ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Barach, Moshe A.; Horton, John

Working Paper

How Do Employers Use Compensation History?: Evidence from a Field Experiment

CESifo Working Paper, No. 6559

Provided in Cooperation with:
Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Barach, Moshe A.; Horton, John (2017) : How Do Employers Use Compensation History?: Evidence from a Field Experiment, CESifo Working Paper, No. 6559, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:
https://hdl.handle.net/10419/167545

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

CESifo WORKING PAPERS

6559
2017

July 2017

How Do Employers Use Compensation History?: Evidence from a Field Experiment

Moshe A. Barach, John J. Horton

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 6559*
Category 4: Labour Markets

How Do Employers Use Compensation History?: Evidence from a Field Experiment

Abstract

We report the results of a field experiment in which treated employers could not observe the compensation history of their job applicants. Treated employers responded by evaluating more applicants, and evaluating those applicants more intensively. They also responded by changing what kind of workers they evaluated: treated employers evaluated workers with 7% lower past average wages and hired workers with 16% lower past average wages. Conditional upon bargaining, workers hired by treated employers struck better wage bargains for themselves. Using a structural model of bidding and hiring, we find that the selection effects we observe would also occur in equilibrium.

JEL-Codes: J010, J300, M500, M510.

Keywords: field experiments, compensation, search and screening.

*Moshe A. Barach**
Georgetown University
McDonough School of Business
USA - Washington D.C. 20057
moshe.barach@georgetown.edu

John Horton
New York University
Stern School of Business
USA - 10012 New York NY
john.horton@stern.nyu.edu

*corresponding author

July 5, 2017
For helpful comments and advice, thanks to Ned Augenblick, Amanda Agan, Amanda Pallais, Joe Golden, John Morgan, Liz Lyons, Noam Yuchtman, Paul Oyer, Stan Veuger, and Steve Tadelis, as well as seminar participants at Berkeley/Haas, and Georgetown MSB.

1 Introduction

There are good reasons for an employer to be interested in a job applicant's past wages. In a competitive labor market, a very recent wage in a similar job is approximately the worker's marginal productivity—precisely what a would-be employer is interested in learning (Kotlikoff and Gokhale, 1992; Oyer et al., 2011; Altonji and Pierret, 2001; Lange, 2007; Kahn and Lange, 2014).[1] Knowing an applicant's productivity can also help the firm avoid the wasteful screening of "over-qualified" applicants, i.e., those with market productivity (as evinced by past wages) higher than the ceiling productivity of the job in question.

Perhaps as a consequence of the screening value of past wages to would-be employers, half of workers in the US report that their current employer learned their wage from their previous job (Hall and Krueger, 2012). Consistent with wage history being used for evaluation, over 80% of workers in the US report that if their employer learned their past wage, they learned it before extending a job offer.[2] Despite the apparent usefulness of wage history information to employers, there is a public concern about its use, in that it can potentially create a path dependence in wages, perhaps impeding wage growth, particularly for women and disadvantaged minorities.

In this paper, we report the results of a field experiment in which treated employers in an online labor market could not observe the wage histories of their applicants, whereas control employers could. The "wage history" for an applicant is the collection of hourly wages for all on-platform contracts started or completed by that applicant, at the time of application.[3]

Our empirical focus is on how the absence of wage history information

[1] Knowing the applicants' productivity would also inform the firm's bargaining strategy, such as what initial offer to make, or how to respond to the worker's initial wage proposal.

[2] We build on the work of Hall and Krueger (2012) and conducted a nationally representative survey, asking whether employers asked about past wage history, and when this occurred. The survey in question was conducted using Google Surveys. It is described in detail in Appendix A.1.

[3] We use the terms "employer" and "employee" or "worker" for consistency with the economics literature, not as a commentary on the legal nature of the relationships created on the platform.

changed the hiring process. We are particularly interested in whether the treatment affected: (1) the extent and intensity of information acquisition by employers, (2) the attributes of the workers evaluated and ultimately hired, (3) wage bargaining, and (4) whether employers still made a hire.

We find that without access to applicant wage histories, employers responded by enlarging the pool of applicants they considered, with treated employers evaluating about 7% more applicants. They also evaluated those applicants more intensively, asking more—and more substantive—questions. Surprisingly, we find no evidence that treated employers put more weight on other individual worker productivity signals available to them, such as past feedback scores, past market experience, and so on. In short, treated employers responded to their information deficit primarily by acquiring more of their own information.

In addition to causing more extensive and intensive evaluation, the treatment also changed what *kind* of workers were evaluated. Treated employers were far more likely to evaluate workers with relatively low past wages. Applicants who were "called-back" by treated employers had about 7% lower average past wages. Employers apparently did not find these lower wage workers wanting after further evaluation, as workers hired by treated employers had about 16% lower past average wages. The observed preference for lower wage workers—which we refer to as "bargain hunting"—is consistent with the framing of the hiring problem found in the personnel and labor economics literatures (Oyer et al., 2011). In short, the treatment made low experience/low wage bid workers seemed like better "deals" relative to high experience/high wage bid workers, as firms with less information infer the workers have closer to the average level of productivity.

For job openings in which a hire was made, we measure the extent of bargaining by comparing the initial proposed wage bid of the worker to the wage that was ultimately agreed upon. We find no evidence that the treatment affected the probability that bargaining occurred (as measured by any difference between the initial bid and the ultimate agreed-upon wage), but we find evidence that when bargaining occurred, workers hired by treated employers

struck a more favorable wage bargain—they were offered and accepted wages which were 9% more of their initial bid compared to those workers bargaining with control employers.

Given the less favorable wage bargain struck by treated employers who still made a hire, a natural concern is that a greater fraction of treated employers might decide to forgo hiring altogether rather than pay higher wages (or, perhaps drop out earlier by deciding not evaluate anyone given their lack of wage history information). This concern was not borne out in the experiment, as treated employers were *more* likely to make a hire. Despite hiring less experienced workers, there is no evidence that treated employers had worse contractual outcomes, as measured by the feedback they gave to hired workers and their reporting to the platform of whether or not the project was a success.

One consistent finding in our results is that effects are stronger for employers that stated, ex ante, they were willing to consider less experienced, less proven workers. These employers are willing to screen more applicants and show a greater willingness to hire less experienced workers when they lack wage history information. This heterogeneity is important, as it suggests different kinds of employers might be differently impacted by policies restricting the use of wage history information.

It is important to note that job applicants in the experiment did not know that their past wage histories might be hidden from certain employers. If workers knew that employers lacked access to wage history, they might change their wage bidding, perhaps turning some of our bargain hunting selection effects into price effects. To address this equilibrium generalizability limitation of the experiment, we estimate a structural model of employer hiring. With our fitted model, we can compute each worker's optimal wage bidding adjustment when facing an employer who lacks wage history information, as well as competition from other bidding workers who are also adjusting. We find that for all workers, regardless of their expected productivity, almost no wage adjustment is optimal, suggesting that our bargain hunting experimental results would hold in equilibrium.

Our paper is a contribution to a larger literature on the role of information

in the hiring process. It is the first paper that we are aware of that explores the role of compensation history in hiring—a particularly potent source of information. It also analyzes a true experiment, which is rare in a literature that has mostly been observational, at least with respect to removing whole "classes" of information rather than just whether an applicant has some attribute or not. The experiment is also timely, in the sense that it approximates policies that are being implemented—or under active consideration—in several conventional labor markets. For example, both NYC and Philadelphia recently passed laws that prevent employers from asking candidates about past compensation.[4] To the extent our results generalize to these settings, these policy proposals would have the intended effect; they would help relatively less experienced workers get their foot in the door (our bargain hunting results), and perhaps help those workers obtain a better wage bargain (our bargaining results), without reducing hiring (our finding of more hiring in the treatment group).

A feature of this paper that differentiates it from the larger literature on information in hiring is that our rich empirical setting allows us to capture margins of adjustment by employers that would be difficult or impossible to detect in other settings. As a case in point, we document the importance of endogenous information acquisition as an employer response to an information deficit.[5] This margin of adjustment has been almost entirely overlooked in extant empirical work. The margin is important because it changes the conception of the hiring problem from a purely statistical selection decision to an economic one, where firms have to trade off costly information acquisition against the resulting improvement in choice quality. Although this added screening effort is not free, it could be justifiable from a social welfare perspective, given that there are likely positive externalities to more screening and hiring of relatively less experienced workers (Terviö, 2009; Pallais, 2014).

The paper is organized as follows: Section 2 describes the empirical context.

[4] "Philadelphia Is About to Ban Employers From Asking Potential Hires About Their Salary History," Fortune Magazine, January 20th, 2017. Accessed online on July 5, 2017, at: `http://fortune.com/2017/01/20/philadelphia-wage-history-employee-salaries/`.

[5] Kuhn and Shen (2013) find that firms' idiosyncratic gender preferences can be overridden by factors such as greater incentive to search broadly for the most qualified candidate.

Section 3 presents the experimental design. Section 4 discusses the related literature. Section 5 presents the results. Section 6 presents a structural model of employer hiring to explore how the experimental results would generalize to an equilibrium in which all employers lacked access to wage history. Section 7 concludes.

2 Empirical context

The context for our experiment is an online labor market. In online labor markets, employers hire workers to perform tasks that can be done remotely, such as computer programming, graphic design, data entry, research, and writing (Horton, 2010). Online labor markets differ in their scope and focus, but common services provided by the platforms include publishing job listings, hosting user profile pages, arbitrating disputes, certifying worker skills, and maintaining feedback systems.

