a strong hire in ways employers did not consider *ex ante*. This alternative mechanism does not yield the three predictions above, but related to our third prediction this mechanism may be more important for higher desirability firms–firms less concerned about a low offer-acceptance likelihood.

Alternatively, enthusiasm for the job as expressed through the signal may imply higher match-specific quality, and higher productivity (for example through less costly effort), conditional on observable characteristics. This alternative mechanism also does not yield our main predictions. As such, we argue that if we find evidence for our three predictions, then uncertainty over acceptance likelihood is an important friction in this setting.

[5]This is similar to the predictions in Lee and Niederle (2015) and the intuition in Coles et al. (2010).

[6]This will be true if the gradient of a given employee's productivity across firm desirability is not too steep.

We test these predictions using detailed résumé data for over 400 job applicants enrolled in one professional master's degree program in the U.S. over five years. This is a specialized program, training students for a particular managerial function. There is a formal campus recruiting component, and a high placement rate with many students working at Fortune 500 companies. We focus on recruiting for internships, as this represents initial entry into this market after enrolling in the degree program. In addition, for over 50% of interns in our sample, their internship converts to a fulltime job with the same employer.[7]

As we describe in Section 3, unlike studies that only observe hires or even interviews for a particular firm, we observe the interview selection process of many firms considering the same group of applicants. We are interested in whether an applicant credibly revealing their interest influences hiring, relative to other similar applicants. Given the comparison group consists of applicants chosen by the firm, they may have higher unobservable quality. Our unique setting and data allow us to use the information on an applicant's interview invitations from other firms as a proxy for what would otherwise be unobservable quality. Rather than having to address issues of selection we directly observe the selection process.[8]

We find evidence consistent with our three predictions. Signals of interest are most valuable for less desirable firms trying to hire high desirability applicants. These firms are over five percentage points more likely—nearly twice as likely—to hire high desirability auction winners, relative to high desirability candidates they invited for an interview. As predicted, this differential does not exist among higher desirability firms, nor among less desirable applicants at less desirable firms. The effect for less desirable firms is increasing in the student's bid value, consistent with our second prediction.

We also find evidence consistent with our third prediction. The auction yields a differential increase in lower-desirability applicants' representation among hires at high desirability firms. This is consistent with high desirability firms hiring auction winners because of revelation of previously unobserved quality, rather than offer-acceptance likelihood. Non-U.S. citizens and applicants identifying as Asian are also

[7]Using data from a highly-ranked MBA program, Kuhnen and Oyer (2016) show firms use summer internships to learn whether students are a good fit for the industry.

[8]For example, see the discussion in Farber, Silverman, and von Wachter (2016) on the ability of résumé audit studies to "estimate important causal effects of worker characteristics that would be exceedingly difficult to estimate from observational data."

much more represented among hired auction winners than hired invitees. This evidence is consistent with the auction benefiting overlooked applicants.[9]

Our evidence suggests the auction represents an improvement in matching relative to the counterfactual in which firms choose all of their interviewees.[10] First, our main result shows the auction is not helping firms by simply identifying the applicants more likely to accept an offer. The auction identifies the high desirability applicants that are likely to accept an offer, and this is what is valued by the less desirable firms. Indeed our results show firms do not value signals from less desirable applicants, presumably because they already infer they have a high acceptance likelihood.

Additionally, we conduct simulations of offers and acceptances comparing scenarios with auctions to counterfactuals without auctions. Consistent with our prediction, the simulations show the auction is a way in which less desirable firms can target higher desirability applicants, while less dramatically increasing their risk of offer rejection. For example, in our simulations, the auction reduces the number of less desirable firms failing to hire by 50% relative to a scenario without auctions in which less desirable firms target high desirability applicants, while yielding on average only slightly lower quality hires. The auction also reduces total number of firms failing to hire in the market.

This paper contributes to a large literature on matching (see, e.g., Roth, 2018), and specifically preference signaling. Several theoretical papers study preference signaling in matching markets.[11]

Few papers empirically examine how matching is improved if applicants can credibly reveal their preferences over vacancies. Coles et al. (2010) find that signals sent through the AEA Economics Job Market mechanism increase the likelihood of obtaining an interview, and especially at liberal arts colleges. They hypothesize that

[9]It is possible that in the absence of auctions these applicants would be more represented among the additional invitees.

[10]Throughout, our discussion of a counterfactual considers the case where the number of interview slots remains the same, but none of the slots are filled via auction.

[11]Coles, Kushnir, and Niederle (2013) show theoretically that allowing applicants to signal to one employer increases the number of matches in the market, increases worker welfare, and has an ambiguous effect on firm welfare. Lee and Schwarz (2017) model firms' interview decisions in two-sided matching markets, and Lee and Schwarz (2007) model this when applicants can credibly signal preferences. Abdulkadiroglu, Che, and Yasuda (2015) show that allowing students to signal their preferences over schools in a centralized school choice market, and using these signals to break ties, leads to an increase in ex ante efficiency. Avery and Levin (2010) develop a model focusing on early applications to a university as a signal of enthusiasm for the college.

one reason may be liberal arts colleges have greater uncertainty over offer acceptance likelihood. Our paper also examines the channel of uncertainty over acceptance likelihood and we make several key contributions. First, we show the signal's differential value by firm characteristics also depends on applicant characteristics – consistent with a signaling model. Second, Coles et al. (2010) have data allowing them to examine interview outcomes, whereas we are able to examine the effect of signals on hiring. These data allow us to provide further confirmatory evidence that this is an important market friction, as well as evidence on the final outcome in the hiring process. Finally, the requirement that firms interview auction winners allows us to analyze whether this mechanism leads firms to hire candidates they initially passed over. This is not possible in the AEA market given that employers are not required to interview signaling applicants.[12]

Finally, our study of interviewees not selected by the firm relates to a growing literature on the use of algorithmic screening or testing in hiring (Cowgill and Tucker, 2018; Hoffman, Kahn, and Li, 2018; Li, Raymond, and Bergman, 2020).[13] More generally, our paper contributes to our understanding of who firms interview and hire, and why.

The paper proceeds as follows. In Section 2 we illustrate our conceptual framework with a stylized model and present the empirical predictions. Section 3 describes the data and the setting, and presents descriptive results. In Section 4 we test our first predictions that the auction is valuable for less desirable firms trying to hire high desirability applicants. Section 5 tests our third prediction and whether the auction improves representation of overlooked candidates among hires moreso at high desirability firms. Section 6 evaluates the impact of the auction on labor market

[12]Lee and Niederle (2014) find signals in an online dating site increase offer acceptances, and Avery, Fairbanks, and Zeckhauser (2003) find early-admission applicants to a university are more likely admitted, all else equal. Avery, Fairbanks, and Zeckhauser (2003) did not have full admissions-relevant information on applicants, and early application also screens based on financial need of the applicants (Kim, 2010). Horton and Johari (2018) show higher quality matches are formed, and more efficiently, when employers can signal the extent to which they are willing to pay for worker experience.

[13]Related work shows the impact on matching when providing employers with additional information on job applicants' quality or work experience. Groh et al. (2015) find little improvement in matching when providing employers with psychometric assessments of unemployed recent university and community college graduates in Jordan. Agrawal, Lacetera, and Lyons (2013) find that providing more information about an applicant in online labor markets improves the applicant's employment outcomes. Pallais (2014) finds a similar result, and evidence that this improves market-level employment and output.

efficiency, and Section 7 concludes.

2 Framework

2.1 Two-Firm Example

We describe in the introduction the general intuition for our predictions. To illustrate these predictions more formally, we introduce a simple two-firm model, further discussed in the online appendix. There are a total of K applicants, and applicants are of type H (high desirability) or L (low desirability), observable to firms and applicants. To focus on uncertainty over applicant acceptance likelihood, we treat wages as exogenous to the model. There are two firms, and all applicants prefer Firm A to Firm B (and Firm B is preferred to remaining unmatched). We also assume that the variability in match-specific productivity within H-type (and L-type) applicants is relatively small (the bounds on the variability are presented formally in the online appendix).

We consider the following game, consistent with our setting's timeline. Applicants submit applications to Firm A and Firm B. Firms review the applications and simultaneously choose one applicant to invite for an interview. Students then decide whether to participate in the auction for an interview with a firm, if they were not invited for an interview. The winner of the auction gets an interview at the firm, and ties in the auction are broken at random. Firms then interview their two applicants: one that they chose and the other who was the auction winner. After the interviews, firms make their offer simultaneously with their competitor. Finally, students decide which offers to accept.

Result: In equilibrium, Firm A hires an H-type applicant that they invited for an interview. Firm B hires an H-type applicant that won the auction.

We solve for the equilibrium in the appendix. However, the intuition is that Firm A will invite and make an offer to its preferred candidate, knowing it will accept. Firm B knows there is some risk associated with making an offer to an H-type candidate, as it may be the same candidate invited by Firm A. As a result, if an H-type candidate wins the auction for an interview at Firm B, it will always prefer that candidate–relative to an invited H-type candidate, as there is little heterogeneity in the match-specific components, and the invited H-type candidate may also be Firm A's invited candidate. As described in the appendix, this result depends on the

simplifying assumption that Firm A's invited H-type candidate does not participate in Firm B's auction, since she knows she will get an offer from her preferred firm. However, we conjecture this result would hold if this assumption were to be relaxed.[14] Because there is no rejection risk associated with making an offer to an invited L-type candidate, as this candidate will not receive an offer from Firm A, the signal would be less useful when coming from L-type candidates.

To summarize, Firm B prefers hiring the H-type auction winner over the H-type invitee, but would prefer an L-type invitee to an L-type auction winner. Firm A does not prefer hiring the H-type auction winner over the H-type invitee.

2.2 Predictions

The above result and discussion lead to the first prediction:

Prediction 1: The differential likelihood that a firm hires an *auction winner* relative to an *invitee* increases in applicant desirability, but less so for more desirable firms.

Our two-firm example is not sufficiently rich to fully capture our second and third predictions, but it is nonetheless illustrative.

Prediction 2: The differential likelihood that a given firm makes an offer to, and hires, an *auction winner* relative to an *invitee* increases in applicant desirability, *and this should be increasing in the student's bid value.*

This prediction does not depend on whether employers observe the bid value. If we had multiple B-type firms in our model, and a B-type firm made an offer to an auction winner with a lower bid, she would be less likely to accept the offer–because she had used more of her points at another firm. As a result, those most likely hired will be those with higher bids. Even if firms do not make multiple offers, they may learn about acceptance likelihood during an interview, better enabling them to target offers to those with higher acceptance likelihoods.

Prediction 3: The difference in applicant desirability between hired invitees and

[14] In a more general model, if H-type candidates invited by A-type firms still participate in auctions at B-type firms, intuition still suggests the signal is more valuable from H-type candidates. If a B-type firm sees an H-type auction winner, the firm knows the auction winner preferred it to other B-type firms where she could have used her signal. This is not true for an invited H-type candidate at a B-type firm. This signal is more valuable for H- than L-type applicants, since it is less likely that invited L-type candidates will receive offers from other B-type firms where they signal, compared to H-type candidates.

hired auction winners increases with firm desirability.

To see the intuition behind this prediction, recall that in our two-firm example Firm A invites an H-type applicant. In a richer model in which A-type firms sometimes hire H-type auction winners, for example because of additional match-specific productivity revealed during the interview, they would be the same type (H) as the hired invitees. In our two-firm example, we also show in the appendix that if the risk of making an offer to an invited H-type applicant is too high, Firm B will invite only L-type applicants. A model in which B-type firms sometimes preferred L-type auction winners to L-type invitees would imply that at B-type firms hired auction winners will sometimes include H-type applicants, whereas hired invitees will be a lower type. This implies Prediction 3 – the difference in applicant desirability between hired invitees and hired auction winners is less negative at A-type firms.

Further enriching the model to allow revelation of ability during interviews presumably would also imply the auction increases A-type firms hiring observable L-type candidates, who reveal they are type H during the interview. We conjecture that under plausible conditions the A-type firms will still invite only H-type candidates, and so their hired invitees will be of observable type H. As above, B-type firms may invite only L-type candidates. In this case, for L types along observable dimensions, the auction yields a differential increase in their representation among hires at A-type firms relative to B-type firms.

Extending to more than two firms makes this model much less tractable. However, we expect the general intuition to hold, as invited H-type candidates at B-type firms will have lower probability of accepting than invited L-type candidates.

3 Setting, Data, and Descriptive Results

3.1 Setting

We examine the job interviews and employment outcomes for students in a professional master's degree program in the U.S. The program takes about two years to complete, training students for a particular managerial field. As a result, in this setting signals are less likely used to show interest in switching fields, though that may be more common in more general master's programs in management where students enter a greater variety of fields. Many graduates are hired into corporate roles at Fortune 500 companies. Students typically are encouraged to have an internship during

their first summer in the program. Our data consist of both internship and fulltime applications, interviews, offers, and hires.

The meeting of applicants and employers is centralized and standardized in this market. The recruiting timeline for the semester is as follows. Employers register to formally recruit on campus, and they specify the number of interview slots they would like for each posting. Applicants then submit résumés. After reviewing the résumés, employers invite applicants to fill 50% of their interview slots.[15] After the employer's 50% of the slots are filled, and non-selected applicants are told they have not been selected, the interview auction takes place for the remaining 50% of the slots.

Applicants can only participate in the auction if they applied to the firm, and were not selected by the employer for an interview.[16] Each applicant is given an allotment of 1,000 points for the entire semester, and can allocate those points as they like. If an applicant's bid is not high enough to obtain an interview, their points are returned to their allotment.[17] The auction is conducted via software where the highest bids are selected, and ties are resolved at random. Employers do not observe the bid amount.[18]

3.2 Data

Our dataset contains two types of information. First, we have detailed information on each student compiled from several sources: their graduate school application, their profile on the job application system, their résumé, and surveys and records of employment outcomes. Second, we observe applications and interviews for nine semesters of recruiting, from Spring 2008 through Spring 2012.

The graduate school application provides information such as race and standardized test scores (e.g., GRE or GMAT). Profiles on the job application system include sex and citizenship status. Résumés for internship or fulltime recruiting provide the student's undergraduate institution, undergraduate major, undergraduate GPA, and

[15] Employers also select alternates who can fill these slots if their first-choice applicants decline the interview.