There has been some research that uses online labor markets as an empirical context. Pallais (2014) conducted a field experiment to demonstrate how much value employers place on past on-platform work experience. Gilchrist et al. (2016) explore the effects of higher wages on output using a field experiment. Stanton and Thomas (2015) show that agencies (which act as quasi-firms) help workers find jobs and break into the marketplace. Agrawal et al. (2016) investigate which factors matter to employers in making selections from an applicant pool, and present some evidence of statistical discrimination; the paper also supports the view of employers selecting from a more-or-less complete pool of applicants rather than serially screening. Horton (2017c) reports results from a large scale minimum wage experiment.

2.1 Transacting on the platform

The process for filling a job opening on the platform is qualitatively similar to the process in conventional labor markets. First, a would-be employer creates

a job post.[6] An employer chooses a job title, writes a job description, and labels the job opening with a category (e.g., "Administrative Support") and required skills. Additionally, employers choose a contractual form (hourly or fixed-price).

Employers also must choose their relative preference for price and quality by selecting from one of the following options: Entry level ("I am looking for workers with the lowest rates."), Intermediate ("I am looking for a mix of experience and value."), and Expert ("I am willing to pay higher rates for the most experienced workers."). The employer's "vertical" preference selection is shown to would-be applicants.[7] za Drawing on the assortive matching in labor markets literature (Rosen, 1982; Sattinger, 1993), firms should tailor their hiring to attract the employees that generate the most match specific productivity for the task at hand. As such, we take firm's ex-ante vertical preferences over price and quality as a strong signal into importance of the project. Firms that indicate they are looking for "entry level" labor are essentially saying "this project isn't essential and so I'm ok with more productivity uncertainty." We will make use of these employer vertical preferences when exploring heterogeneity in the effects of the treatment.

Once the employer submits his or her job opening, it is reviewed by the platform and then posted publicly to the marketplace. Would-be applicants generally learn about job openings via electronic searches. Potential applicants can see the details of the job opening, as well as some information about the associated employer. If the worker chooses to apply, he or she submits a wage bid (for hourly jobs) or a total project bid (for fixed-price jobs) and a cover letter.

Employers can also seek out workers themselves, inviting workers to apply to their opening. To help employers find and evaluate workers, the platform hosts persistent worker "profiles." A profile page shows details about the worker's work history on the platform, skills, education, availability to take on more

[6]Employers also choose whether to make it public or private. Public jobs can be seen by all workers on the platform, while only invited applicants can see private jobs.

[7]See Horton and Johari (2015) on the effects of this feature on applicant sorting.

work, and other information that he or she wants to share.[8] A worker also lists his or her "profile rate," which is an hourly wage. Although it is self-reported, it is usually close to the typical wage that the worker earns, and employers consider it when deciding who to invite to their openings. The profile rate is a useful measure for our purposes, as it is recorded even if the worker has no work history on the platform. Furthermore, unlike average past wages, the profile rate is not "dragged" down by wages from jobs completed far in the past.

After applying, the applicant immediately appears in an interface the employer has for tracking applicants. This interface shows the applicant's bid, name, picture, self-reported skills, and a few pieces of platform-verified information, such as total hours-worked and average feedback rating from previous projects (if any). For these past projects, employers could, historically, see how many hours the worker worked on that project and, critically, his or her past wage. We will discuss how the treatment affected their ability to see this information when discussing the experimental design.

The employer can screen his or her applicants by asking questions and organizing interviews. After this screening, employers can decide to make an offer (or offers). Although employers typically extend an offer at the same wage as the worker's original wage bid, about 11% of workers are hired at a wage lower than the proposed wage due to back-and-forth wage bargaining.

Once hired, hours-worked are recorded using proprietary software that workers install on their computers. At the conclusion of the contract, both parties give a reason for ending the contract (typically that the project was completed successfully) and provide both written and numerical feedback about each other.

3 Experimental design

The experiment was conducted in late 2014 by the platform. All employers that posted a job opening during a 14-day period were allocated to the experiment.

[8] See Horton (2017a) on the importance of worker capacity information on probability of match formation.

The unit of randomization was the individual employer. All allocated employers were assigned to either the treatment group ($n = 2,974$) or the control group ($n = 2,948$). The sample sizes were determined by the platform ex ante. If an assigned employer posted an additional job opening, this job opening also received the treatment assignment of the initial job opening. However, we only use the first job opening by each employer in our analysis, as the treatment could have affected the probability of posting additional openings or the characteristics of any subsequent job opening.

We also restrict the sample to hourly job openings, as the bidding and hiring process of fixed price jobs is qualitatively different. Our sample is further restricted to only public jobs, which any applicant could bid on.[9] The change in the interface available to employers was not explained to treated employers; the interface simply changed.[10] Applicants to job postings were not aware of the experiment, and hence the possibility that the employer might not have access to their past on platform compensation history. Given that this wage history is visible to workers on their own profiles—and that it was historically available to employers—most workers presumably applied believing it would be available.

To assess balance, the means for a collection of pre-randomization attributes with respect to job opening characteristics, employer characteristics, and the composition of the applicant pool are shown in Appendix A.2. The experimental groups are well-balanced, which is unsurprising, as the software used to allocate employers to treatment cells has been used many times and has proven reliable.

[9] As a robustness check, we run our analysis on private jobs which are composed only of applicants expressly invited by the employer with whom the employer previously worked. There are no treatment effects on private jobs. These analyses are available from the authors upon request.

[10] We monitored employer discussion forums (which are generally not very active) and there was almost no discussion of the experiment. One employer did post about it, and others responded suggesting it was most likely a bug.

3.1 Employer's view of the applicant pool

A stylized representation of the employer's evaluation interface for an hourly job opening is shown below. Note that the employer can see the applying worker's name, hourly rate wage bid, average feedback rating, and on-platform experience, measured in hours of work completed. Critically, there is no information in this interface about the past hourly wage earned by the worker.

Name	Hourly Wage	Stars	Hours	Country
Ada H.	$18.00/hour	4.5	123	UK
Paul Y.	$15.00/hour	4.2	89	India

From this list, employers could "view" an application by clicking on it. An employer viewing an application would see that applicant's past work history. How this work history was presented differed by the employer's treatment assignment: in the control group, employers could see the past hourly wage associated with each past job held by the worker, but in the treatment, they could not. For example, a work history item for an applicant would be presented to a treated employer as:

Job Title:	Lead data scientist
Contract Type:	Hourly
Total:	$451.34
Time:	December 2014 - present

whereas a control employer viewing the same applicant would see:

Job Title:	Lead data scientist
Contract Type:	Hourly
Hourly Wage:	$17.00
Total:	$451.34
Time:	December 2014 - present

Note that the control employer could see that the worker worked at $17.00/hour, but a treated employer could not. Treated employers could not circumvent this restriction by searching for the worker and finding his or her information elsewhere on the platform—the employer's treatment assignment restricted access to this information everywhere. The availability of wage information in the control and the absence of such information in the treatment was the only treatment-specific difference in the interfaces and information presented to employers.

3.2 Constructing measures of evaluation

We measure whether an applicant was "viewed" by the employer, sent a message by the employer (i.e., "called-back"), asked a question by the employer, or (planned to be) interviewed by the employer "face-to-face," by scheduling a video teleconference session. To illustrate different measurements, consider an employer who received 6 applications:

Name	Wage Bid	Stars	Hours-worked	Country
Molly M.	$10.13/hour	4.6	563	Philippines
Ada H.	$6.15/hour	4.5	123	UK
Eliot G.	$6.10/hour	4.1	20	Russia
Julia M.	$7.16/hour	4.3	75	US
Paul H.	$8.27/hour	4.2	89	India
Emma G.	$7.16/hour	4.3	75	US

The employer clicked to learn more about Paul, Julia and Molly, and so the number of applicants "viewed" is three. After learning more about those three applicants, the employer also sent a message to Paul and Julia, making the number messaged just two. A message in this context can be thought of as an invitation to interview for a job opening, similar to a "call back" in the audit study literature.

Employers communicate with applicants through a platform-provided messaging system. All of the messages back and forth between one employer and one applicant are considered a message "thread." We search these message threads for a number of measures of employer evaluation and construct indicator variables of these measures: (1) setting up a face-to-face meeting by exchanging Skype IDs, (2) ending a sentence with a question mark, and (3) starting a sentence with a question word i.e., "when, where, why, or how."

Returning to our example, if the employers' message to Julia was:

> Hi Julia - Nice application, looking forward to working together!

This message would not have a question word, a question mark, or any attempt to set up a face-to-face meeting. In contrast, consider a message to Paul:

> Hi Paul — I'm interested in your application. Do you have much experience with my kind of project? When did you last use Python? My Skype ID is x12889—please get in touch to schedule an interview.

In the case of the message to Paul, we would have a question mark, a question word ("when"), and a Skype ID exchanged to set up a face-to-face meeting. We also search the text for words that indicate the nature of the messages (i.e., simply coordinating logistics or asking probing questions)—we will discuss those mesures later, when we present the experimental results.

3.3 Summary statistics on the hiring process

Summary statistics on hiring and screening in the control group are presented in Table 1. On average, 35 applicants apply to each job opening, and 1 of these applicants is invited to apply to the job by the employer, leaving about 34 "organic" applicants who apply to a job without being invited. Employers only view 7 of the applications submitted to the job by organic applicants, and only message, i.e., "call-back," about 2 of these applicants.

Turning to the interviewing phase, employers specifically ask at least one question to about 62% of the applicants they message. About half of applicants who are messaged are asked to conduct a "face-to-face" interview, at least as measured by appearance of the "Skype" keyword. On average, this hiring process leads to about 40% of job openings posted being filled within 6 months of being posted (on average 0.58 applicants are hired, as some jobs hire more than one applicant). This is not dissimilar to the traditional labor market—there is only a 44% chance of a job posting being filled within 30 days, and it is suspected that many of these openings are never filled.[11]

[11] Report by CEBR, http://press.indeed.com/wp-content/uploads/2015/01/Time-to-fill-jobs-in-the-US.pdf

Table 1: Per-opening summary statistics for the control group ($n = 2,948$)

Statistic	Mean	St. Dev.	Min	Median	Max
Applicants	35.105	43.296	0	22	639
Workers invited to apply	3.668	21.230	0	0	1,007
Invited workers that applied	1.414	4.504	0	0	175
Organic applicants	33.691	43.036	0	20.5	639
Applications viewed	7.321	9.257	0	5	122
Organic applications viewed	6.671	9.014	0	4	116
Organic applicants messaged	1.797	3.684	0	1	91
Org. appl's "questioned"	1.121	2.050	0	0	36
Org. appl's face-to-face interviewed	0.890	1.890	0	0	23
Applicants hired	0.580	1.039	0	0	26

Notes: This table reports summary statistics on the applicant pool characteristics and employer evaluation of their applicant pool in the control group. All reports are on a per-opening basis. For example, the the row labeled "Applicants" reports statistics on the number of applicants of different type. "Invited" workers are those that the employer sought out and asked to apply for the opening. "Organic" applicants are workers that applied without being invited. An application is "viewed" if the employer clicked on a worker's application to learn more about the applicant.

4 Related work and theoretical background

There is a large literature in labor and personnel economics on how information affects the hiring process. The economic problem of hiring is conceptually straightforward: the firm compares the marginal increase in revenue from a worker's labor to what they will have to pay to obtain that labor (Oyer et al., 2011). What makes this challenging from the firm's perspective is that a worker is, in a sense, an "experience good," and so the employer has to make an inference about productivity, relying on whatever signals they have available (Spence, 1973; Holzer, 1987). Even if the wage is bargained over, the bargaining solution will depend on the expected surplus generated, which in turn depends on the applicant productivity (Nash Jr., 1950; Binmore et al., 1986). Short of some extremely high powered incentive contract, there is no escape from the productivity inference problem.

Most of the empirical literature has focused on how the presence or absence of pieces of information seem to change the hiring process, as evinced by the

realized change in the characteristics of workers that are hired. A seminal paper on the effects of removing information from the hiring process is Goldin and Rouse (2000), who show that orchestras switching to blind auditions increased their fraction of female orchestra members. Presumably most of this change was due to blind auditions making taste-based discrimination infeasible. But to the extent that prior discrimination against female musicians was statistical, the results could indicate a shift towards relying on other signals—namely actual music ability as captured by a performance. This kind of "signal substitution," typically detected as change in the demographic composition of the workers evaluated or hired, is a common focus in the literature (Bertrand and Mullainathan, 2004; Autor and Scarborough, 2008).

A more recent example of information being removed from the hiring process comes from efforts to remove criminal history information from the application process—so called "ban the box" laws (Henry and Jacobs, 2007). To test the effects of a "ban the box" imposition, Agan and Starr (2016) use applicant call-back rates of fictitious applications sent before and after a "ban the box" policy went into effect. They find evidence that the policy caused signal substitution, with employers giving black applicants lower call-back rates after they could no longer directly screen on criminal records.

Shoag and Veuger (2016) reach a different conclusion, finding that "ban the box" policies seemed to work as intended, increasing employment among workers likely to be discriminated against on the probability of having a criminal record. Specifically, they find an increases in employment among African-Americans age 18-64. However, Doleac and Hansen (2016) find that the increase in hiring was concentrated among older black men, and that younger low-skill black (and Hispanic) males were harmed, which is consistent with a statistical discrimination story where the employer focus is on *recent* criminality.

In a paper conceptually similar to Shoag and Veuger (2016), Shoag and Clifford (2016) look at the effects of the rising use of consumer credit scores in the hiring process, finding that these likely increased total employment but harmed groups likely to have relatively poor credit scores.[12] Interestingly, the

[12]These findings are also consistent with Herkenhoff et al. (2016).

Shoag and Clifford (2016) finding of increased overall employment when credit scores are used highlights a concern with policies that alter what information employers have—if more information can increase hiring, less information might reduce it. As is common in this literature, the focus is on signal substitution as evinced by compositional changes in who is hired, or receives a call-back (in the case of Agan and Starr (2016)).

In our empirical setting, we can also look for signal substitution, but we do not have to rely solely on the attributes of hired workers. Instead, we can examine the employer's actual selection process, going from the applicant pool, to viewed applicants, to interviewed applicants, and finally to the hired applicant(s), if any.

Although this "full applicant pool" empirical setting is relatively rare, there are a few recent papers that can report on selection from the entire applicant pool. For example, Hoffman et al. (2015) studies the quasi-experimental introduction of a class of information into the hiring process, namely a signal from a job test. Cowgill (2017) also examines the selection process starting from the pool of applicants, comparing human and machine learning approaches to screening. Burks et al. (2015) assess the benefits to firms of hiring through employee referrals by looking at the pools of referred and non-referred applicants. Horton (2017b) considers the effects on application selection when employers have access to algorithmic recommendations about which workers to recruit. Barach et al. (2017) investigates when, how, and why employers choose to use algorithmic recommendations in hiring.

In the existing literature on the effects of information in hiring, researchers are mostly silent on the possibility that employers, when denied some class of information, might respond by collecting their own information. This absence of research attention is not because this margin is viewed as implausible—the idea of endogenous information acquisition has a long history in economics (Arrow, 1996; Grossman and Stiglitz, 1980; Stigler, 1961). Rather, little attention has been paid to endogenous information acquisition for the simple reason that in prior work, only the hiring outcomes are observed, not the hiring process itself. A key advantage of our empirical setting is that we can address this shortcoming

in the literature by directly measuring proxies for information acquisition by employers.

Endogenous information acquisition has been a feature of work used to analyze auctions (e.g., Milgrom and Weber (1982)), voting (e.g., Martinelli (2006); Persico (2004)), and medical patient decision-making (e.g., Kőszegi (2003)), among many other applications. In the labor literature, search-focused models have workers acquiring information about wages offered by different employers (Mortensen, 1970; McCall, 1970; Mortensen and Pissarides, 1999). Far less attention has been paid to information acquisition on the demand side of the labor market, though there are some exceptions. For example, there is cross-sectional work examining how employer and industry characteristics affect the time employers spend on recruiting and screening applicants (Barron et al., 1985, 1989).

5 Empirical results

We will present experimental results chronologically with respect to the hiring process, going from initial screening to post-hire contractual outcomes. As we have a true experiment, we will always present results as simple means comparisons at the job opening level, though when useful, we will switch to a regression framework.

5.1 Employer evaluation and information acquisition

Although the treatment does not affect the employer's initial view of the applicant pool, a treated employer might decide to click on, or "view" more or fewer applicants after observing less information from each applicant he or she views. In the top line of Table 2, labeled "*Measures of employer interest*," we can see that treated employers on average view another 0.45 applications from a baseline of 7 applications per opening, or about 7% more applicants than in the control.

In our experimental context, the equivalent of an interview call-back is the employer "messaging" an applicant. In Table 2, in the panel labeled "*Mea-*

Table 2: Effects of hiding applicant wage history on various outcomes

	Control	Treatment	Difference	% Change
Measures of employer interest (number of applicants)				
Viewed	6.67 (0.17)	7.12 (0.18)	0.45 (0.24)*	6.76
Measures of employer evaluation/elicitation (number of applicants)				
Messaged	1.80 (0.07)	1.93 (0.07)	0.13 (0.10)	7.09
Questioned (Q Word)	1.12 (0.04)	1.27 (0.05)	0.15 (0.06)**	13.40
Questioned (Q Mark)	1.19 (0.04)	1.31 (0.05)	0.12 (0.06)**	10.14
Face-to-Face Sched.	0.89 (0.03)	0.95 (0.04)	0.06 (0.06)	6.30
Characteristics of called-back (i.e., messaged) applicants				
Bid amount	13.96 (0.37)	13.12 (0.35)	-0.85 (0.51)*	-6.05
Profile wage rate	13.49 (0.31)	12.74 (0.31)	-0.75 (0.44)*	-5.55
Avg 6-month wage	11.70 (0.29)	10.87 (0.28)	-0.82 (0.40)**	-7.03
Min 6-month wage	9.38 (0.25)	8.68 (0.24)	-0.70 (0.34)**	-7.45
Max 6-month wage	14.91 (0.40)	13.96 (0.38)	-0.96 (0.55)*	-6.41
Previous hours worked	1189.01 (41.22)	1167.15 (38.79)	-21.86 (56.59)	-1.84
Prior billed jobs	30.63 (0.97)	29.02 (0.80)	-1.61 (1.25)	-5.25
Avg Feedback	4.71 (0.01)	4.71 (0.01)	-0.00 (0.01)	-0.09
Characteristics of hired applicants				
Bid amount	11.80 (0.40)	10.66 (0.50)	-1.15 (0.64)*	-9.73
Profile wage rate	12.11 (0.40)	11.02 (0.50)	-1.09 (0.64)*	-8.98
Avg 6-month wage	10.65 (0.54)	8.94 (0.43)	-1.70 (0.69)**	-16.01
Min 6-month wage	8.34 (0.35)	7.02 (0.37)	-1.31 (0.51)***	-15.77
Max 6-month wage	13.77 (0.91)	11.64 (0.55)	-2.12 (1.06)**	-15.42
Previous hours worked	1123.25 (56.21)	1212.01 (84.09)	88.76 (101.14)	7.90
Prior billed jobs	35.02 (1.70)	33.17 (1.45)	-1.86 (2.24)	-5.30
Avg Feedback	4.72 (0.01)	4.71 (0.02)	-0.01 (0.02)	-0.16
Job opening outcomes				
Hire made?	0.40 (0.01)	0.43 (0.01)	0.03 (0.01)**	7.23
Feedback (1-10)	8.72 (0.10)	8.31 (0.36)	-0.41 (0.38)	-4.68
Contract rated a success?	0.55 (0.02)	0.54 (0.03)	-0.01 (0.04)	-2.08
Wage Bargaining | hire made)				
Any bargaining?	0.11 (0.01)	0.10 (0.01)	-0.00 (0.02)	-3.91
Wage-to-bid | bargaining	0.84 (0.03)	0.93 (0.03)	0.08 (0.04)*	9.65