[16] If there are remaining slots after the auction, applicants can sign up through a free-for-all period.

[17] While we do not observe the auction dates, interview dates are distributed throughout the course of a month, and so it is reasonable that auctions follow a similar distribution. We discuss below some implications of these dynamics for the impact of the auction.

[18] The recruiting schedule given to the recruiters who arrive on campus does not designate those chosen by the employers, and those who obtained a slot via the auction.

employment history. Uniquely, we observe the résumé itself, not only the data associated with the résumé.

For each résumé-listed job, we classify the occupation using the five-digit Standard Occupational Classification (SOC) code. We determine which résumé-listed jobs were held prior to the start of the master's degree program. We then determine whether applicants had any prior experience in the occupation in which their master's program specializes. We also collect data on whether each pre-master's program employer is in The Fortune 1000, or The Global 2000 list of largest public companies from Forbes, in the year 2008.[19] All additional résumé information, such as awards and honors, leadership in university activities, participation in professional clubs, and volunteering, was also coded.[20]

When employers screen applicants to select interviewees they have the résumé and a short application. Thus, we observe the same information as the employers as well as additional information from the admissions record (e.g. standardized test score).

Data on employment outcomes include internship and fulltime employer, and salary for each of these positions.

The second type of information contains detailed matched job posting-applicant data. For each job posting-applicant pair, we observe whether the firm invited the student for an interview, whether the student participated in the interview auction, and whether the employer ultimately interviewed the student for the posting.[21]

From the administrative and survey data we observe whether the applicant received an offer or was hired.

Some employers post openings for more than one type of position, and so students may apply to one or several positions for each employer within a semester. We will refer to each employer's position as a job posting or an employer interview schedule.

[19]We do not code individuals as having Fortune 1000 or Global 2000 experience if their listed position at one of these companies was sales or food preparation. This suggests they may have been working in a particular establishment of a large retail or restaurant company as a waiter or retail clerk. We wish to distinguish this from corporate work experience.

[20]There are two résumés that were inadvertently not originally fully coded. In order to maintain consistency, we have opted not to retroactively code these résumés.

[21]We code individuals as obtaining their interview through an invitation if they were invited or they obtained their interview after being selected as an alternate by the firm. There are a small number of interviewees who were neither invitees nor auction winners, and presumably obtained their interview through the free-for-all period if remaining slots existed after the auction. We include these in our classification with auction winners as well, and they comprise approximately 4% of our sample of interviewees.

We focus on internship recruiting for several reasons. First, internship recruiting represents these students' entry into this market after enrolling in the degree program. Second, individuals participating in the formal fulltime recruiting program will be a selected sample. For over 50% of interns in our sample, their internship converts to a fulltime job. A nontrivial percentage of students who accept jobs at their internship employer do not apply to any other jobs posted through career services.[22]

Our data consist of 182 employer interview schedules (job postings), with an average of 20 per semester (Table 1).[23] Our sample consists of interview schedules in which at least one person received their interview through an invitation, and at least one through the auction. This comprises 83% of all campus recruiting interview schedules. We confirm that roughly 50% of the firm's interview slots are filled by invitations with the remainder filled by the auction (Table 1).[24] On average, approximately one applicant is hired per schedule, but a nontrivial proportion of job postings do not result in a hire. We examine the reasons behind this high rate of unfilled vacancies, despite the auction, in Section 6.1.

For students in this program, campus recruiting is an important part of their job search. Nearly 90% of all students hired for internships, regardless of whether this was through campus recruiting, apply to one of the job postings in our sample.[25] Partly since there are many more students than firms, 44% of these applicants are ultimately hired by a firm in our sample.

3.3 Employer and Job Seeker Heterogeneity and Differences between Invitees and Auction Winners

To test our predictions, it is of use to have a unidimensional measure of desirability both for employers and applicants. For applicants, we examine the total number of firms inviting a student for an interview as a proxy of the student's desirability. Firms

[22]The set of firms recruiting may also be selected, as these are firms that did not fulfill their fulltime needs through their interns. Firms may be better at screening applicants for fulltime jobs given the selected sample of applicants. Alternatively, they may apply different standards in screening or hiring for fulltime relative to internship recruiting, consistent with results from a study of the labor market at a prestigious MBA program (Kuhnen and Oyer, 2016).

[23]Recruiting is more concentrated in the Fall than in the Spring, and all of our results include year-by-semester fixed effects.

[24]Roughly 75% of the auctions in our sample have non-winning bids.

[25]More than 80% of our sample of applicants end up being hired for an internship through some method.

will unlikely have information on a candidate's actual number of invited interviews, but they may try to infer a student's desirability. For this reason, we use the predicted number of interview invitations for each candidate as a measure of an applicant's desirability. For student i applying to jobs in semester t, we estimate:

$$TotalInvites_{it} = X\beta + \gamma TotalApplications_{it} + \delta_t + \epsilon_{it}$$

where X includes student characteristics described in Table 2 and δ_t are year-by-semester fixed effects. We use $X\hat{\beta}$ as our measure of predicted total interview invitations, and we standardize it so it has mean zero and standard deviation of one among the sample of all applicants.[26]

Table 2, column 3 shows the coefficients from the regression. The predicted number of interview invitations is higher for non-Asian applicants, students with pre-Master's experience similar to their degree program, students with experience at Fortune 1000 or Global 2000 firms, high GPA students, students with volunteering experience on their résumé, and business majors. Interestingly, the coefficient on GRE/GMAT is not statistically significant from zero. This information is not observable to employers. The lack of significance implies these test scores do not have additional predictive content beyond that contained in the résumés.

Table 2 and Figure 1 show the characteristics of students obtaining their interview through an invitation and those obtaining their interview through the auction, and the differences are consistent with the regression predicting interview invitations. This suggests auctions are identifying a different pool of applicants with some observable characteristics, such as experience and academic quality, which are conceivably less desirable to firms. In Section 5, we test whether the auction increases representation of these candidates among hires, moreso at high desirability firms, consistent with our third prediction.

Despite these differences, we show below there is considerable overlap in the distribution of applicant desirability between invitees and auction winners. This is consistent with congestion in this market, and a greater number of potentially desirable applicants than interview slots (Figure 3a, Table 1). On average, firms invite roughly seven applicants for interviews, but receive 42 applications per job posting.

[26] We also construct the index using a Poisson regression with exposure equal to total applications per person, which yields similar results.

We find that winning an auction, conditional on participating, is not correlated with GMAT/GRE scores, unobservable to employers. This alleviates concerns that winning the auction is a signal of ability rather than a signal of acceptance likelihood.

As a unidimensional proxy for a firm's desirability we create an index measuring the firm's percentile in the semester's intern salary distribution, averaged over the nine semesters in our data.[27] Taking the average over many semesters avoids the concern that salary may depend on whether the hire was an auction winner or invitee. As we discuss below, as an alternative measure, we construct a firm desirability index based on the relative number of applications the firm received.

For much of the analysis, we use an indicator for whether the firm's desirability is at least equal to the median among firms recruiting that semester.[28]

Descriptive Evidence

Figure 2 shows the distribution of student desirability for auction winners is shifted to the left of that of invited applicants. However, roughly 40% of auction winners have desirability index around or above the mean. As such, we can test our first prediction, comparing the likelihood of hiring auction winners and invitees, conditional on applicant desirability.

Figure 2 also shows the distribution of desirability for invitees is quite similar for above- and below-median desirability firms. This suggests that these firms are inviting similar candidates. As our conceptual framework suggests, this will be particularly challenging for less desirable firms trying to hire high desirability applicants. Figure 3b shows more direct evidence that firms are concentrating interview invitations on a smaller group of applicants. Further, when a below-median desirability firm invites a high desirability candidate for an interview, the candidate has on average 3.3 in-

[27]Some students report base salary as well as bonus and relocation benefit separately. However, these are reported less consistently, and so we use only the base salary to construct the index. We observe salary only if the firm hires a student that semester, and many firms recruit without hiring. We calculate the firm's percentile in the salary distribution each semester, and then take the average across all semesters. This allows us to include firms that recruit without hiring in one semester, but do hire in another semester. As a robustness check, we create the same salary index but instead use the fulltime salaries reported by the firm's hires. The correlation between the two indices is roughly 0.7.

[28]In nominal dollars the average base internship salary between 2008 and 2012 is roughly $4800 per month. In real dollars, average monthly salary is $800 (23%) higher at high desirability firms relative to low desirability firms.

vitations from above-median desirability firms.[29] This suggests that low desirability firms face a high risk of rejection when making offers to high desirability invitees.[30] The similarity in invitations between low and high desirability firms is also consistent with our stylized model. We show that if the productivity difference between H- and L-type applicants is large enough, Firm B is indifferent between inviting H- or L-type applicants as it is guaranteed an H-type auction winner, and it will hire that applicant.

If offers were costless, identifying interested applicants would be less important, as firms could make offers until someone accepts. Two facts in our data suggest offers are costly. If offers were costless, when an offer is rejected we would expect firms to make another offer. However, we see that an additional rejected offer is associated with an increase of less than one total offer (Appendix Figure A1). Alternatively, firms with more rejected offers may have fewer vacancies and fewer total offers. Second, of the firms making offers in the subset of semesters for which we observe all offers, more than a quarter do not hire. Fifty percent of these make only one offer and 90% make one or two offers.[31]

Figure 4 shows descriptive evidence for our first prediction that lower desirability firms are more likely to hire high desirability auction winners relative to high desirability students they invited for interviews, but this is not true for the higher desirability firms.[32] Among lower desirability students, firms are more likely to hire invitees than

[29]Over 90% of high desirability invitees at below-median desirability firms have an invitation from an above-median desirability firm (Appendix Figure 3).

[30]Interviews take place throughout the month, and the typical spacing of the process's various stages (opening announced; followed by the period in which applicants submit resumes; etc.) suggest these are distributed similarly to interview dates. We also find in our data that, on average, less desirable firms are more likely to have interviews later in the month. This may be another way in which the auction is helpful for less desirable firms. It is credibly revealing interest at a time when students have good information about their other opportunities. At later points in the month, they know how many invitations they have had, and whether they have second round interviews. As a result, if someone signals interest later on, this may be more valuable and informative of their likelihood of accepting. Unfortunately, we lack sufficient precision to make conclusive statements about differential effects by interview week.

[31]If the cost of an unfilled vacancy is not very high, or firms believed they could fill the vacancy with high probability at another university, they may be less likely to make offers to non-invited candidates if they had slight preferences for invitees conditional on observables. This implies we would be less likely to find support for our predictions. We understand the number of other campuses is likely small given the number of other similar programs in terms of specific area and program reputation. Further, the likelihood of filling vacancies will not necessarily be higher at these other universities, as firms will be competing with other firms for applicants there as well.

[32]For the purposes of this figure, we define high desirability students as students with a desirability

auction winners, across all bins of firm desirability. These results are consistent with lower desirability firms benefiting more from a mechanism that identifies interested, high-quality applicants in a market where they face substantial competition. In the next section, we more formally test this relationship, including job-posting fixed effects.

4 The Differential Benefit of Auctions to Less Desirable Firms

4.1 Main Results

Prediction 1

We start by empirically examining the first prediction – the differential likelihood that a given firm makes an offer to, and hires, an *auction winner* relative to an *invitee* increases in applicant desirability, but less so for higher desirability firms.

We test this prediction with the following econometric specification on the sample of interviewed applicants:

$$\begin{aligned} Hire_{f_t s_t} = \beta_1 Auction_{f_t s_t} + \beta_2 Auction_{f_t s_t} * Q_{s_t} + \beta_3 Auction_{f_t s_t} * Q_{s_t} * Q_{f_t} \\ + \beta_4 Auction_{f_t s_t} * Q_{f_t} + \beta_5 Q_{f_t} * Q_{s_t} + \beta_6 Q_{s_t} + \gamma_{f_t} + \epsilon_{f_t s_t} \end{aligned} \tag{1}$$

The dependent variable indicates whether applicant s, applying to job posting f, posted in semester t, is hired by the firm. The variable Q_{s_t} refers to applicant desirability for applicant s in semester t, Q_{f_t} refers to desirability of the job posting f that is posted during semester t, and *Auction* is an indicator for an auction winner (relative to an invited interviewee). We include job posting (interview schedule) fixed effects (γ_{f_t}), which also capture any semester fixed effects as the job posting is specific to semester t. We test whether individuals obtaining interviews through the auction ($Auction_{f_t s_t}$) are differentially more likely to be hired relative to individuals invited by the firm for an interview, and how this varies with applicant and firm desirability. Our measure of Q_{f_t} is an indicator for whether the desirability of the firm who posted job f is above the median that semester, using the firm desirability measure defined

index greater than or equal to the mean among the sample of applicants. These plots are based on binscatter estimation, additionally adjusting for semester fixed effects and total number of students interviewed for that job posting.

in the previous section. We will refer to job postings as firms for short. The variable Q_{s_t} is a measure of student s's desirability during semester t, based on the predicted interview invitations received by the student.

The coefficient β_2 measures how the differential likelihood of hiring an auction winner changes with student desirability, for less desirable firms. The coefficient β_3 measures whether the differential likelihood of hiring an auction winner increases less in Q_{s_t} for more desirable firms. The coefficient $\hat{\beta}_2$ should be positive, and the coefficient $\hat{\beta}_3$ should be negative based on Prediction 1.[33] Our main results show equation (1) estimated separately for above- and below-median desirability job postings, and we test whether the coefficients on $Auction * Q_{s_t}$ are statistically different across these regressions. We report bootstrapped standard errors based on 400 replications, as Q_{s_t} is a generated regressor.[34]

In the first column of Table 3, we include both above- and below-median desirability firms, and find that the differential likelihood of hiring an auction winner increases with the student's desirability, but the effect is not statistically significant. We then estimate the regressions separately for low- and high desirability firms, to test Prediction 1.