Notes: This table reports means errors across experimental groups. Next to each mean, standard errors are reported in parentheses. The "% Change" column is the percentage change in the treatment, relative to the control. Significance stars are calculated using p-values for a two-sided t-tests of the null hypothesis of no difference in means across groups. Significance indicators: $p \leq 0.10 : *$, $p \leq 0.05 : **$, and $p \leq .01 : ***$.

sures of employer evaluation/elicitation," the means for several outcomes are reported. Treated employers called-back about 7% more applicants, though this estimate is imprecise and not conventionally significant. Treated employers ask at least one question (as measured by a question word) to an additional 0.15 applicants per job opening, which corresponds to a 13% increase in the number of applicants questioned. The increase in questioning as measured by question mark is similar in magnitude.

As the information acquisition outcomes are counts of applicants, we can potentially gain more precision from a regression with the appropriate link function, and so we estimate a Poisson count regression

$$y_j = \lambda(\beta_0 + \beta_1 \textsc{WageHistHid}_j + \mathbf{X}_j\gamma + \epsilon_j), \tag{1}$$

where y_j is some outcome of interest, $\textsc{WageHistHid}_j$ is the treatment indicator, and $\mathbf{X}_j$ is a collection of pre-randomization job opening and employer characteristics.[13] We plot the coefficients on $\textsc{WageHistHid}_j$ in Figure 1. For each point estimate, a 95% CI is shown. For each outcome, we plot the coefficient using the full sample, labeled "Pooled" and for each of the three employer vertical preference levels, "Beginner," "Intermediate," and "Expert."

Reassuringly, the regression coefficients shown in Figure 1 for the "Pooled" sample give marginal effects similar to similar to those found in the means comparison: treated employers called-back 8.9% more applicants, used at least one question mark in 12.1% more message threads (and at least one question word in 15.6% more message threads), and set up 8.2% more face-to-face interviews as control employers.

Turning to the employer vertical preference sub-samples, we can see that the overall increase in information acquisition in the treatment is primarily driven by employers interested in hiring low and medium expertise applicants, with no evidence of a treatment effect for employers with "high" vertical preferences.

[13] We control for the category of the job opening, prior jobs billed by the employer, the employer's prior spend on the platform, the number of applications to the job openings, the number of recommended applications to the job opening, the average bid, and an indicator if the employer requested specific skills.

Figure 1: Effects of the treatment on extensive margin measure of employer information acquisition

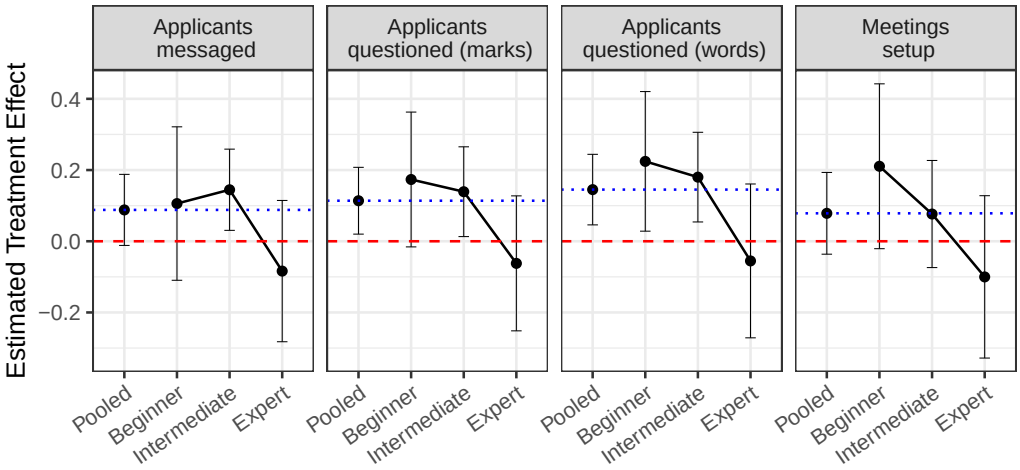

Note: This figure shows the relationship between measures of information acquisition and the treatment separated by requested expertise of the worker. The level of observation is the job opening. The model and controls are the same as used in the estimation of Equation 1, except that each employer vertical preference group is fit separately. Heteroskedasticity-robust standard errors are used for 95% confidence intervals.

The effects presented so far are all extensive margin effects, but we are also interested in intensive margin effect i.e., of the applicants they called-back, did treated employers evaluate them differently? Although we know the treatment changed the quantity of applicants called-back, we can still test whether employers engaged in more evaluation on a per-applicant basis by estimating the following application-level logit model:

$$y_{i,j} = \text{logit}^{-1}(\beta_0 + \beta_1 \textsc{WageHistHid}_j + \mathbf{X}_{i,j}\gamma + \epsilon_j) \Big| \textsc{CalledBack}_{i,j} = 1, \tag{2}$$

where $y_{i,j}$ is some outcome for applicant i to job opening j, such as whether they were asked a question, and $\mathbf{X}_{i,j}$ is a collection of pre-randomization job opening and applicant characteristics.[14] The sample is restricted to applicants the employer called-back.

Figure 2 plots odds ratios from applicant-level logit estimates of Equation 2, both for the entire sample (labeled "Pooled") and each of the vertical preference tiers. Coefficients are exponiated to be interpretable as odds ratios. The three panels of the figure report estimates for the question marks measure, the question words measure, and the face-to-face meeting setup.

As with the extensive margin estimates, treatment effects appear to be concentrated among employers with low and medium vertical preferences. For example, employers with the low vertical preferences are 1.51 times more likely to question an applicant they message (using the "question words" measure). In contrast, employers with the highest vertical preferences show no treatment effects. Across groups, there is no evidence of an intensive margin difference in face-to-face interviewing.

One potential explanation for the treatment effect on information acquisition being limited to employers who are looking for "entry level" and, to a lesser

[14] We control for the category, prior jobs billed by the employer, the employer's prior amount spent on the platform, the number of applications to the job openings, the number of recommended applications to the job opening, an indicator if the employer requested specific skills, the log of the applicant's bid, the log of the applicant's tenure, the number of prior jobs worked by the applicant, and the applicants' prior feedback.

Figure 2: Effects of the treatment on intensive margin information acquisition measures, by employer vertical preferences

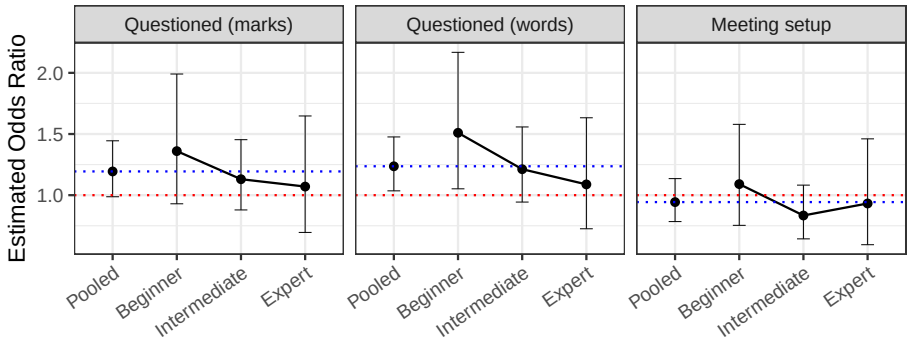

Note: This figure shows the relationship between intensity of information acquisition and the treatment, by the employer's reported vertical preference. The model and controls are the same as used in the estimation of Equation 2, except that each employer vertical preference group is fit separately. Heteroskedasticity-robust standard errors are used for 95% confidence intervals.

extent, "intermediate level" workers is that the value of locating and hiring a suitable worker is already high enough that those employers seeking "expert" labor are already engaging in higher levels of information acquisition. It is also possible that employers looking for "entry level" are closer to the "no surplus" margin described by Clemens and Wither (2014); when the treatment reduces cheap information, these employers now need to acquire more information to be willing to make a hire.

5.2 Characteristics of called-back applicants

In the absence of wage history information, treated employers might change the kind of applicants they evaluate. The panel of Table 2 labeled "*Characteristics of called-back applicants*" compares the mean characteristics of applicants who are called-back, by experimental group. The workers called-back by treated employers had lower wage bids, lower profile rates, and lower past wages. The effects are substantial—called-back workers in the treatment bid about 6% less, had 6% lower profile rates, and 7% lower average past wages. These bargain hunting effects are strongest on the average past wage measures—

which is calculated using precisely the wage measures that were hidden from employers in the treatment.

5.3 Probability of hiring and the characteristics of hired workers

Treated employers could have abandoned their job openings if their information deficit lowered the expected value of hiring below their reservation value. However, we find the opposite, with employers being somewhat more likely to make a hire. In the panel of Table 2 labeled "*Job opening outcomes*," we see that the treatment increased hiring by about 3 percentage points, from a baseline hire rate in the control group of 40%.

As we saw in Section 5.2, called-back applicants in the treatment have substantially lower past wages. This bargain hunting selection also carried over to hiring. The panel labeled "*Characteristics of hired applicants*" in Table 2 shows the strong shift towards workers with lower past wages in the treatment group: their hourly wage bids were nearly 10% lower and average past wages were about 16% lower.[15] These differences are substantially larger than the differences in mean attributes of the called-back applicants.

Of course, as we noted earlier, the treated group did have about a 7.2% higher fill rate, and so some of the decrease in the past wage could be due to selection. However, it is highly unlikely for the change to be purely due to selection—the additional filled job openings would have to pay hired workers *-137%* of the mean wage of hired workers in the control to get a 16% reduction overall.

The treatment induced bargain hunting among employers with respect to which workers were called-back and ultimately hired. To study how these effects differ by employer's vertical hiring preferences, we estimate an applicant-level

[15] We found no indication that treated employers were more likely to hire someone they had worked with in the past. However, this would be quite rare in our data, because if an employer already knew someone, they would have likely just contacted them directly with a private job opening, and we do not include private job openings in our sample.

selection model

$$\textsc{CalledBack}_{i,j} = \beta_0 + \beta_1 \log(\textsc{ProfileRate}_{i,j}) + \beta_2 \textsc{WageHistHid}_j + \beta_3 \left(\log(\textsc{ProfileRate}_{i,j}) \times \textsc{WageHistHid}_j\right) + \epsilon_{i,j} \Big| \textsc{Viewed}_{i,j} = 1, \quad (3)$$

for the whole sample, as well as for the three employer vertical preference tier sub-samples.

Figure 3 plots the predicted probability of being called-back versus the applicant profile rate for both treatment and control, faceted by the employer's vertical preference type. In the leftmost panel, which shows results for the full sample, we can see the treatment shifting call-back probability towards workers with relatively lower profile rates). In the low ("Beginner") and medium ("Intermediate") vertical preference panels, we mostly see a shift in evaluation towards relatively lower profile rate workers, whereas in the highest tier, we simply see slightly less hiring of the highest profile rate workers.

5.4 Reliance on other signals of productivity

One way in which policies that remove information from the hiring process can backfire is if employers put more weight on some other, correlated signal in their screening. In our setting, we can directly look at this signal substitution by estimating a model of the employer's selection decision. We compare the effect of five salient signals on the employer's probability of calling back a viewed applicant by the employer's treatment status. The five signals we analyze are the applicant's profile wage rate, the applicant's mean prior feedback score, the number of previous jobs an applicant has completed, the applicant's prior earnings, and the applicant's tenure on the platform.

To make comparing the effect on interviewing across signals of different types easier, we transform each viewed applicant's signal into a z-score which is normalized within a job opening. For example, if a job received only two applicants, one with 1 day of tenure and another with 2 days, we would give them tenure z-scores of -0.71 and 0.71, respectively.

Figure 3: Employer probability of calling back a worker conditional on the applicant log profile rate, the treatment assignment, and the employer vertical preference

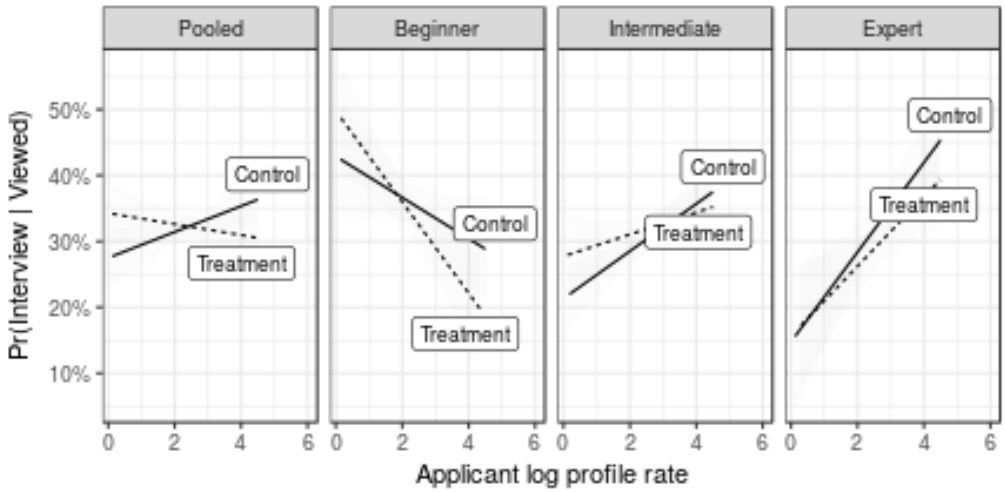

Note: This figure plots predicted values from a linear model which regresses if the employer called-back an applicant, conditional on viewing the applicant on the log of the applicant's profile wage rate by treatment status of the employer, separated by requested expertise of the worker.

Figure 4 reports the regression coefficients on each signal from regressions run separately for treatment and control employers, and by employer vertical preference, from a regression of the form:

$$\begin{aligned}\textsc{CalledBack}_{i,j} &= \beta_0 \textsc{WageBidZ}_{i,j} + \beta_1 \textsc{ProfileWageZ}_{i,j} \\ &+ \beta_2 \textsc{PriorFeedbackZ}_{i,j} + \beta_3 \textsc{NumPriorJobsZ}_{i,j} + \\ \beta_4 \textsc{PriorEarningsZ}_j + \beta_5 \textsc{TenureZ}_{i,j} + \gamma_j + \epsilon_{i,j} &\Big| \textsc{Viewed}_{i,j} = 1, \qquad (4)\end{aligned}$$

where γ_j is a job opening fixed effect. As we have shown, most other treatment effects differed by employer vertical preference, and so by looking for signal substitution by their preferences, we potentially have a more powerful test.

In Figure 4 from top to bottom, we can see that employers value prior jobs, prior earnings, and prior feedback: for each signal, for all vertical tiers and for both the treatment and control, the coefficients are positive. Employers do not rely very strongly on the applicant's profile wage rate, conditional on all the

Figure 4: Applicant characteristics on the probability call-back, by vertical preference tier

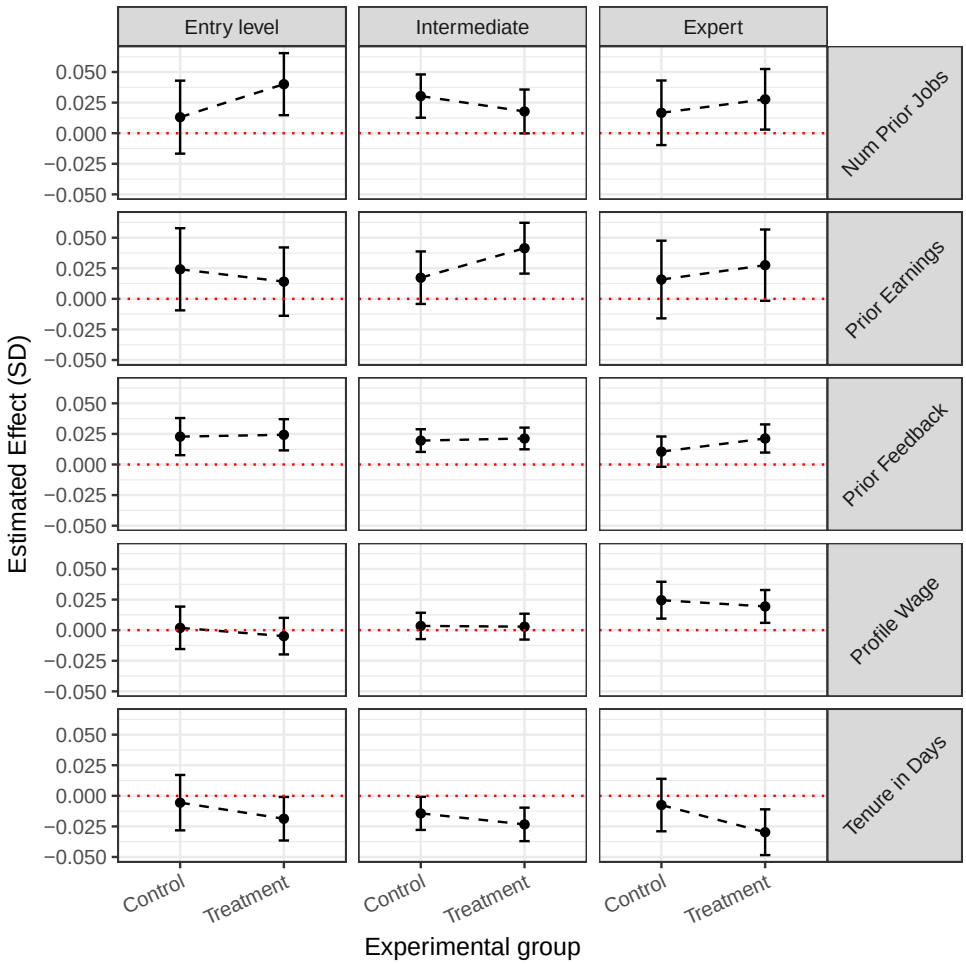

Note: This figure shows the coefficients on the non-wage bid parameters in Equation 4. The sample consists of viewed applications. Independent variables are transformed into z-scores normalized within job. Estimates represent change in standard deviations in probability of interviewing an applicant with 1 standard deviation increase in signal. Estimates are from linear probability models with job-level fixed effects. Heteroskedasticity-robust standard errors clustered at the job-level are used for 95% confidence intervals.