Throughout the table, the coefficient on *Auction* is small and not statistically significant. This implies that for average desirability applicants, firms are similarly likely to hire auction winners and invitees. However, consistent with our prediction, the differential likelihood that less desirable firms hire auction winners increases with student desirability, and the effect is statistically significant at the five-percent level. For students with desirability one standard deviation above the mean, auction winners are 4.7 percentage points more likely to be hired than invitees. Consistent with our prediction, this is not true for higher desirability firms, and the coefficients on $Auction * StudentDesirability$ are statistically different for low and high desirability firms (columns 2 and 3, respectively).[35]

[33] Auction winners are a selected group in that they were not invited for an interview by the firm. If applicant quality is observable to firms, but partially unobservable to the researcher, the absence of a firm's invitation may imply lower unobservable quality. This should bias us towards finding auction winners are less likely hired than invited applicants, conditional on Q_{s_t}, but without obvious implications for β_2 and β_3.

[34] We conservatively report standard errors based on the standard nonparametric bootstrap, as the clustered bootstrap based on job postings – where we draw a sample of job postings with each bootstrap replication – yields slightly smaller standard errors on $\hat{\beta}_3$.

[35] We also estimate these regressions with student fixed effects, which adds an additional 390 fixed

Columns 4 and 5 are similar to columns 2 and 3, respectively, but include all components used to construct the desirability index as additional controls, rather than the index itself. This yields similar, though slightly larger, magnitudes. For students with desirability one standard deviation above the mean, less desirable firms are 5.5 percentage points more likely to hire auction winners than invitees ($p\text{-}value = .06$). This differential implies lower desirability firms are 1.8 times more likely to hire auction winners with desirability of one standard deviation above the mean, relative to invitees.[36] For students with desirability 1.5 standard deviations above the mean, firms are 7.8 percentage points more likely to hire auction winners than invitees ($p\text{-}value = .03$). Again, we do not see this differential among higher desirability firms.

Taken together, and consistent with our prediction, the auction helps low desirability firms to hire high desirability applicants. Invited candidates may have rejected offers, or these less desirable firms may not make them offers given the perceived risk of rejection. These risks are lower among applicants who used their costly signal. Higher desirability firms value this attribute of the auction much less, given they perceive higher offer acceptance rates among invited applicants.

The auction is not simply providing value by identifying applicants with fewer opportunities, which could conceivably be accomplished easily by firms in a counterfactual environment without auctions. Our results show the auction is providing value by identifying the high desirability applicants who are likely to accept an offer. These applicants may be much more challenging for firms to identify on their own, without an auction.

The magnitudes suggest high desirability firms are less likely to hire auction winners at most levels of applicant desirability. As a result, these firms may prefer abandoning the auctions, since it is requiring them to interview candidates they are less likely to hire.

---

effects. This yields similar results though they are not statistically significant. The differential coefficient on $Auction * StudentDesirability$ in high desirability firms is -.04 (relative to -.055 in columns 2 and 3).

[36]The average likelihood of being hired by a lower desirability firm is .07, for invitees with desirability .5 to 1.5 standard deviations above the mean.

Prediction 2

Next, we test our second prediction – the differential likelihood that a given firm makes an offer to, and hires, an *auction winner* relative to an *invitee* increases in applicant desirability, *and this should be increasing in the student's bid value.* To test this prediction, we estimate (1) and include interactions with the student's bid value. We test this prediction using data on bid values available for six out of nine of our semesters.

Table 3, Column 6 shows that for less desirable firms, the differential value of the auction among high desirability candidates is increasing in the amount the student bid in the auction (coefficient on Auction*Student Desirability*Bid Value). For students with desirability one standard deviation above the mean, auction winners who bid one standard deviation above the mean bid of winners, are 14 percentage points more likely to be hired relative to invitees of the same desirability ($p\text{-}value \leq .05$). Consistent with our prediction, this differential is significantly smaller for auction winners with lower bids, and also not significant for higher desirability firms.

Together, the results in this section show that consistent with Prediction 1, the auction provides value by helping less desirable firms to hire high desirability applicants who are likely to accept. Consistent with Prediction 2, we find that these firms are most likely to hire the auction winners who send the strongest signals of their acceptance likelihood. This evidence suggests that identifying high desirability applicants who are likely to accept an offer is an important friction in this labor market. These applicants exist, and the auction makes it easier to identify them.

Robustness of the results

While salary is one measure of firm desirability, it may not capture other relevant measures of desirability unrelated to compensation. As a robustness check, we use the number of applications a firm receives as a measure of firm desirability. Specifically, we create an index analogous to the salary index, which is the firm's average percentile in the applications distribution across all semesters in the data. Appendix Table A4 shows similar results when using this applications index as our measure of firm desirability. The differential likelihood of hiring an auction winner increases with applicant desirability, but statistically significantly less so for firms that get more applicants ($p\text{-}value \leq .05$).

As a second robustness check, we re-estimate specification (1), but define high and low desirability firms instead using an index based on fulltime salaries. Appendix Table A3 columns 1 to 4 show evidence consistent with Prediction 1, though not statistically significant.[37] We also estimate specification (1), using the firm's desirability (average percentile in the salary distribution) rather than an indicator for above-median desirability. Appendix Table A3 columns 5 to 7 show this also yields a pattern consistent with Prediction 1 – the differential likelihood of hiring an auction winner increases with student desirability substantially less for higher desirability firms. However, this difference is not statistically significant.

We note that our sample years (2008-2012) are during and immediately after the Great Recession in 2007-2009. If firms are less interested in hiring during this time, then high desirability invitees may have fewer other offers than they would during an expansion. This may imply less desirable firms face lower rejection risk from high desirability invitees, thus reducing the value of the signal. As a result, we may find a lower value of the signal than we would if our sample period coincided with an economic expansion.

4.2 Offers

The framework discussed in Section 2 implies Prediction 1 should be true for offers, in addition to hires, if the cost of making an offer is high enough. If less desirable firms are restricted in the number of offers they would like to extend, perhaps only interested in making one offer, they should be more likely to make that offer to a high desirability auction winner relative to a high desirability invitee. On the other hand, if offers are costless, then less desirable firms may be equally likely to make offers to high desirability invitees, but since those candidates are more likely to reject such an offer, the firm is more likely to hire high desirability auction winners.

We use additional survey data that contains information on all offers received by students, including rejected offers. Unfortunately, these data are only available for four of our nine semesters, and so the sample size is greatly reduced.[38] Appendix Table

[37] There are some firms who are missing the fulltime salary index, but not the intern salary index, since they never hire for a fulltime job in the data. Appendix Table A3 shows results both excluding these firms, and including them by imputing whether they were a high desirability firm based on their intern salary index.

[38] Roughly 56% of student/semester observations report more than one offer, conditional on reporting at least one offer.

A1 shows evidence for Prediction 1 when considering offers: the differential likelihood of making an offer to an auction winner increases with applicant desirability, but less so for more desirable firms. This is consistent with less desirable firms valuing signals from high desirability applicants, and offers being costly.

If less desirable firms are targeting offers to individuals more likely to accept, conditional on student desirability, there may be little difference in the offer acceptance likelihood between auction winners and invitees. There are no statistically significant differences in acceptance likelihoods, though we are underpowered with only about 70 observations receiving offers from less desirable firms (restricting to offers and to seasons where we observe offers; Appendix Table A5). If higher desirability firms are making offers with less concern for acceptance likelihood, acceptance likelihoods should be higher for auction winners than invitees. Indeed, the average desirability auction winner is more likely to accept than the average desirability invitee at more desirable firms ($p\text{-}value \leq .05$), and the magnitude suggests this increases with student desirability, but it is not statistically significant.

We have no clear prediction that retention should differ for auction winners relative to invitees, conditional on being hired as an intern and on applicant desirability. There are several model extensions that might generate the prediction that hired auction winners have higher retention than hired invitees.[39] Because this is an important empirical question, we test this but we cannot reject that hired auction winners are equally likely to stay for a fulltime job as hired invitees; however, this may be driven by our smaller sample size (Appendix Table A5).[40]

5 Representation of Overlooked Candidates Among Hires

In this section we test the third prediction that the difference in applicant desirability between hired invitees and hired auction winners increases with firm desirability.

We test this with the following specification, restricted to hires:

[39]For example, differences in acceptance likelihood during intern recruiting may be magnified during fulltime recruiting, and later in one's career, as individuals have more time to generate other offers.

[40]Because we do not have fulltime job outcomes for the last two cohorts in our data, our sample size is only 136 observations.

$$(Q_{s_t}|Hire_{f_t s_t} = 1) = \beta_1 Auction_{f_t s_t} + \beta_2 Auction_{f_t s_t} * Q_{f_t} + \beta_3 Q_{f_t} + \gamma_t + \epsilon_{f_t s_t}$$

where as above, Q_{s_t} is the student's desirability measure, Q_{f_t} is an indicator for whether the firm's desirability measure is above the median that semester, and γ_t are semester fixed effects. Prediction 3 implies the coefficient $\hat{\beta}_2$ should be negative. If $\hat{\beta}_1 + \hat{\beta}_2 < 0$, this implies that at high desirability firms hired auction winners have lower desirability measures than hired invitees. In other words, high desirability firms are hiring different types of candidates through the auction – applicants with lower ex-ante desirability. As we discussed, this would be consistent with the auction allowing some individuals with lower desirability measures to reveal they are high desirability candidates.

We find evidence consistent with Prediction 3. Auction winners have lower desirability measures than invitees, and this is similarly true among low and high desirability firms (Table 4, column 1; Figure 5). However, among hires, auction winners have similar desirability as invitees at lower desirability firms (column 2, coefficient on Auction; Figure 5). This shows that less desirable firms invite high desirability applicants, but they do not hire them. At higher desirability firms, the desirability of hired auction winners is 0.6 standard deviations below that of hired invitees, and this difference is statistically more negative than at lower desirability firms (column 2, coefficient on Auction*High Q Firm). This evidence is consistent with the information revelation story above–when high desirability firms hire auction winners they include observable L-types who showed themselves to be H-types during the interview.

Table 4, Panel B shows results where the dependent variables are the components of the student desirability index that appeared as the most significant predictors of the number of interview invitations (in Table 2, column 3). While precision is lower, we see similar patterns. At less desirable firms, hired auction winners have similar experience at prestigious companies as hired invitees, as well as similar likelihood of having a high undergraduate GPA. However, at high desirability firms hired auction winners are less likely to have this prestigious experience than hired invitees, and are less likely to have a high undergraduate GPA. However, the differences relative to less desirable firms are not significant.

One notable applicant characteristic for which we do not find evidence of Predic-

tion 3 is the case of applicants identifying as Asian, suggesting differences in predicted interviews between Asian and non-Asian applicants (Table 2) are not affecting hires as our framework predicts. The results in Table 4 show that at lower desirability firms, hired auctions winners are more likely Asian than hired invitees, and we cannot rule out the same at high desirability firms. When these candidates have an opportunity to interview with the firm because of the auction, it does result in hires, and yields a different composition of hired auction winners and hired invitees.[41] We do not see similar effects when restricting to firms that only allow applications from U.S. citizens (Appendix Table A2). Our findings are consistent with auctions allowing employers to learn about foreign credentials and work experience with which they may be less familiar, and consistent with employers engaging in statistical discrimination rather than "taste-based" discrimination.[42]

To understand whether there are similar effects for non-US citizens, we focus on job postings that allow applications from non-US citizens. Roughly 80% of the postings in our sample are restricted to U.S. citizens or permanent residents, while 20% also specify that students with an F-1 Visa may apply.[43] Given that auctions cannot benefit non-U.S. citizens in these restricted postings, we focus on the non-restricted postings.[44] We do not have power to estimate these specifications separately for low- and high-desirability firms.

Auction winners for jobs without citizenship restrictions are 31 percentage points less likely to be U.S. citizens than invited applicants (Table 4).[45] Among hires, auction

[41]The similarity of the coefficients in the interviewees and hires regressions suggests the compositional difference in hires is explained by the compositional difference in interviewees – Asian applicants' have higher representation in the pool of auction winners than in the pool of invitees. As a result, this seems more likely than an alternative explanation in which Asian auction winners revealed themselves to be better matches than Asian invitees.

[42]See, for example, the discussion on statistical versus taste-based discrimination in Oreopoulos (2011), who finds lower callback rates for applicants with foreign sounding names in a résumé audit study.

[43]This information is missing for one of the semesters, comprising about 8% of the postings in our main sample. Approximately 12% of restricted postings specify they are open only to US citizens, potentially because this is required by law or contract.

[44]Compliance with these application restrictions is nearly perfect; only 0.6% of applications to restricted postings are from students who are neither U.S. citizens nor permanent residents.

[45]Given that some of the restricted postings are open only to U.S. citizens, we continue to use an indicator for U.S. citizenship, which is zero for permanent residents. Only 3% of the sample of applicants are permanent residents, and using an indicator for citizen or permanent resident yields very similar results.

winners are 43 percentage points less likely to be U.S. citizens than invited applicants. Firms hire non-U.S. citizens upon having the opportunity to interview them, again consistent with the auction allowing firms to learn about foreign experience and credentials with which they may be less familiar.

The results in this section show that less traditional candidates are more represented among hired auction winners than hired invitees at high desirability firms. We also find non-U.S. citizens and those identifying as Asian are more represented among hired auction winners than among hired invitees (not limiting to high desirability firms).

6 The Effect of Auctions on Labor Market Efficiency

This section evaluates the effect of interview auctions on unfilled vacancies and quality of hires. First, we analyze the likelihood a firm hires from their interview pool, and whether firms with lower hiring rates might benefit from greater use of the auction. Second, we compare outcomes when an auction is in place to a counterfactual in which firms invite all of their interviewees.