other signals, though the profile rate is highly correlated with the bid. Tenure is also an exception, with longer tenure having a negative effect. However, it is important to remember that these coefficients effects are all conditioned on all other effects. As such, the negative sign is unsurprising, as having joined the platform a long time ago is not a signal of quality per se, unlike, for example, having received good feedback.

Comparing within panels, there is no obvious pattern with respect to the treatment, offering no strong evidence of signal substitution. For example, employers with low vertical preferences in the treatment put more weight on the number of prior jobs, but medium-preference employers do not. Similarly, low-preference employers in the treatment seem to value earnings slightly less, while medium- and high-preference employers value it more, at least according to the magnitudes—the differences themselves are far from significant. Taken together, there is no strong evidence of signal substitution.

5.5 Prevalence and outcome of wage bargaining

We consider two aspects of wage bargaining: (1) whether any occurred among hired workers, as measured by a difference between the hired wage and what the worker initially proposed, and (2) the outcome of the wage bargaining, as measured by the ratio of the realized wage to the bid. In Table 2, the panel labeled "*Wage bargaining | hire made*" reports both outcomes. There is no strong evidence of a change in the fraction of worker/employer pairs that negotiate, but some evidence that hired workers in the treated group strike better wage bargains, conditional upon bargaining. However, given that we know the treatment encouraged the hiring of lower wage workers, this could simply be a selection effect. In a regression framework, we can potentially try to control for these composition change.

Column (1) of Table 3 reports the estimates of the linear regression:

$$\textsc{AnyBargaining}_j = \beta_0 + \beta_1 \textsc{WageHistHid}_j + \epsilon_j \big| \textsc{Hired}_j = 1.$$

Matching the means comparison result, we can see that the coefficient on the

treatment indicator is essentially a precisely estimated 0—there is no evidence that the treatment affected the prevalence of bargaining. From the constant term, we can see that for about 11% of filled job openings in which a hire was made, bargaining does occur. In those cases, we can look at whether the wage bargain is affected by the treatment, subject to the selection caveats described above.

Table 3: Effect of the treatment on the existence and outcomes of hired worker wage bargaining

	Dependent variable:		
	AnyBargaining	WageToBidRatio	
	(1)	(2)	(3)
Wage history hidden, WageHistHid	−0.005	0.089**	0.281**
	(0.016)	(0.043)	(0.122)
Applicant profile rate in logs (LPR)			0.008
			(0.034)
WageHistHid × LPR			−0.093*
			(0.054)
Constant	0.108***	0.836***	0.818***
	(0.012)	(0.029)	(0.083)
Observations	1,424	150	150

Notes: This table reports regressions where the outcomes are measures of bargaining. In Column (1), the outcome is whether any wage bargaining occurred. The sample is restricted to employers that made a single hire. In Columns (2) and (3), the outcome is the ratio of the realized wage to the initial wage bid. The sample for these two regressions are only those hires for which some bargaining ocurred. Heteroskedastic robust standard errors are reported. The top and bottom .5% of wage-to-bid ratios are dropped. Significance indicators: $p \leq 0.10 : *$, $p \leq 0.05 : **$, and $p \leq .01 : ***$..

Selection caveats aside, Column (2) reports the estimates of the linear regression:

$$\textsc{WageToBidRatio}_j = \beta_0 + \beta_1 \textsc{WageHistHid}_j + \epsilon_j \Big| \textsc{AnyBargaining}_j = 1.$$

The coefficient on the treatment indicator is positive and significant, with the treatment increasing the wage-to-bid ratio by about 9%, from a baseline ratio

of 0.84.

Given that the treatment increased employer interest in relatively low-wage workers, a natural question is whether the bargaining effects are concentrated among those workers. In Column (3), we interact the treatment indicator with the applicant's log profile rate. The effects on the interaction term are negative and conventionally significant, implying that workers with relatively low wages saw the largest increase in their bargained wage. We can also see this in the larger coefficient on the $\textsc{WageHistHid}_j$ indicator. The coefficient on the log profile rate is a precisely estimated zero, which undercuts the notion that lower profile workers simply have better bargaining outcomes on the ratio measure (which would be a concern given that the treatment causes bargain-hunting).

These bargaining results might explain why hired workers in the treatment group had 16% lower past wages, but only about 10% lower hired wages—the bargaining effects give those hired workers somewhat higher-than-expected wages, given their history.

5.6 Contractual outcomes

The treatment induced employers to hire workers with lower past average wages. If these hired workers were "worse," this might have reduced the quality of the work the employer received. We cannot separately identify these two effects, but we do have two measures of contract outcomes to look at: (1) an indicator whether the employer reported the job opening being completed successfully, and (2) the numerical feedback the employer left for the worker (1 to 10 scale). In Table 2, in the panel labeled "*"Job outcomes*," we can see that average feedback was slightly lower in the treatment, but a 95% confidence interval would comfortably include zero.

6 Equilibrium considerations

In the experiment, workers did not know that employers would lack access to their wage history. With an actual market-wide policy change, workers would know that their wage history is hidden, and could respond. Workers could

respond in many ways, such as altering what jobs they apply to or how they present themselves to employers. However, perhaps the most consequential way of responding would be to adjust their wage bids up or down. This could potentially turn some of the compositional effect we observe into a price effect.

Given our findings, the most likely wage response would be for relatively low wage workers to adjust their wage bids up, and for relatively high wage workers to adjust their bids down. However, the desirability of any one worker increasing his or her wage bid depends on how much competition they face on a per-job basis from similarly situated workers.

To determine what is likely to happen in equilibrium, we model both the employer's hiring problem and the workers' bidding decision. We begin by modeling the employer's hiring decision. If all employers were homogeneous and workers differed only in their productivity and wage bids, then the utility employer j derives from hiring worker i is $\pi_j = u_0 - w_i/y_i$, where u_0 is the value the employer derives from the completed project, w_i is the wage paid to employee i, and y_i is the productivity of employee i. Profit maximizing employers thus care about maximizing $\log y_i - \log w_i$. We assume that the employer error in inferring productivity, plus any un-modeled heterogeneity in worker job-specific productivity, is captured by ϵ_i, which is distributed iid extreme value for all values of i. The employers choice function follows a logit choice function and can be modeled as a discrete choice problem, with the probability of hiring worker i being:

$$\Pr(\textsc{Hired}_i = 1) = \frac{e^{x_i'\beta}}{\sum_a e^{x_a'\beta}}, \tag{5}$$

where a indexes the other applicants to that job opening and $x = [\log \hat{y}, \log w]$. We will assume that the firm makes a single hire.

6.1 Employer preferences

In the data, we observe a collection of w_i for each worker, as this is the worker's hourly wage bid for each opening. We do not observe perceived productivity,

y_i, directly, but we estimate it from workers' wage bids. Consistent with the market being approximately competitive, we assume that worker's bids are proportional to their expected marginal productivity given their attributes, plus some idiosyncratic error. For predictors, we use the worker's average feedback to date, log mean wage over the last six months, log cumulative prior earnings, log cumulative prior hours, number of past contracts, tenure, and whether they are affiliated with an agency and all the pair-wise interactions of these predictors. We then label each worker with the prediction from this model.

With measures of both w_i and y_i, we can estimate Equation 5 by maximum likelihood. We restrict the sample to job openings where exactly one hire was made, and where the number of applicants was two or more.[16] We estimate the β coefficients separately for treatment job openings and for the control job openings. Table 4 presents the estimated β coefficients for both the control and treatment groups. Although the differences in parameter point estimates are not conventionally significant, we can see that treated employers put relatively more weight on the wage bid and less on the perceived productivity, which is consistent with the bargain hunting effect we observed.

Table 4: Estimated Coefficent Vector

	$\hat{\beta}_y$	$\hat{\beta}_w$
Control	0.493(0.126)	−0.158(0.117)
Treatment	0.429(0.119)	−0.233(0.111)

Notes: This table reports the estimated β coefficent vector for treatment and control job postings. Estimates are generated from Equation 5 using maximum likelihood estimation.