6.1 Are Firms Using the Auction Optimally?

Our conceptual framework suggests that low desirability firms may be able to hire desirable applicants by making offers to auction winners and that this is less necessary for high desirability firms. In our framework, we assumed that firms know whether they are high desirability firms. However, if offers are costly, failure to hire may be prevalent for lower desirability firms to the extent they mistakenly believe and act as if they are higher desirability firms. We find that firms in the middle of the desirability distribution have the lowest likelihood of hiring from their interview pool (Appendix Figure A2a). These are the firms we would expect to have the most uncertainty over their place in the distribution. Our evidence is consistent with these firms behaving as though they are of higher desirability than their actual desirability level.[46]

Further consistent with these middle-desirability firms behaving as though they are high desirability firms, they are less likely to hire invitees, relative to both more

[46]The plot in Appendix Figure A2a is based on binscatter estimation, adjusting for total applicants interviewed by the firm. This addresses the possibility that medium-desirability firms interview more applicants, which could explain the lower hire likelihood. The figure also controls for semester fixed effects.

and less desirable firms (Appendix Figure A2b). However, relative to less desirable firms, they are more likely to make offers to invitees relative to auction winners, and to invitees with higher desirability (Appendix Figure A2c, A2d). Taken together, the results suggest medium-desirability firms are getting rejected by high desirability invitees, and would benefit from using the auction results to a greater extent.

6.2 Comparing Match Rate and Hiring Quality With and Without Auctions

Our framework suggests auctions help less desirable firms to hire high desirability applicants. In the absence of auctions, less desirable firms may choose between two types of strategies. They may make offers to high desirability applicants and face a lower likelihood of filling the vacancy. Alternatively, they may reduce their risk of offer rejection by targeting less desirable candidates. To further study this question and the role of the auction, we simulate offers and acceptances under both sets of strategies. We then compare the match rate and the quality of the hires from these counterfactuals without auctions, where firms invite all candidates, to a scenario with auctions. Specifically, we compare three scenarios: Scenario A, where auctions exist, and two scenarios without auctions, B1 and B2, which differ in our assumption on the behavior of low-desirability firms.

First, for Scenarios B1 and B2, counterfactuals without auctions, we identify counterfactual invitees as the nearest-neighbor matches to each of the firm's invitees. To compare to Scenario A, we constrain the number of nearest-neighbor matches for a given posting to be equal to the number of auction winners for a given posting.[47]

Next, we generate firm preferences over the candidates in their interview list. We assume firms agree that the best students are those with above-median predicted interview invitations, based on all applicants. In the scenarios without the auction, firms have random preferences over students within these blocks of applicant desirability.

[47] In some cases the number of invitees was greater than auction winners, and we drop the nearest-neighbor matches to the lowest desirability invitees. In other cases the number of invitees was less than the number of auction winners, and we add the highest desirability applicants still in the pool. It is possible that the availability of the auction affects interview choices, and that with the auction less desirable firms invited more desirable candidates than they would have absent the auction. However, we do not expect our choice of counterfactual invites will have a large effect on the results, given the way in which we specify offer strategies.

In Scenarios B1 and B2, our non-auction counterfactuals, we assume high desirability firms make their first offer to the highest ranked high desirability candidate on their interview list. But whereas we assume in B1 that low desirability firms make their first offer to the highest ranked of the high desirability applicants, we assume in B2 that low desirability firms make their first offer to the highest ranked of the low desirability applicants (to avoid competing with high desirability firms).[48] While there may be an equilibrium without auctions in which some less desirable firms make offers to high desirability and some to low desirability candidates, this should lead to an intermediate outcome between the outcomes we present. For simplicity we focus on results assuming less desirable firms all make offers to the same type of applicants in Scenarios B1 and B2.

We assume students agree on high and low desirability firms based on their average percentile in the salary distribution, but with independent and uncorrelated preferences within those blocks.[49]

Finally, for Scenario A, we use the actual invites and auction winners on the firm's interview schedule. Based on the prediction from our conceptual framework, we assume that less desirable firms first rank the high desirability auction winners, then high desirability invitees, and then low desirability interviewees regardless of whether they were auction winners or invitees. We assume high desirability firms first rank the high desirability invitees, then high desirability auction winners, and then low desirability interviewees. We assume that within a block of firm quality, students prefer the firms where they won the auction.

Simulations

We conduct 250 simulations of the auction and non-auction counterfactuals. In each simulation, we keep the interviewed applicants fixed, but we draw different firm pref-

[48] Assuming that a given firm must pay the same wage to a new hire, regardless of their desirability, there will be some parameters for which all high desirability firms make offers to the highest ranked high desirability candidate on their list. Specifically, this will be true if there is a large enough difference in the productivity of high desirability relative to low desirability candidates, and the number of high desirability firms relative to high desirability applicants is not too high. If this is true, then in the scenarios without auctions, less desirable firms will know there is an upper bound on the probability of hiring a high desirability applicant.

[49] Coles, Kushnir, and Niederle (2013) model a pure coordination game in which firms have independent and uncorrelated preferences, and students agree on high- and low desirability firms but have independent and uncorrelated preferences within those blocks.

erences over the applicants and different student preferences over the firms. Firms all make offers at the same time. If a student receives multiple offers, she accepts the offer from her highest-ranked firm. If a firm's offer is rejected in this first round, we allow for a second round of offers.

If the firm did not make any offers in the data, we do not allow them to make offers in the counterfactual. Because of this, we restrict to the semesters with information on all offers received.[50]

The auction (Scenario A) reduces the number of less desirable firms that fail to hire by more than 50%, relative to the scenario without auctions B1 (Figure 6 and Appendix Table A6).[51] There is also a 16% reduction in firms failing to hire relative to the scenario without auctions B2, in which less desirable firms make offers to less desirable candidates. This likely reflects two attributes of the auction. First, the auction addresses a coordination problem by distributing applicants across firms. Second, the auction identifies applicants more likely to accept an offer as they have used their costly signal at the firm.

Consistent with these potential mechanisms, the auction also reduces the number of high desirability firms that fail to hire. And consistent with our prediction, the reduction is much smaller than the reduction for less desirable firms targeting high desirability applicants. These results show the auction is reducing the total number of firms failing to hire in the market.

Less desirable firms are less likely to hire in Scenario B1, but conditional on hiring the mean quality of the hire is larger by 0.2 standard deviations relative to Scenario A. However, relative to the scenario without auctions B2, in which less desirable firms make offers to less desirable candidates, the quality of the hire is larger by 0.7 standard deviations in the scenario with auctions. We see this as an underestimate of

[50]We make several additional adjustments. If we observe firms hiring multiple candidates in the data, we allow them to make that many offers in the first round. If a student receives an offer from a firm that is not participating in this formal recruiting process, we assume it arrives simultaneously with the other offers.

[51]For comparison, we implement these simulations assuming students are indifferent among firms in the same block of desirability, regardless of whether they were invited or won the auction. This yields a 39% reduction in less desirable firms that fail to hire, smaller than the 54% reduction when assuming students prefer firms where they won the auction. The reduction in firms failing to hire, even when assuming students are indifferent, may reflect that the auction reduces coordination problems by distributing applicants across firms. However this may also reflect an advantage from making offers to auction winners, if they are on fewer interview schedules than similar desirability invitees.

the gains from the auction, because conditional on observable quality, auction winners may be higher productivity workers due to idiosyncratic fit with the firm.

Our framework suggests auctions will help less desirable firms hire high desirability applicants. In the absence of the auction, these firms would face a high rejection rate if they made offers to high desirability candidates. If instead they made offers to less desirable candidates to increase their likelihood of hiring, this would reduce the quality of hires. Our counterfactual shows these improvements, in likelihood of hiring and quality of hires, are large.

7 Conclusion

Frictions in the matching of workers and firms can lead to suboptimal outcomes, with consequences for individuals, firms, and markets. Understanding mechanisms that can reduce these search frictions is important for improving worker and firm outcomes. We study a labor market with interview auctions, a mechanism that allows workers to credibly reveal their interest in a particular position.

We test three main predictions from a conceptual framework that builds on a signaling model. Using data on over 400 applicants from one degree program over five years, our main finding is that auctions are valuable to less desirable firms trying to hire high desirability candidates. These firms are over five percentage points more likely to hire desirable auction winners relative to the candidates they invite, nearly doubling the likelihood. We also show that the differential likelihood of hiring auction winners increases with the bid of the auction winner, even though employers do not observe the actual bid values. If these students are difficult to identify in the pool of applicants, this suggests a counterfactual in which firms chose all invitees would not yield similar results. Auctions appear to be providing firms with additional information, helping them to identify a pool of high desirability workers that is more likely to accept offers.

Second, and consistent with our prediction, we find greater representation of overlooked candidates among hired auction winners than among hired invitees, and moreso among high desirability firms. Additionally, we see greater representation of applicants identifying as Asian and non-U.S. citizens among hired auction winners than among hired invitees more broadly, rather than just at high desirability firms. Finally, our simulations show substantial improvements in likelihood of hiring and in

quality of hires relative to markets without auctions. The number of less desirable firms failing to hire falls by 50% with an auction, relative to a scenario without auctions in which less desirable firms target high desirability applicants. The auction also reduces the total number of firms failing to hire in the market.

The results suggest that introducing mechanisms that allow applicants to credibly signal their interest could help improve worker-firm matching, and improve outcomes for applicants who tend to be overlooked in résumé screening. We believe this has potentially important implications for online job posting and job application sites, as well as local job centers.[52]

References

[1] Abdulkadiroglu, Atila, Yeon-Koo Che, and Yosuke Yasuda (2015): "Expanding 'Choice' in School Choice," *American Economic Journal: Microeconomics*, Vol. 7(1).

[2] Agrawal, Ajay, Nicola Lacetera, and Elizabeth Lyons (2016): "Does Standardized Information in Online Markets Disproportionately Benefit Job Applicants from Less Developed Countries," *Journal of International Economics*, Vol. 103 (November).

[3] Avery, Christopher, Andrew Fairbanks, and Richard Zeckhauser (2003): The Early Admissions Game: Joining the Elite, Cambridge: Harvard University Press.

[4] Avery, Christopher and Jonathan Levin (2010): "Early Admissions at Selective Colleges," *American Economic Review*, Vol. 100.

[5] Budish, Eric, Gerard P. Cachon, Judd B. Kessler, and Abraham Othman (2017): "Course Match: A Large-Scale Implementation of Approximate Competitive Equilibrium from Equal Incomes for Combinatorial Allocation," *Operations Research*, Vol. 65(2).

[6] Coles, Peter, John Cawley, Phillip B. Levine, Muriel Niederle, Alvin E. Roth, and John J. Siegfried (2010): "The Job Market for New Economists: A Market Design Perspective," *Journal of Economic Perspectives*, Vol. 24(4).

[52] For example, an online platform with job postings could provide users with two opportunities per month to send a signal to employers.

[7] Coles, Peter, Alexey Kushnir, and Muriel Niederle (2013): "Preference Signaling in Matching Markets," *American Economic Journal: Microeconomics*, Vol. 5(2).

[8] Cowgill, Bo and Catherine Tucker (2020): "Algorithmic Fairness and Economics," *Columbia Business School Research Paper*, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3361280.

[9] Delong, Thomas and Vineeta Vijayaraghavan (2006): "SG Cowen: New Recruits," Harvard Business School No. 9-402-028, Boston, MA. Harvard Business School Publishing.

[10] Faberman, Jason and Marianna Kudlyak (2019): "The Intensity of Job Search and Search Duration," *American Economic Journal: Macroeconomics*, Vol. 11(3).

[11] Farber, Henry S., Dan Silverman, and Till von Wachter (2016): "Determinants of Callbacks to Job Applications: An Audit Study," *American Economic Review: Papers & Proceedings.* Vol. 106(5).

[12] Glassdoor (2015): "50 HR and Recruiting Stats That Make You Think," Accessed Online 8/23/2019, https://www.glassdoor.com/employers/blog/50-hr-recruiting-stats-make-think/.

[13] Groh, Matthew, David McKenzie, Nour Shammout and Tara Vishwanath (2015): "Testing the Importance of Search Frictions and Matching Through a Randomized Experiment in Jordan," *IZA Journal of Labor Economics*, Vol. 4:7.

[14] Hoffman, Mitch, Lisa B. Kahn, and Danielle Li (2018): "Discretion in Hiring," *Quarterly Journal of Economics*, Vol. 133(2).

[15] Kim, Matthew (2010): Early Decision and Financial Aid Competition Among Need-Blind Colleges and Universities," *Journal of Public Economics*, Vol. 94(5-6).

[16] Kuhnen, Camelia and Paul Oyer (2016): "Exploration for Human Capital: Evidence from the MBA Labor Market," *Journal of Labor Economics*, Vol. 34(2).

[17] Lee, Soohyung and Muriel Niederle (2015): "Propose with a Rose? Signaling in Internet Dating Markets," *Experimental Economics*, Vol. 18.

[18] Lee, Robin and Michael Schwarz (2007): "Signalling Preferences in Interviewing Markets," *Dagstuhl Seminar Proceedings*. Dagstuhl, Germany.

[19] Lee, Robin and Michael Schwarz (2017): "Interviewing in Two-Sided Matching Markets," *RAND Journal of Economics*, Vol. 48(3).

[20] Li, Danielle, Lindsey Raymond, and Peter Bergman (2020): "Hiring as Exploration," Working Paper.

[21] Marinescu, Ioana and Roland Rathelot (2018): "Mismatch Unemployment and the Geography of Job Search," *American Economic Journal: Macroeconomics*, Vol. 10(3).

[22] Oreopoulos, Philip (2011): "Why do Skilled Immigrants Struggle in the Labor Market? A Field Experiment with Thirteen Thousand Resumes," *American Economic Journal: Policy*, Vol. 3(4).

[23] Pallais, Amanda (2014): "Inefficient Hiring in Entry-Level Labor Markets," *American Economic Review*, Vol. 104(11).

[24] Roth, Alvin E. (2018): "Marketplaces, Markets, and Market Design," *American Economic Review*, Vol. 108(7).

[25] Spence, Michael (1973): "Job Market Signaling," *Quarterly Journal of Economics*, Vol. 87(3).

Figure 1: Characteristics of Auction Winners Relative to Invitees

(a) Undergraduate GPA

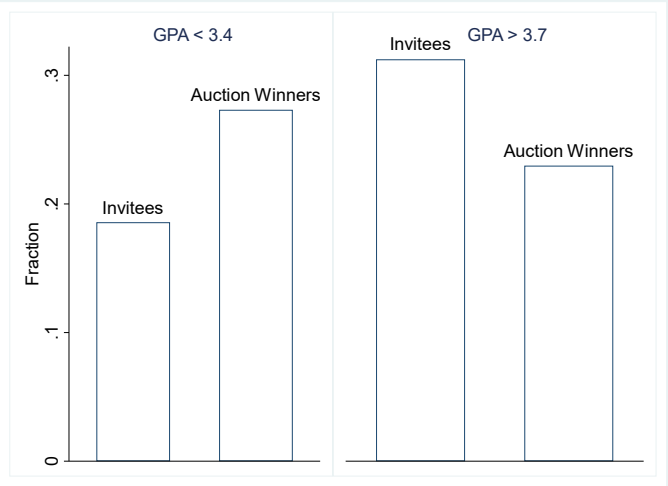

(b) Pre-Master's Experience at Fortune 1000/Forbes Global 2000 Company

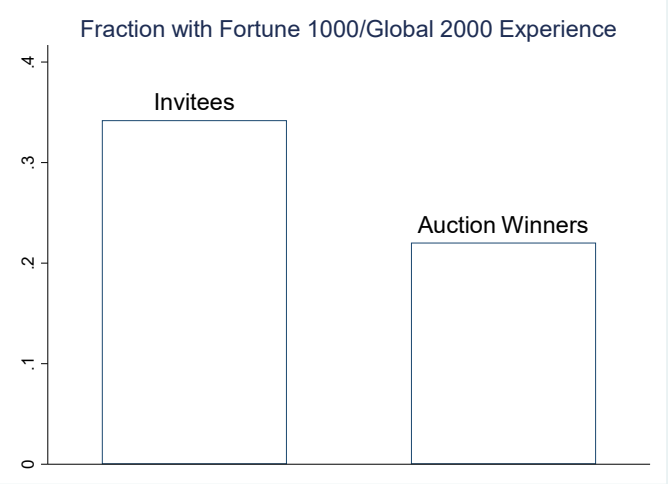

(c) Race/Origin

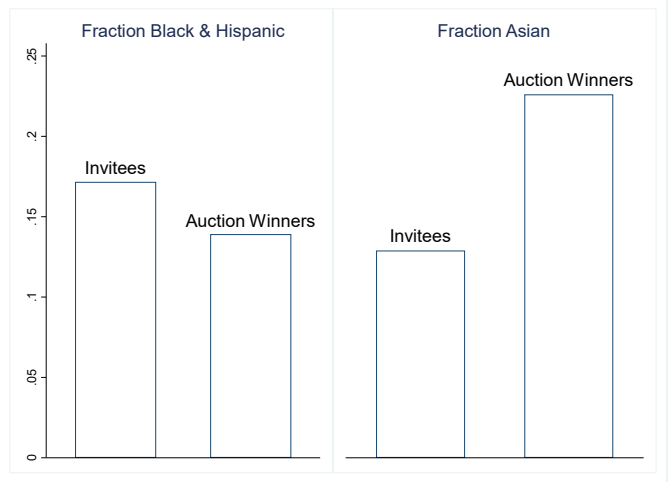

Notes: Plots show the fraction of invitees and auction winners with low and high GPA, previous experience at a Fortune 1000/Forbes Global 2000 company, and by race/origin category. See text for details.

Figure 2: Distribution of Desirability for Invitees Relative to Auction Winners

(a) Below-Median Desirability Firms

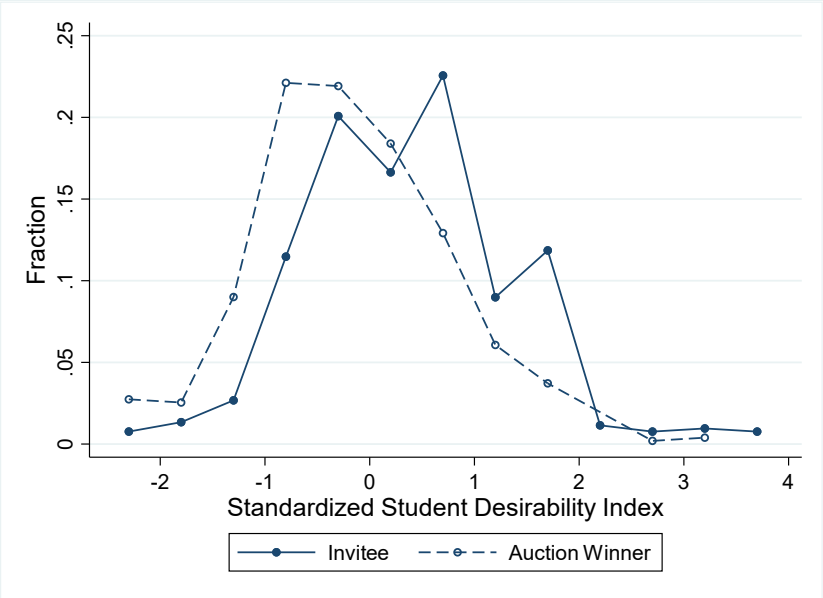

(b) Above-Median Desirability Firms

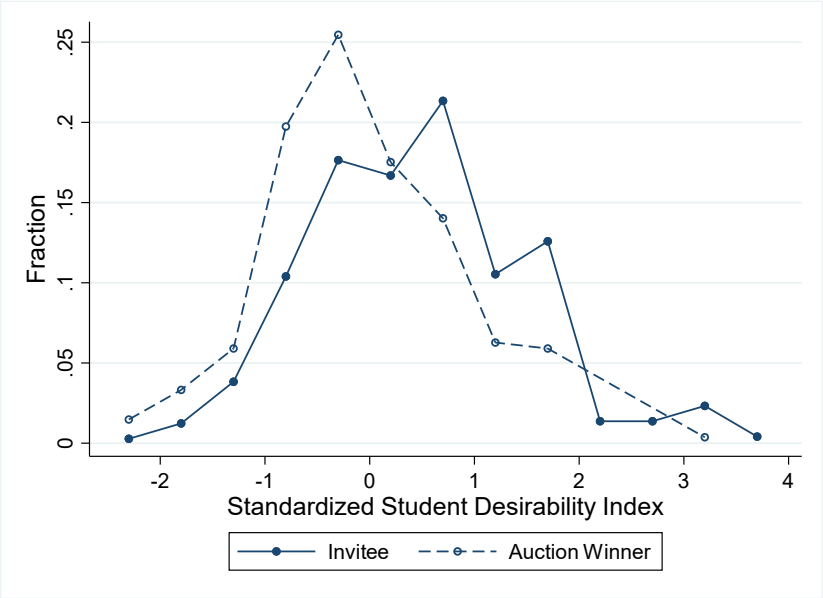

Notes: Plots show histograms of the student desirability index for invited interviewees (invitees) and auction winners. The student desirability index is based on a prediction of total interview invitations, using applicant characteristics. The index is standardized to be mean zero with a standard deviation of one among the sample of all applicants. See text for details.

Figure 3: Congestion and Competition Among Firms

(a) Applications and Interview Invitations per Posting

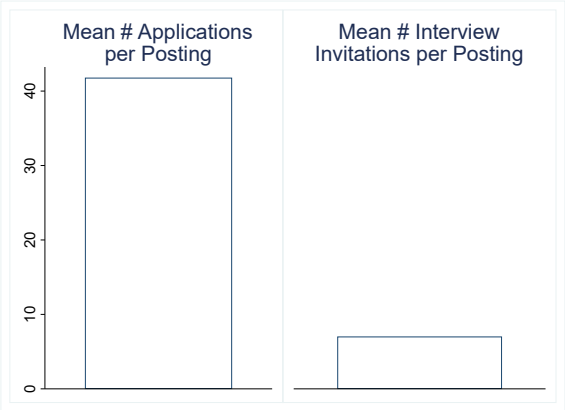

(b) Total Interview Invitations per Applicant and Invited Interviewee

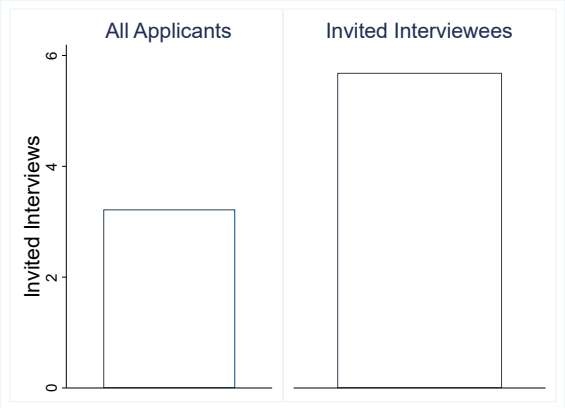

Notes: The left-hand side of Figure (b) shows the mean number of invited interviews per applicant. The right-hand side of Figure (b) shows the mean number of invited interviews, among the sample of job posting-interviewee pairs for which the interviewee was invited.

Figure 4: Likelihood of Hiring Auction Winners Relative to Invitees, by Student and Firm Desirability

(a) Higher-Desirability Candidates: Desirability Index ≥ Mean

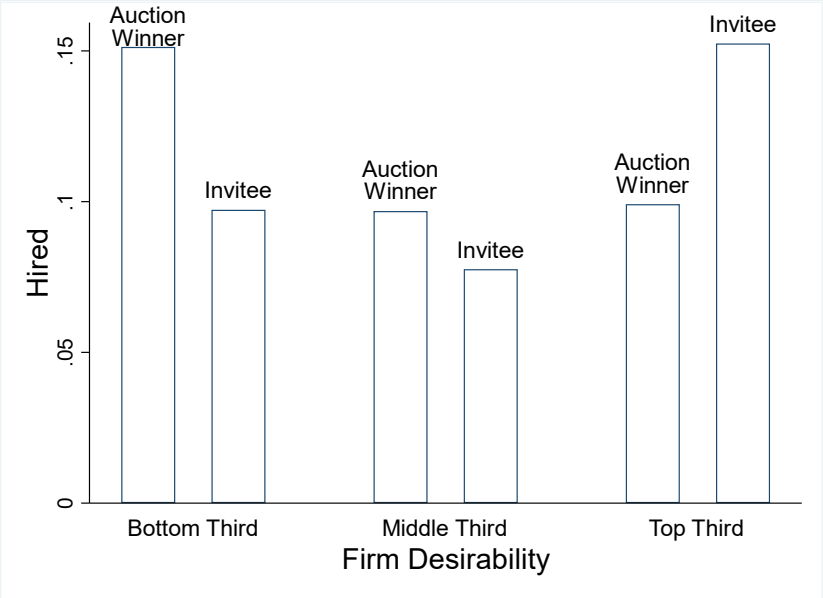

(b) Lower-Desirability Candidates: Desirability Index < Mean

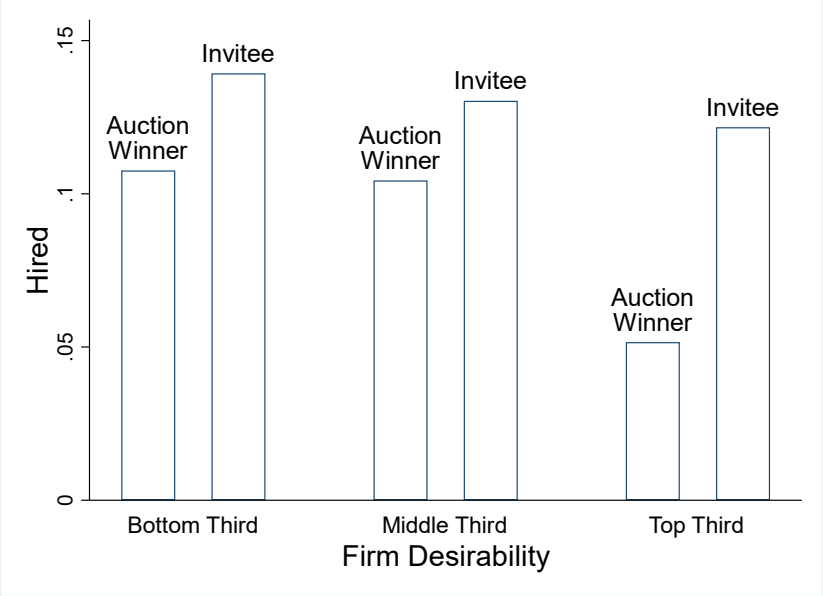

Notes: Bars show the results from a binscatter estimation, in which the observations underlying the bars are at the interviewee, job posting, semester level. The bins are based on the terciles of firm desirability (firm percentile in the salary distribution averaged over all semesters), for the sample of all interviewed applicants. The dependent variable is whether the student is hired for the job posting. The binned estimation additionally controls for semester fixed effects and total students interviewed for the job posting.