Using these estimated coefficients we can compute for each worker his or her probability of being hired in both the treatment and the control group, given

[16] We also estimated the model by adding a "not hire" option to the choice set, giving it a value of β_0. This approach performed very poorly, yielding a massive β_0. The likely reason is that job openings that go unfilled are likely due to idiosyncratic factors pertaining to the employer rather than a poor collection of applicants.

that workers wage bid using Equation 5. A worker's predicted change in hiring probability for each job opening they apply to is thus:

$$\Delta p_i = \Pr(\textsc{Hired}|w_i, \hat{\beta}_{WHH=1}) - \Pr(\textsc{Hired}|w_i, \hat{\beta}_{WHH=0}). \tag{6}$$

Note that we are making use of the coefficients from the employer's fitted choice model, where $\hat{\beta}_{\text{WHH=1}}$ is the coefficient for the treatment group where the wage history is hidden and $\hat{\beta}_{\text{WHH=0}}$ are the estimated coefficients for the control. In the top panel of Figure 5, the solid curve shows the change in probability of being hired for workers of various predicted productivities due to employers being prohibited from observing past wages. We can clearly observe the bargain hunting: the effect of the treatment on probability of being hired is bigger for workers with relatively lower estimated productivity than it is for workers with relatively high estimated productivities.

6.2 Wage bid adjustment

We now consider how workers would adjust their wage bids when they know they are facing an employer with "treatment preferences" and when competing with other similarly situated job applicants. We model workers as adjusting their wage bid by an own-productivity dependent multiplier function, $m(\hat{y}_i)$. As a function form for $m()$, we assume

$$m(\hat{y}_i) = e^{\gamma_0 + \gamma_1 \log \hat{y}_i}, \tag{7}$$

where γ_0 and γ_1 are parameters to estimate. If both γ_0 and γ_1 are zero, then $m(\hat{y}_i) = 1$, and there is no adjustment. Given our "bargain hunting" findings, we expect that relatively low productivity workers will bid up, but that relatively high productivity workers will bid down, implying $\gamma_0 > 0$ but $\gamma_1 < 0$.

To estimate $m(\hat{y}_i)$, we must first describe the worker's wage bidding problem. For a given job opening, they choose a bid w_i that maximizes their payoff,

Figure 5: Estimated effects on individual worker's probability of hire

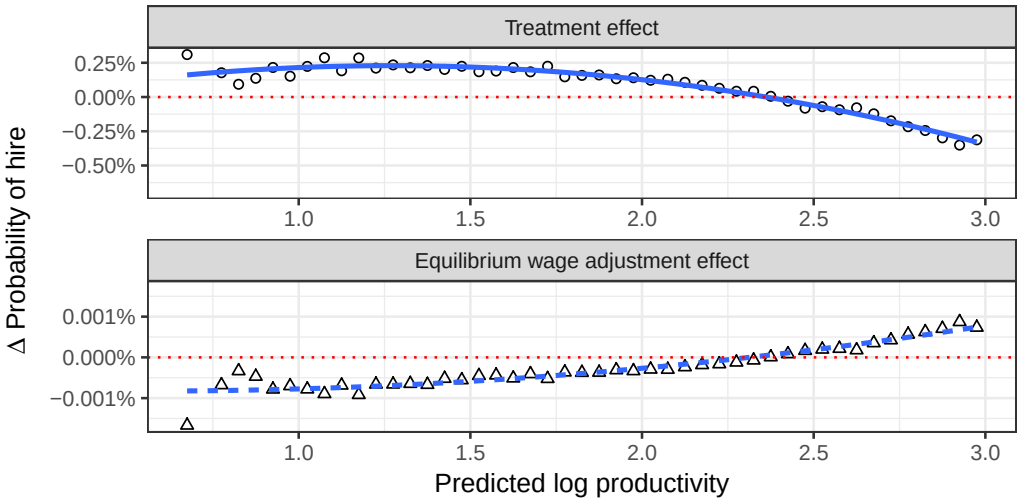

Note: This figure shows the relationship between worker productivity and the change in probability of being hired due to (1) the employer's ability or lack thereof to observe past wages, in the top panel, and (2) the worker's equilibrium wage adjustment, in the bottom panel. Plotted points are averages within $\log \hat{y}_i$ bins of size .005. In the top panel, the solid curve plots a quadratic line of best fit of change in probability of being hired for workers of various predicted productivities due to employers being prohibited from observing past wages. In the bottom panel, the dashed curve plots a quadratic line of best fit of change in probability of being hired for workers of various predicted productivities due to worker's equilibrium wage adjustment due to employers that cannot observe past wages.

or

$$\arg\max_{w_i} \quad p(w_i)(w_i - c_i)h,$$

where h is the hours of work required, $p(w_i)$ the probability of hire as a function of the wage bid and c_i is the opportunity cost of an hour of work. The optimal wage bid satisfies

$$w^* = -\frac{p(w^*)}{p'(w^*)} + c_i.$$

In our experiment, employers being unable to observe past wages shifts the probability of being hired at a given wage, $p(w_i)$, curve for each worker. For relatively low wage workers, the $p(w_i)$ curve was shifted out, leading to an increase in probability of being hired. For relatively high wage workers, the $p(w_i)$ curve was shifted in, leading to a decrease in probability of being hired For a small shift in the curve caused by the experiment, the optimal shift in the wage bid can be calculated by taking the partial derivative with respect to $p(w_i)$ and reorganizing:

$$\mathrm{d}w^* = -\frac{\mathrm{d}p(w)}{p'(w)}.$$

We can now see that the optimal shift in bid is equal to the change in the probability of being hired scaled by the inverse slope of the demand curve. We can also express the optimal adjustment in the wage bid, as a percentage change, or

$$\frac{\mathrm{d}w^*}{w^*} = \frac{\mathrm{d}p}{p(w) - c_i p'(w)}.$$

Empirically, this optimality condition gives us a first order condition for each worker:

$$\mathbb{E}\left[\frac{\mathrm{d}p}{p - c_i p'(w)} - \frac{\mathrm{d}w}{w}\right] = 0. \tag{8}$$

To use Equation 8, we need estimates of $p'(w)$ and c. For $p'(w)$, we can exploit the matched worker-employer nature of our data and the fact that workers vary their wage bids, despite the fact that their productivity is presumably fixed during relatively short periods of time, such as during our experiment. In Appendix A.4, using both worker and job opening fixed effects, we obtain an estimate of $p'(w) = -0.011$. For the reservation wage, c_i, we rely on the literature on the relationship between actual and reservation wages and use $c_i = 0.9w_i$, but we try a range of values.[17]

We estimate the the γ vector from Equation 7 by minimizing the square error in the sample analog of Equation 8. In short, for a given γ, we calculate the change in each worker's wage bid, then compute the change in individual hire probability (which depends on the wage bids of all other workers). We then compute the total squared error from Equation 8. The γ parameters that minimize this error are $\gamma_0 = 0.0078$ and $\gamma_1 =$ -0.0014.

As γ_0 is positive and γ_1 is negative, for low levels of worker productivity, the worker's equilibrium wage adjustment multiplier will be greater than 1, leading to an increase in the worker's optimal equilibrium bid. But as productivity increases, the worker's equilibrium wage adjustment multiplier will decrease but still remain positive (except at very high productivity levels, $\log \hat{y}_i \geq 6$). However, the implied adjustments are small. Workers with estimated log productivities of 0.66 adjust their wage bids up by 0.69% while workers with estimated productivities of 2.19 adjust their wages up by 0.47% and workers with estimated productivities of 2.53 adjust their wages up by 0.43%.

The bottom panel of Figure 5 shows how little of an effect wage bidding adjustments have on hire probability. As expected, worker's equilibrium wage adjustment decreases the probability of being hired for relatively low produc-

[17] Krueger and Mueller (2016) use data from a survey of recently unemployed workers in New Jersey. The data they collected shows that workers report reservation wages which are about 98% of the previous wages the week they become unemployed. Brown and Taylor (2013) use data from the British Household Panel Survey which asked recently unemployed workers both for their reservation wage and the wage they expect to make upon returning to work. The ratio between log reservation wage and log expected wage is .929. An older study by Lancaster and Chesher (1983) analyzed both the 1974 "National Survey of the Unemployed" and the 1973 "Men out of Work (Oxford)" surveys, and found that the ratio of reservation wages to expected wages were .89 and .84 in the two surveys respectively.

tivity workers and increases the probability of being hired for relatively high productivity workers. However, these effects are vanishingly small relative to the change induced by the treatment, shown in the top panel of the same figure. For example, a worker with a predicted log productivity of 1 is about 0.19% more likely to be hired when employers are unable to observe past wage histories. A maximizing worker will increase her bid by about 0.64%, reducing her probability of being hired by only about 0.00074%.

We conclude that almost none of the bargain hunting we observed in the experiment would be eroded in equilibrium. Given the lack of changes in the wage bidding, we expect that we would also observe the same increase in evaluation and hiring.

7 Conclusion

This paper highlights how employers react when wage history is removed. We find that they react by acquiring more information, expanding their evaluation on both the extensive and intensive margins. We find no evidence that they simply abandon their search or rely more heavily on other signals. The strongest finding of the paper is a marked shift towards evaluating and hiring lower wage workers.

Subject to the caveat of generalizing to other contexts, our findings suggest that policies that limit employer access to wage history would more or less have the intended effects, benefiting those with relatively low wages. These workers would benefit both from being more likely to be evaluated by employers, and perhaps also by being able to strike a better wage bargain. However, we also show that not all employers are equally pliable with respect to whom they consider: employers with relatively "low" vertical preferences were the employers who responded by more information acquisition and a shift in hiring. This finding suggests that perhaps policies prohibiting asking about wage history could be usefully targeted towards entry level positions.