Figure 5: Characteristics of Interviewed and Hired Candidates, Invitees vs. Auction Winners

(a) Below-Median Desirability Firms

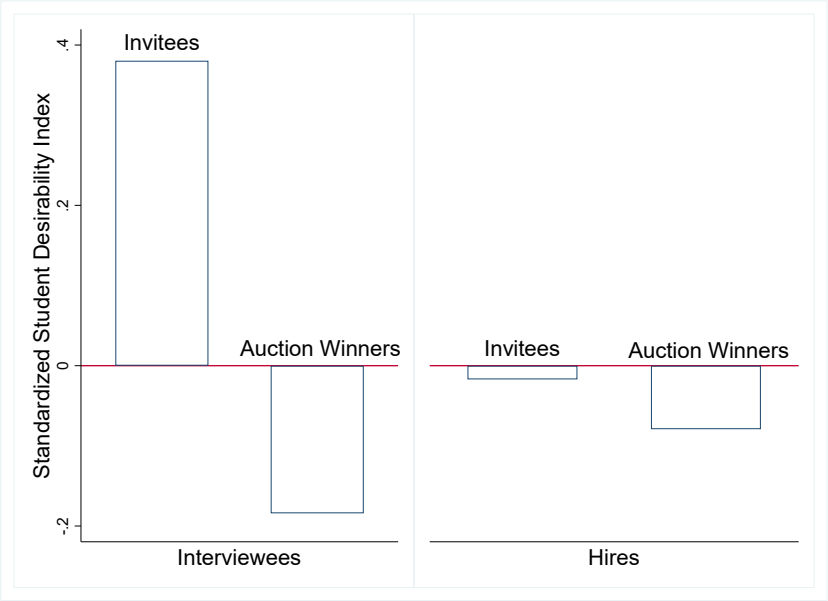

(b) Above-Median Desirability Firms

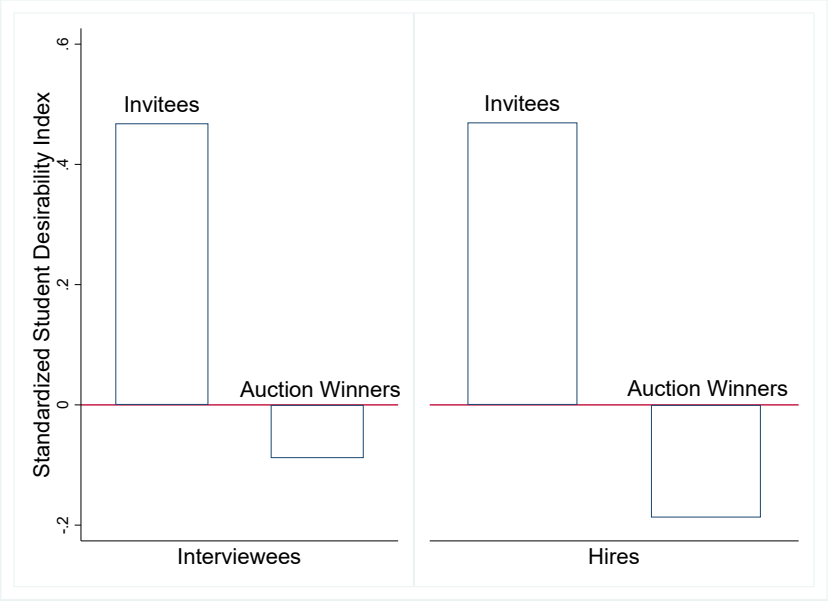

Notes: Plots show the mean student desirability index for invited applicants and auction winners. The left plot shows differences among interviewees. The right plot shows differences among hires. See text for details, including construction of the student desirability index and the firm desirability measure.

Figure 6: Counterfactual Simulations - Auctions Relative to Scenarios without Auctions

Below-Median Desirability Firms

(a) Failure to Hire

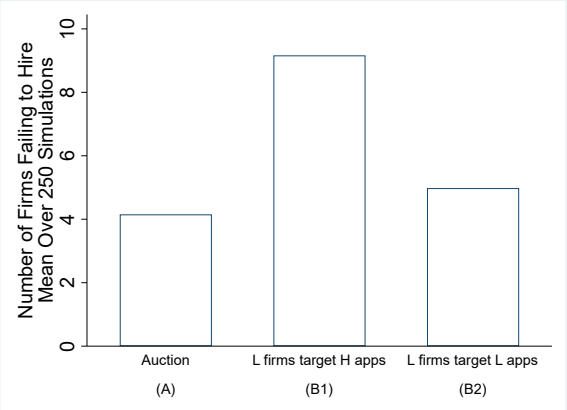

(b) Desirability of Hired Applicants

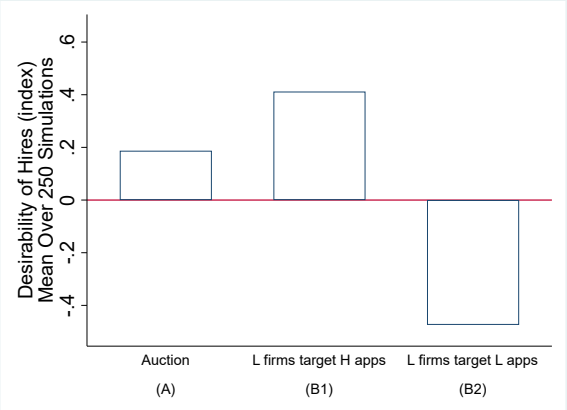

Above-Median Desirability Firms

(c) Failure to Hire

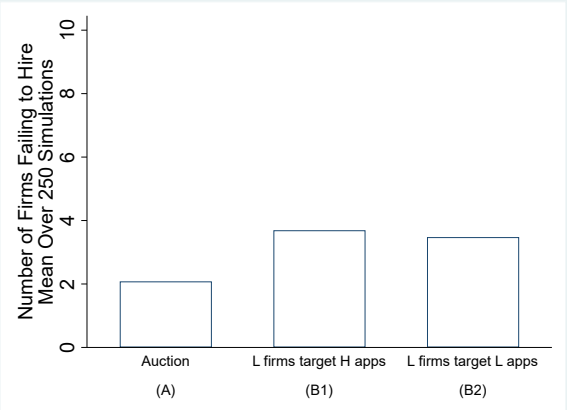

(d) Desirability of Hired Applicants

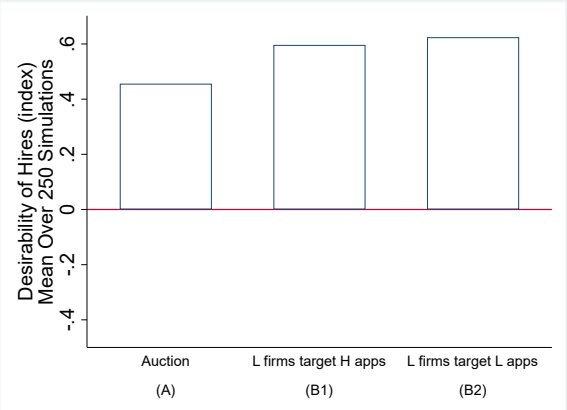

Notes: This table shows results from 250 simulations of scenarios with auctions, and without auctions under two different strategies. The strategy "L firms target H apps" (scenario B1) refers to the strategy in which less desirable firms make their first offers to the highest ranked high desirability applicants on their interview schedule. The strategy "L firms target L apps" (scenario B2) refers to the strategy in which less desirable firms make their first offers to the highest ranked low desirability applicants on their interview schedule. In the auction (scenario A), less desirable firms make their first offers to the highest ranked high desirability auction winners, and then to the highest rank of the high desirability invitees, and then to the highest ranked low desirability applicants (auction or invitee). We show firms failing to hire among firms who made offers in the data (31 low desirability firms and 36 high desirability firms), and quality of hires after two rounds of offers. See text for details.

Table 1: Summary Statistics: Job Postings and Applicants

Total interview schedules with auctions	182
Interview schedules with auctions per semester	20.2
	[9.2]
Interview schedule characteristics, schedules with auctions	
Applicants per schedule	41.8
	[18.0]
Campus interviews per schedule	13.0
	[5.6]
Campus interviews per schedule, invited	7.0
	[3.8]
Campus interviews per schedule, not via invitation	6.0
	[3.0]
Number of hires of campus interviewees per schedule	1.0
	[1.1]
Proportion of schedules with zero hires per semester	0.38
	[.49]
Interviewee characteristics	
Total applications sent	19
	[8.2]
Total interviews	7.4
	[3.7]
Total interviews, invited	4.3
	[3.4]
Total interviews, not via invitation	3
	[2.2]

Notes: Standard deviations in brackets. Statistics are for schedules with auctions with at least one invited and one non-invited interviewee. Interviewee characteristics are presented for the sample of job posting-interviewee pairs.

Table 2: Predicting the Number of Interview Invitations Per Student

	Mean and Standard Deviation		Linear Prediction
	Invitees	Auction Winners	Y = Total Interview Invitations
Female	0.74	0.73	-0.204
	[.44]	[.44]	(0.203)
Black or Hispanic	0.17	0.14	0.0325
	[.38]	[.35]	(0.322)
Asian	0.13	0.23	-0.887***
	[.34]	[.42]	(0.288)
Two or More Years Since Obtaining Bachelor's	0.31	0.27	-0.0300
	[.46]	[.45]	(0.194)
US Citizen/Permanent Resident	0.93	0.86	0.140
	[.26]	[.35]	(0.291)
Pre-Master's Occupation Similar to Degree	0.56	0.45	0.898***
	[.5]	[.5]	(0.189)
Fortune 1000/Forbes 2000 Pre-Master's	0.34	0.22	0.813***
	[.47]	[.41]	(0.216)
Bachelor's Institution: Carnegie Research 1	0.67	0.62	0.584
	[.47]	[.48]	(0.359)
3.4 < Undergraduate GPA ≤ 3.7	0.5	0.5	0.254
	[.5]	[.5]	(0.234)
Undergraduate GPA > 3.7	0.31	0.23	0.853***
	[.46]	[.42]	(0.304)
Undergraduate Latin Honors	0.15	0.09	0.843*
	[.36]	[.29]	(0.450)
Any Awards or Honors	0.61	0.53	0.156
	[.49]	[.5]	(0.188)
Any Leadership in Univ. Activities	0.76	0.69	0.0755
	[.43]	[.46]	(0.216)
Team Sports	0.08	0.06	0.340
	[.28]	[.23]	(0.400)
Any Professional Clubs	0.59	0.49	0.287
	[.49]	[.5]	(0.194)
Any Volunteering	0.52	0.46	0.366**
	[.5]	[.5]	(0.186)
Business/HR/Communications Major	0.45	0.38	0.444**
	[.5]	[.49]	(0.209)
Psychology Major	0.35	0.41	-0.331
	[.48]	[.49]	(0.233)
GMAT/GRE Concorded Score	532.34	532.66	0.001
	[94.81]	[99.17]	(0.001)
Total Applications Per Person	19.18	18.86	0.177***
	[8.24]	[8.2]	(0.013)
Observations	1277	1084	608
R-Squared			0.487

Notes: *** p-value ≤ .01, ** p-value ≤ .05, * p-value ≤ .1. Standard deviations in brackets in Columns 1 and 2. Column 1 contains mean characteristics for those obtaining their interview through the firm's invitation, and column 2 for individuals obtaining their interview through the auction. Standard errors in column 3 are clustered at the student level, and shown in parentheses. The regression in column 3 includes semester fixed effects, and also includes an indicator for other race, highly selective bachelor's institution, any honor societies, and any fraternity/sorority listed on the resume (coefficients not shown). Missing values are replaced with zero, and we include an indicator for whether the value is missing. See paper for details.

Table 3: Likelihood of Hire Conditional on Interview, by Interview Source

Y_{fs} = *Applicant s Hired for Internship f*	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Interview from Auction (Auction)$_{fs}$	-0.004	0.006	-0.013	0.011	-0.011	-0.039	0.051
	(0.012)	(0.021)	(0.015)	(0.022)	(0.016)	(0.038)	(0.034)
Standardized Student Desirability Index$_{s}$	-0.007	-0.028**	0.007				
	(0.008)	(0.013)	(0.011)				
Auction$_{fs}$* Student Desirability$_{s}$	0.009	0.041**	-0.014	0.044**	-0.019	-0.013	-0.017
	(0.011)	(0.019)	(0.015)	(0.020)	(0.015)	(0.042)	(0.039)
Auction$_{fs}$* Student Desirability$_{s}$ * Bid Value$_{fs}$						0.026**	0.004
						(0.012)	(0.011)
Auction$_{fs}$ * Bid Value$_{fs}$						0.012	-0.016
						(0.012)	(0.011)
Observations	2,361	1,034	1,273	1,034	1,273	713	847
R-Squared	0.096	0.118	0.073	0.143	0.088	0.185	0.122
Firm Desirability	All	Low	High	Low	High	Low	High
Student Controls	No	No	No	Yes	Yes	Yes	Yes
Job Posting Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: *** p-value ≤ .01, ** p-value ≤ .05, * p-value ≤ .1. Each observation is an applicant *s* , job posting *f* pair, for a job posting and applicant in semester *t* . We include only pairs for which the applicant *s* has an interview, and is on the firm's interview schedule. The variable *Interview from Auction* denotes applicants who won the interview auction for that interview schedule and the relatively small number who were neither invitees nor auction winners, but presumably obtained the interview through the free-for-all period after the auction. The omitted category is interview by invitation, which includes applicants initially invited for an interview by the firm, and those obtaining their interview after being selected by the firm as an alternate. High desirability (low desirability) firms refers to firms for which their desirability index is above (below) the median for the firms in the regression sample for that semester. The firm's desirability index is their average percentile in the salary distribution across all semesters. The variable Student Desirability is based on a linear prediction of the total number of interview invitations a student receives that semester, and is mean zero with standard deviation one among all applicants. Columns 4 and 5 include the components of the index as linear regressors rather than the index itself. We present bootstrapped standard errors based on 400 replications. See paper for details.

Table 4: Differences Between Invitees and Auction Winners, by Firm Desirability

Panel A: Dependent Variable: Standardized Student Desirability Index

	Interviewees	Hires
Auction Winner	-0.584***	-0.119
	(0.076)	(0.181)
Auction*High Q Firm	-0.009	-0.500*
	(0.081)	(0.267)
High Q Firm	0.090*	0.447**
	(0.053)	(0.203)
Observations	2,307	180
R-Squared	0.173	0.144

Panel B: Components of Index as Dependent Variable

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. Variable:	Fortune 1000/Forbes 2000 Experience		Bachelor's GPA in Top Quartile		Asian		US Citizen	
	Interviewees	Hires	Interviewees	Hires	Interviewees	Hires	Interviewees	Hires
Auction	-0.130***	-0.050	-0.078**	0.035	0.118***	0.167*	-0.310***	-0.432***
	(0.035)	(0.101)	(0.037)	(0.104)	(0.031)	(0.095)	(0.052)	(0.157)
Auction*High Q Firm	-0.001	-0.145	-0.024	-0.145	-0.044	-0.079		
	(0.039)	(0.133)	(0.036)	(0.149)	(0.034)	(0.120)		
High Q Firm	-0.023	-0.066	0.044*	0.098	-0.046**	-0.070	0.124**	0.190
	(0.025)	(0.100)	(0.026)	(0.102)	(0.019)	(0.075)	(0.048)	(0.206)
Observations	2,307	180	1,885	146	2,222	168	438	40
R-Squared	0.064	0.084	0.049	0.028	0.059	0.110	0.295	0.414

Notes: Student Desirability Index is the predicted number of interview invitations for a student, based on a linear regression of total interview invitations on student characteristics, total applications, and semester fixed effects. The prediction is based only on the student characteristics. The variable is standardized to be mean zero and with standard deviation of one for the sample of all applicants. High Q Firm refers to firms for which the average salary percentile across all semesters is above the median for the firms in the regression sample for that semester. Regressions also include semester fixed effects. Standard errors are clustered at the student level. In Panel B the number of observations varies for the different dependent variables in columns 1-6 due to missing values. In columns 7-8, when the dependent variable is an indicator for U.S. citizen, the number of observations falls as these columns are limited to job postings that allow applications from non-U.S. citizens. See text for details.

Preference Signaling and Worker-Firm Matching: Evidence from Interview Auctions Online Appendix

Ron A. Laschever and Russell Weinstein*

July 29, 2021

1 Additional Details on Framework and Model Equilibrium

As introduced in Section 2, we consider a simple two-firm model, with Firm A and Firm B. There are a total of K applicants, of which λK applicants are of type H (high desirability) and $(1-\lambda)K$ are of type L (low desirability). To focus on uncertainty over applicant acceptance likelihood, we treat wages as exogenous to the model. We assume that firms post wages, and thus pay the same wage to any hire, regardless of whether they hire an H or an L candidate; and that Firm A posts a high wage and Firm B posts a low wage. Firms observe the wages posted by their competitor. Each firm has one vacancy, and can make only one offer.

We assume firms have uncorrelated, idiosyncratic preferences over applicants of a given type. For example, for Firm A and H-type applicant i, $\pi_A(H_i) = v_A(H) + \varepsilon_{A,i} - w_A$, where $\varepsilon > 0$. Firm A pays wage w_A and receives two components from hiring an applicant that contribute to its profit - a common component $v_A(H)$ which it receives from hiring any H-type applicant, and a firm-employee match-specific component $\varepsilon_{A,i}$

*Laschever: Compass Lexecon. E-mail: rlaschever@compasslexecon.com; Weinstein: School of Labor and Employment Relations and Department of Economics, University of Illinois at Urbana-Champaign. E-mail: weinst@illinois.edu. The opinions expressed in this article are the authors' own and do not reflect the views of Compass Lexecon.

from hiring applicant i. Firms fully observe the common component, as well as their match-specific component for each applicant, but do not observe the match-specific components of their competitor. Furthermore, we assume that for H-type applicants the match-specific component ε is relatively small and that it is bounded above by $\frac{v_j(H)-w_j}{\lambda K-1}$ for Firm $j = A, B$.

We assume that conditional on hiring, the profit from hiring any H-type applicant is larger than the profit from hiring the firm's most-preferred L-type applicant. Specifically, we assume $v_j(H) > v_j(L) + \bar{\varepsilon}_{jL} > \pi_j(0)$ for $j = A, B$, where $\bar{\varepsilon}_{jL}$ denotes the match-specific component from hiring the firm's most-preferred L applicant. Productivity of an L applicant is denoted $v_j(L)$, and profits when the vacancy is unfilled are denoted $\pi_j(0)$. This condition implies that $\bar{\varepsilon}_{jL}$ is bounded above by $v_j(H) - v_j(L)$. Finally, we make an assumption on the upper-bound of the ratio of profits from hiring an H-type relative to an L-type: $\frac{v_B(H)+\bar{\varepsilon}_{BH}-w_B}{v_B(L)-w_B} < \frac{\lambda K}{\lambda K-1}$. We also show our results hold when this assumption is relaxed.

Without loss of generality, we denote the utility an applicant derives from a match as unidimensional and dependent on wage, and therefore all applicants prefer Firm A as its wage is greater than Firm B's wage: $u(w_A) > u(w_B) > 0$, where we normalize the utility of being unmatched to zero.

We consider the following game, consistent with our setting's timeline. Applicants submit applications to Firm A and Firm B. Firms review the applications and simultaneously choose one applicant to invite for an interview. Students then decide whether to participate in the auction for an interview with a firm, if they were not invited for an interview. The winner of the auction gets an interview at the firm, and ties in the auction are broken at random. Firms then interview their two applicants: one that they chose and the other who was the auction winner. After the interviews, firms make their offer simultaneously with their competitor. Finally, students decide which offers to accept.

We solve for an equilibrium using backward induction, starting with the last step, in which students accept offers from their preferred firm if they have multiple offers. Thus, the student getting an offer from Firm A will accept that offer, and the student getting an offer from Firm B will accept only if they did not get an offer from Firm A. Given that any applicant will accept Firm A's offer, it is clear, based on our assumptions above, that Firm A inviting an L-type applicant is a dominated strategy. As a result, in the analysis below we focus our attention on solving for an equilibrium

given that Firm A invites an H-type applicant, and we assume Firm B knows Firm A will invite an H-type applicant (but not which specific H-type applicant). We also assume that the invited interviewee at Firm A does not participate in the auction for Firm B since she knows she will get an offer from Firm A, and this is also known by Firm B. As a result, Firm B knows that any applicant participating in its auction has not received an offer from Firm A.

Knowing the acceptance strategies, and specifically that all applicants will accept Firm A's offer with certainty, if Firm A invites an H-type applicant it will always make an offer to that applicant regardless of the winner of the auction. If Firm B makes an offer to an H-type invitee, there is some risk this is the same person that was invited by Firm A, in which case the offer made by Firm B would be rejected. Thus, if Firm B has an H-type invitee and an H-type auction winner, Firm B will make an offer to the auction winner, since the applicant will accept with certainty and we assume Firm B faces relatively little match-specific heterogeneity among the H-type applicants. Specifically, Firm B will make an offer to the auction winner if the following condition holds:

$$\left(1 - \frac{1}{\lambda K}\right) * (v_B(H) + \bar{\varepsilon}_{BH} - w_B) < v_B(H) - w_B + \varepsilon_{BH,winner}$$

where $\bar{\varepsilon}_{BH}$ is the firm-employee match-specific component associated with Firm B's most preferred H-type applicant, λ is the fraction of applicants that are type H, and $(1 - \frac{1}{\lambda K})$ is the probability that the H-type invitee was not also the H-type invitee at Firm A. The expression on the left hand side of the inequality is the expected profit from making an offer to the firm's most-preferred H-type applicant. The right hand side of the inequality is the profit from making an offer to an H-type auction winner, with match-specific component $\varepsilon_{BH,winner}$. This condition will be satisfied if it is satisfied for the case where the auction winner's ε is equal to zero (the extreme lower bound), resulting in the following condition:

$$\left(1 - \frac{1}{\lambda K}\right) * (v_B(H) + \bar{\varepsilon}_{BH} - w_B) < v_B(H) - w_B$$

The inequality is satisfied given our aforementioned assumption that $\bar{\varepsilon}_{BH}$ is bounded above by $\frac{v_B(H) - w_B}{\lambda K - 1}$.

If Firm B has invited an H-type applicant and the auction winner is type L, Firm

B will make an offer to the auction winner if the profit from hiring the L-type auction winner is greater than the expected profit from hiring the most preferred H type:

$$(1 - \frac{1}{\lambda K}) * (v_B(H) + \bar{\varepsilon}_{BH} - w_B) < v_B(L) - w_B + \varepsilon_{BL,winner} \tag{1}$$

where $\varepsilon_{BL,winner}$ is the match-specific component associated with the L-type auction winner. If (1) holds for an applicant whose match-specific component is at the extreme lower bound of zero, then Firm B will always prefer the L-type auction winner:

$$(1 - \frac{1}{\lambda K}) * (v_B(H) + \bar{\varepsilon}_{BH} - w_B) < v_B(L) - w_B \tag{2}$$

This is satisfied given our aforementioned assumption that $\frac{v_B(H)+\bar{\varepsilon}_{BH}-w_B}{v_B(L)-w_B} < \frac{\lambda K}{\lambda K-1}$.

If Firm B has invited an L-type applicant and the auction winner is type L, Firm B will make an offer to the invited candidate since they each will accept with certainty, and if Firm B invites an L-type applicant it will choose to invite its most preferred applicant.

Next, knowing the offer strategies, applicants make decisions about participating in the auction. If Firm A invites an H-type applicant, given the offer strategy, it will always make an offer to that applicant, and applicants have no incentive to participate in the auction.

If Firm B invited an H-type applicant, given the offer strategy, Firm B will make an offer to whomever wins the auction. As a result, all L-types participate in the auction at Firm B. If Firm A invited an H-type applicant, all H-type applicants will also participate in the auction at Firm B, as argued above. If Firm B invited an L-type applicant, given the offer strategy, Firm B will make an offer to the L-type invitee rather than the L-type auction winner, and so L-types will not participate in the auction. All H-type applicants will participate in Firm B's auction.

Knowing the auction participation strategies, firms decide whether to invite an H- or L-type applicant. As we have discussed, Firm A will invite an H-type applicant, since its preferred applicant will accept with certainty. If Firm B invites an H-type applicant, this will result in hiring whomever wins the auction, either an H- or L-type applicant. If Firm B invites an L-type applicant, L-type applicants do not participate in the auction, and so the firm is guaranteed to have an H-type auction winner. Thus, inviting an L-type applicant will result in hiring an H-type auction winner. This is

preferred to inviting an H-type applicant, and having a non-zero probability of hiring an L-type auction winner (e.g. if an L-type rather than an H-type wins the auction), given our assumption above that $\varepsilon_{jL} < v_B(H) - v_B(L)$.

We note that our assumption on the ratio of profits from hiring an H-type relative to an L-type (equation (2)) is not the sharpest (largest-possible) bound in the sense that it can be further increased while resulting in a similar outcome. Consider the case in which there were some values of $\varepsilon_{BL,winner}$ such that if those applicants were to win the auction, Firm B would prefer the risk of making an offer to an H-type invitee (i.e. equation (2) does not hold). If Firm B invited an H-type applicant, all L-type applicants would still participate in the auction. Since applicants do not know their own match-specific components, they do not know if the firm would prefer making them an offer to the risk of making an offer to the H-type invitee. Thus, if Firm B invites an H-type applicant, it is not guaranteed to have an H-type auction winner, whereas inviting an L-type applicant guarantees an H-type auction winner. At the extreme, if we had alternatively assumed there is no $\varepsilon_{BL,winner}$ such that equation (1) will hold, then if Firm B invites an H-type, L-type applicants have no incentive to participate in the auction. In this case, Firm B will be indifferent between inviting an H- or L-type applicant, because either will guarantee an H-type auction winner. These alternative assumptions still yield the result that Firm B will hire an H-type auction winner.

We have shown that in equilibrium, Firm A will never hire an auction winner, Firm B will always hire an H-type auction winner and prefers this to hiring an H-type invited candidate. Firm B will never hire an L-type candidate. The signal is useful to the less preferred firm when trying to hire a desirable candidate, and less useful to the preferred firm, and less useful when coming from a less desirable candidate. This is consistent with the predictions outlined in Section 2.

Appendix Figure A1: Relationship between Offers Rejected and Total Offers

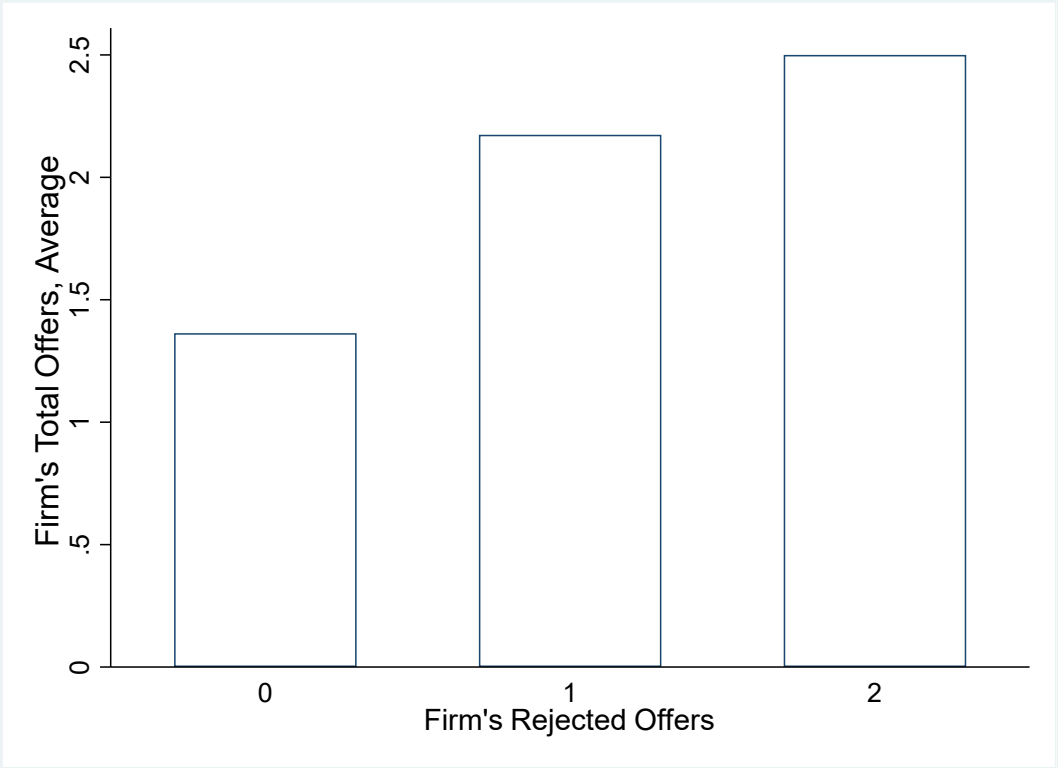

Notes: This plot shows the average number of total offers made by firms, for each value of rejected offers on the x-axis. Among firms that made offers, nearly 87% had two or fewer offers rejected. We do not show results for greater numbers of offers rejected, as the sample size per value of rejected offers is very small for these values. See paper for details.

Appendix Figure A2: Unfilled Vacancies, Offers to Auction Winners, and Competition from Other Firms

(a) Likelihood that Interviewee is Hired, by Firm Desirability

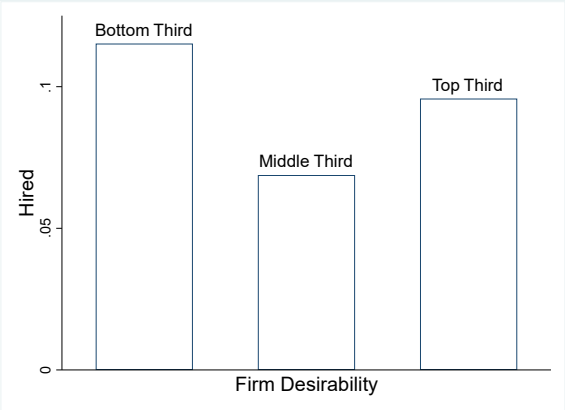

(b) Likelihood that Interviewee is Hired, by Firm Desirability and Interview Source

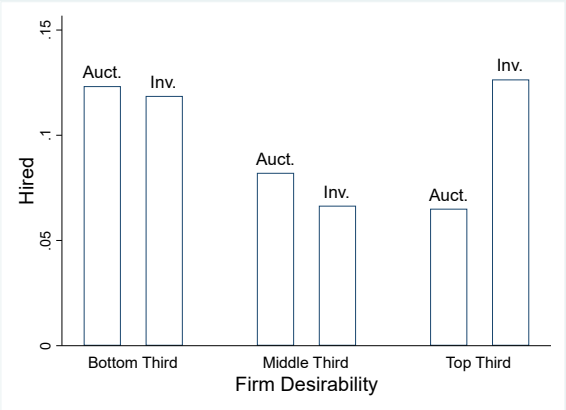

(c) Likelihood that Interviewee Gets an Offer, by Firm Desirability and Interview Source

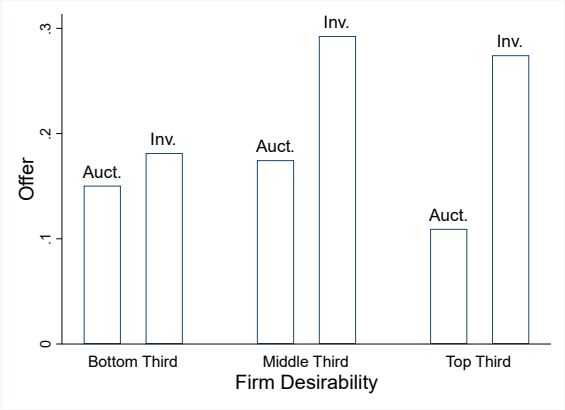

(d) Desirability Index of Invitees Receiving Offers, by Firm Desirability

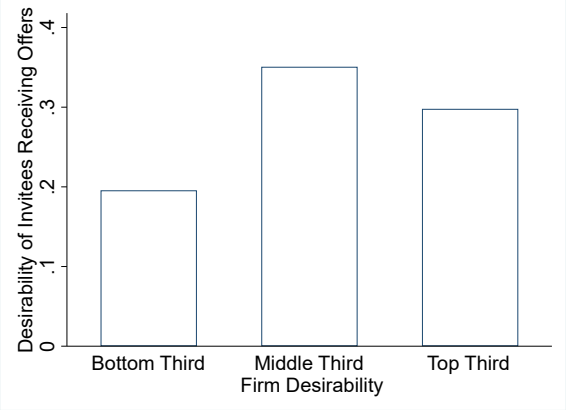

Notes: Each plot shows results of binscatter estimation, with observations binned by firm desirability (average percentile in the salary distribution, across all semesters). Observations are at the applicant *s*, job posting *f* level, where the applicant and job posting are in semester *t*. The binned estimation in (a) through (c) additionally controls for semester fixed effects and total students interviewed for the job posting. For (c) and (d), the sample is restricted to recruiting during the 2010-2011 and 2011-2012 academic years, as these students received surveys asking about all offers. We further restrict to respondents. For (d) the sample is limited to invitees receiving an offer from job posting *f* that semester. See text for details.

Appendix Figure A3: Interview Invitations from High Desirability Firms, Among High Desirability Invitees at Lower Desirability (Below-Median) Firms

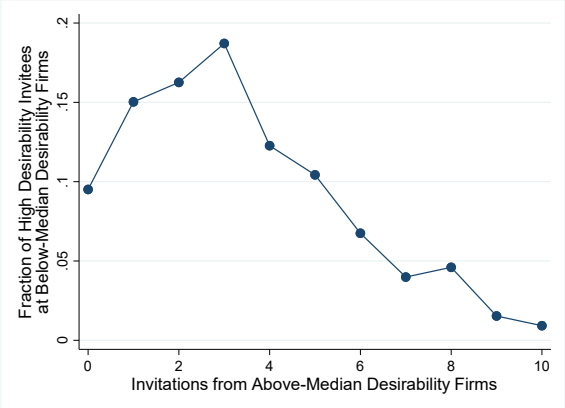

Notes: This figure shows a histogram of the number of invited interviews from above-median desirability firms, among high desirability invitees at below-median desirability firms. High desirability invitees are defined as applicants with desirability index of at least zero (i.e. above average desirability).

Appendix Table A1: Likelihood of Hire and Offer Conditional on Interview, by Interview Source

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Hired		Offer		Offer	
Interview from Auction (Auction)$_{fs}$	0.015	0.013	-0.007	-0.077**	-0.002	-0.079**
	(0.028)	(0.023)	(0.037)	(0.030)	(0.040)	(0.032)
Standardized Student Desirability Index$_{s}$	-0.021	0.001	0.002	0.053**		
	(0.017)	(0.018)	(0.026)	(0.024)		
Auction$_{fs}$ * Student Desirability$_{s}$	0.060**	0.005	0.071*	-0.049	0.066	-0.077**
	(0.027)	(0.027)	(0.041)	(0.032)	(0.047)	(0.034)
Observations	459	551	459	551	459	551
R-Squared	0.127	0.083	0.177	0.148	0.216	0.217
Firm Desirability	Low	High	Low	High	Low	High
Student Controls	No	No	No	No	Yes	Yes
Job Posting Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Note: *** p-value ≤ .01, ** p-value ≤ .05, * p-value ≤ .1. Each observation is an applicant *s,* job posting *f* pair, for applicant *s* applying in semester *t* , and job posting *f* posted in semester *t* . Regressions are limited to years in which students received a survey asking for information about all offers (2010-2011 and 2011-2012 academic years), and to students who responded to the survey. Columns 5 and 6 include the components of the student desirability index as linear regressors rather than the index itself. We present bootstrapped standard errors based on 400 replications. See paper for details.

Appendix Table A2: Differences in Proportion Asian Between Invitees and Auction Winners, Among Firms that only Allow Applications from US Citizens

	Asian	
	Interviewees	Hires
Auction	0.0168	0.00581
	(0.0288)	(0.104)
Auction*High Desirability Firm	0.0274	0.0975
	(0.0313)	(0.132)
High Desirability Firm	-0.0209	-0.0107
	(0.0179)	(0.0921)
Observations	1,679	119
R-Squared	0.044	0.090
Firms	Only allow US Citizens	

Notes: Regressions also include semester fixed effects. Standard errors are clustered at the student level. Firms with citizenship restrictions are those that restrict applications to individuals who are US Citizens or permanent residents. High Desirability Firm refers to firms for which the average salary percentile across all semesters is above the median for the firms in the regression sample for that semester. See text for details.

Appendix Table A3: Likelihood of Hire Conditional on Interview, by Interview Source

Y_{fs} = Applicant s Hired for Internship f	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Interview from Auction (Auction)$_{fs}$	0.030	-0.027	0.019	-0.027	0.044	0.076*	0.049
	(0.018)	(0.017)	(0.018)	(0.017)	(0.042)	(0.044)	(0.043)
Standardized Student Desirability Index$_s$	-0.009	-0.003	-0.012	-0.003	-0.040	-0.036	-0.042
	(0.011)	(0.012)	(0.011)	(0.013)	(0.027)	(0.028)	(0.026)
Auction$_{fs}$*Student Desirability$_s$	0.013	0.001	0.017	0.001	0.054	0.056	0.060*
	(0.017)	(0.018)	(0.017)	(0.018)	(0.037)	(0.038)	(0.035)
Auction$_{fs}$*Student Desirability$_s$*Firm Desirability$_f$					-0.073	-0.078	-0.082
					(0.060)	(0.059)	(0.058)
Auction$_{fs}$*Firm Desirability$_f$					-0.083	-0.124*	-0.088
					(0.064)	(0.067)	(0.067)
Student Desirability$_s$*Firm Desirability$_f$					0.053	0.048	0.056
					(0.043)	(0.047)	(0.043)
Differential Likelihood of Hiring an Auction Winner Relative to Invitee with Student Desirability 1 SD Above Mean, for							
Firms with Desirability 1 SD Below Mean					0.040	.046**	
					[.025]	[.022]	
Firms with Desirability 1 SD Above Mean					-0.027	-0.033	
					[.021]	[.022]	
Observations	972	1,213	1,103	1,226	2,307	2,185	2,329
R-Squared	0.093	0.103	0.089	0.104	0.097	0.100	0.098
Firm Desirability	Low	High	Low	High	All	All	
Firm Desirability based on Fulltime or Intern Salary	Fulltime (FT)		FT, Imputing		Intern	FT	FT, Imputing

Notes: *** p-value ≤ .01, ** p-value ≤ .05, * p-value ≤ .1. See notes to Table 3. For specifications where firm desirability is based on fulltime salary and we impute, we impute missing firm desirability based on the firm's desirability measure based on the intern salary. We present bootstrapped standard errors based on 400 replications. See paper for details.

Appendix Table A4: Likelihood of Hire Conditional on Interview, by Interview Source
Firm Desirability Based on Number of Applications Received

Y_{fs} = *Applicant s Hired for Internship f*	(1)	(2)	(3)
Interview from Auction (Auction)$_{fs}$	-0.009	-0.001	-0.105*
	(0.019)	(0.018)	(0.054)
Standardized Student Desirability Index$_{s}$	-0.020*	0.002	-0.085**
	(0.011)	(0.011)	(0.040)
Auction$_{fs}$*Student Desirability$_{s}$	0.036**	-0.014	0.121**
	(0.016)	(0.017)	(0.052)
Auction$_{fs}$*Student Desirability$_{s}$*Firm Desirability$_{f}$			-0.187**
			(0.087)
Auction$_{fs}$*Firm Desirability$_{f}$			0.167*
			(0.089)
Student Desirability$_{s}$*Firm Desirability$_{f}$			0.130*
			(0.067)
Observations	1,049	1,312	2,361
R-Squared	0.081	0.107	0.100
Firm Desirability based on Applications	Low	High	All

Notes: *** p-value ≤ .01, ** p-value ≤ .05, * p-value ≤ .1. See notes to Table 3. For each firm and semester, we calculate the number of applications the firm received. We then calculate the firm's percentile in the distribution of applications received for that semester, and then average across all semesters. High desirability (low desirability) firms refers to firms for which this average percentile is above (below) the median for the firms in the regression sample for that semester. We present bootstrapped standard errors based on 400 replications. See paper for details.

Appendix Table A5: Offer Acceptances and Other Offers Received, by Interview Source

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	*Offer Accepted* $_{fs,}$ Conditional on Offer				*Stay for FT Job* $_{fs,}$ Conditional on Internship			
Interview from Auction (Auction)$_{fs}$	0.132	0.293**	0.055	0.403*	0.065	-0.087	-0.117	0.354
	(0.146)	(0.131)	(0.182)	(0.223)	(0.128)	(0.190)	(0.271)	(0.443)
Standardized Student Desirability Index$_{s}$	-0.202**	-0.106	-0.146	-0.099	-0.142	-0.032	-0.228	-0.045
	(0.098)	(0.071)	(0.128)	(0.111)	(0.111)	(0.077)	(0.178)	(0.256)
Auction$_{fs}$*Student Desirability$_{s}$	0.158	0.125	0.047	0.094	0.074	0.114	0.119	0.258
	(0.161)	(0.137)	(0.205)	(0.255)	(0.146)	(0.168)	(0.246)	(0.617)
Firm Desirability$_{f}$	0.609	0.912**			0.591	-0.409		
	(0.477)	(0.437)			(0.478)	(0.720)		
Observations	69	83	69	83	70	66	70	66
R-Squared	0.149	0.230	0.609	0.518	0.175	0.138	0.712	0.628
Firm Desirability	Low	High	Low	High	Low	High	Low	High
Semester FE	Yes	Yes	No	No	Yes	Yes	No	No
Job Posting FE	No	No	Yes	Yes	No	No	Yes	Yes

Notes: *** p-value ≤ .01, ** p-value ≤ .05, * p-value ≤ .1. We present bootstrapped standard errors based on 400 replications. Observations are at the applicant *s* , job posting *f* level, where the applicant and job posting are in semester *t* . Regressions in columns 1-4 are limited to years in which students received a survey asking for information about all offers (2010-2011 and 2011-2012 academic years), and to students who responded to the survey, and received an offer from job posting *f* in that semester. Regressions in columns 5-8 exclude students who started the program in Fall 2011 or Spring 2012, since we do not observe fulltime jobs for those students. See text for details.

Appendix Table A6: Counterfactual Exercises: Auctions Relative to Scenarios without Auctions

	Low Desirability Firms			High Desirability Firms		
	Auction	No Auction		Auction	No Auction	
Number of Firms Failing to Hire	4.2	9.2	5.0	2.1	3.7	3.5
	[1.7]	[1.9]	[1.5]	[1.3]	[1.5]	[1.4]
Mean Quality of Hire	0.2	0.4	-0.5	0.5	0.6	0.6
	[.1]	[.1]	[.1]	[.1]	[.1]	[.1]
Low Desirability Firms Make Offers to Low or High Desirability Applicants?	High	High	Low	High	High	Low
Scenario	A	B1	B2	A	B1	B2

Notes: This table shows results from 250 simulations of scenarios with auctions, and without auctions under two different strategies. In the auction, less desirable firms make their first offers to their highest ranked high desirability auction winners, then to their highest ranked high desirability invitee, and then to the highest ranked low desirability applicants (auction or invitee). We show firms failing to hire among firms who made offers in the data (31 low desirability firms and 36 high desirability firms), and quality of hires, after two rounds of offers. Standard deviations in brackets. See text for details.