References

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

Agrawal, Ajay K, Nicola Lacetera, and Elizabeth Lyons, "Does information help or hinder job applicants from less developed countries in online markets?," *Journal of International Economics*, 2016, *103*, 1–12.

Altonji, Joseph G and Charles R Pierret, "Employer learning and statistical discrimination," *The Quarterly Journal of Economics*, 2001, *116* (1), 313–350.

Arrow, Kenneth J, "The economics of information: An exposition," *Empirica*, 1996, *23* (2), 119–128.

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

Barach, Moshe A., Ming D. Leung, and Sibo Lu, "Foot in the door or head and shoulders above the rest? How algorithmic recommendations affect employer search and hiring behavior," *Working paper*, 2017.

Barron, John M, Dan A Black, and Mark A Loewenstein, "Job matching and on-the-job training," *Journal of Labor Economics*, 1989, *7* (1), 1–19.

_, John Bishop, and William C Dunkelberg, "Employer search: The interviewing and hiring of new employees," *The Review of Economics and Statistics*, 1985, *67* (1), 43–52.

Bertrand, Marianne and Sendhil Mullainathan, "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination," *The American Economic Review*, 2004, *94* (4), 991–1013.

Binmore, Ken, Ariel Rubinstein, and Asher Wolinsky, "The Nash bargaining solution in economic modelling," *The RAND Journal of Economics*, 1986, pp. 176–188.

Blake, Thomas and Dominic Coey, "Why marketplace experimentation is harder than it seems: The role of test-control interference," in "Proceedings of the fifteenth ACM conference on Economics and Computation" ACM 2014, pp. 567–582.

Brown, Sarah and Karl Taylor, "Reservation wages, expected wages and unemployment," *Economics Letters*, 2013, *119* (3), 276 – 279.

Burks, Stephen V, Bo Cowgill, Mitchell Hoffman, and Michael Housman, "The value of hiring through employee referrals," *The Quarterly Journal of Economics*, 2015, *130* (2), 805—839.

Clemens, Jeffrey and Michael Wither, "The minimum wage and the great recession: Evidence of effects on the employment and income trajectories of low-skilled workers," *Working Paper*, 2014.

Cowgill, Bo, "Automation and subjective decisionmaking: Theory and evidence from resume screening," *Working Paper*, 2017.

Doleac, Jennifer L and Benjamin Hansen, "Does "ban the box" help or hurt low-skilled workers? Statistical discrimination and employment outcomes when criminal histories are hidden," *Working Paper*, 2016.

Gilchrist, Duncan S, Michael Luca, and Deepak Malhotra, "When 3+ 1 > 4: Gift structure and reciprocity in the field," *Management Science*, 2016, *62* (9), 2639–2650.

Goldin, Claudia and Cecilia Rouse, "Orchestrating impartiality: The impact of "blind" auditions on female musicians," *The American Economic Review*, 2000, *90* (4), 715–741.

Grossman, Sanford J and Joseph E Stiglitz, "On the impossibility of informationally efficient markets," *The American Economic Review*, 1980, *70* (3), 393–408.

Hall, Robert E and Alan B Krueger, "Evidence on the incidence of wage posting, wage bargaining, and on-the-job search," *American Economic Journal: Macroeconomics*, 2012, *4* (4), 56–67.

Henry, Jessica S and James B Jacobs, "Ban the box to promote ex-offender employment," *Criminology and Public Policy*, 2007, *6* (4), 755.

Herkenhoff, Kyle, Gordon Phillips, and Ethan Cohen-Cole, "The impact of consumer credit access on employment, earnings and entrepreneurship," *Working Paper*, 2016.

Hoffman, Mitch, Lisa B Kahn, and Danielle Li, "Discretion in hiring," *Working Paper*, 2015.

Holzer, Harry J, "Hiring procedures in the firm: their economic determinants and outcomes," *Working Paper*, 1987.

Horton, John Joseph, "Online Labor Markets," *Internet and Network Economics*, 2010, p. 515.

_ , "Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution," *Working Paper*, 2017.

_ , "The effects of algorithmic labor market recommendations: Evidence from a field experiment," *Journal of Labor Economics*, 2017.

_ , "Price floors and employer preferences: Evidence from a minimum wage experiment," *Working Paper*, 2017.

_ **and Ramesh Johari**, "At what quality and what price?: Eliciting buyer preferences as a market design problem," in "Proceedings of the Sixteenth ACM Conference on Economics and Computation" ACM 2015, pp. 507–507.

Kahn, Lisa B and Fabian Lange, "Employer learning, productivity, and the earnings distribution: Evidence from performance measures," *The Review of Economic Studies*, 2014, *81* (4), 1575–1613.

Kőszegi, Botond, "Health anxiety and patient behavior," *Journal of Health Economics*, 2003, *22* (6), 1073–1084.

Kotlikoff, Laurence J and Jagadeesh Gokhale, "Estimating a firm's age-productivity profile using the present value of workers' earnings," *The Quarterly Journal of Economics*, 1992, *107* (4), 1215–1242.

Krueger, Alan B. and Andreas I. Mueller, "A contribution to the empirics of reservation wages," *American Economic Journal: Economic Policy*, February 2016, *8* (1), 142–79.

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

Lancaster, Tony and Andrew Chesher, "An econometric analysis of reservation wages," *Econometrica*, 1983, *51* (6), 1661–1676.

Lange, Fabian, "The speed of employer learning," *Journal of Labor Economics*, 2007, *25* (1), 1–35.

Martinelli, César, "Would rational voters acquire costly information?," *Journal of Economic Theory*, 2006, *129* (1), 225–251.

McCall, John Joseph, "Economics of information and job search," *The Quarterly Journal of Economics*, 1970, *84* (1), 113–126.

McDonald, Paul, Matt Mohebbi, and Brett Slatkin, "Comparing Google consumer surveys to existing probability and non-probability based internet surveys," *Google White Paper*, 2012.

Milgrom, Paul R and Robert J Weber, "A theory of auctions and competitive bidding," *Econometrica*, 1982, *50* (5), 1089–1122.

Mortensen, Dale T, "Job search, the duration of unemployment, and the Phillips curve," *The American Economic Review*, 1970, *60* (5), 847–862.

_ **and Christopher A Pissarides**, "New developments in models of search in the labor market," *Handbook of Labor Economics*, 1999, *3*, 2567–2627.

Nash Jr., John F., "The bargaining problem," *Econometrica*, 1950, pp. 155–162.

Oyer, Paul, Scott Schaefer et al., "Personnel economics: Hiring and incentives," *Handbook of Labor Economics*, 2011, *4*, 1769–1823.

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

Papke, Leslie E and Jeffrey Wooldridge, "Econometric methods for fractional response variables with an application to 401 (k) plan participation rates," *Journal of Applied Econometrics*, 1996, *595* (11), 619–632.

Pennebaker, J.W., R.J. Booth, R.L. Boyd, and M.E. Francis, "Linguistic inquiry and word count: LIWC2015.," 2015.

Persico, Nicola, "Committee design with endogenous information," *The Review of Economic Studies*, 2004, *71* (1), 165–191.

Rosen, Sherwin, "Authority, control, and the distribution of earnings," *The Bell Journal of Economics*, 1982, pp. 311–323.

Sattinger, Michael, "Assignment models of the distribution of earnings," *Journal of Economic Literature*, 1993, pp. 831–880.

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

_ **and Stan Veuger**, "No woman no crime: Ban the cox, employment and upskilling," *Working Paper*, 2016.

Spence, Michael, "Job market signaling," *The Quarterly Journal of Economics*, 1973, *87* (3), 355.

Stanton, Christopher and Catherine Thomas, "Landing the first job: The value of intermediaries in online hiring," *The Review of Economic Studies*, 2015, *83* (2).

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

Terviö, Marko, "Superstars and mediocrities: Market failure in the discovery of talent," *The Review of Economic Studies*, 2009, *76* (2), 829–850.

A Online Appendix: Not for publication

A.1 Survey evidence employer compensation history usage

Following Hall and Krueger (2012), we ran two nationally representative surveys to shed light on match formation and wage bargaining in the conventional labor market. We had two goals for our surveys. First, we wanted to determine how frequently firms ask about applicant compensation history, and when in the process they ask. If wage history is asked only after a job offer is made, the firm can only be using it for bargaining purposes. In contrast, if the firm asks before making an offer, they can use it during the screening process and during the bargaining process. The latter matches the scenario in our empirical context, in the control group. Second, we wanted to determine how frequently the worker is the first one to make the wage offer, as is the case in our empirical context.

We ran our surveys on Google Surveys, an online marketing research service that compares favorably to other Internet-based panels (McDonald et al., 2012). In our first survey, we asked subjects:

> *In the last job that you interviewed for, did the employer ask about your past wage/salary history?*

with answer options of:

- *No*
- *Yes, before I was offered a job*
- *Yes, after I was offered a job*

The results of this survey are publicly available.[18] For this survey, we received responses from 391 subjects with demographic weights. Of those, 115 reported they were asked about their wage history, or 29.4%. Among those asked, 82.6%

[18] https://surveys.google.com/reporting/survey?survey=tfqrbh2keackwznfzkwtzgp45a.

Figure 6: Survey evidence on when compensation history is asked about and the order of wage bargaining

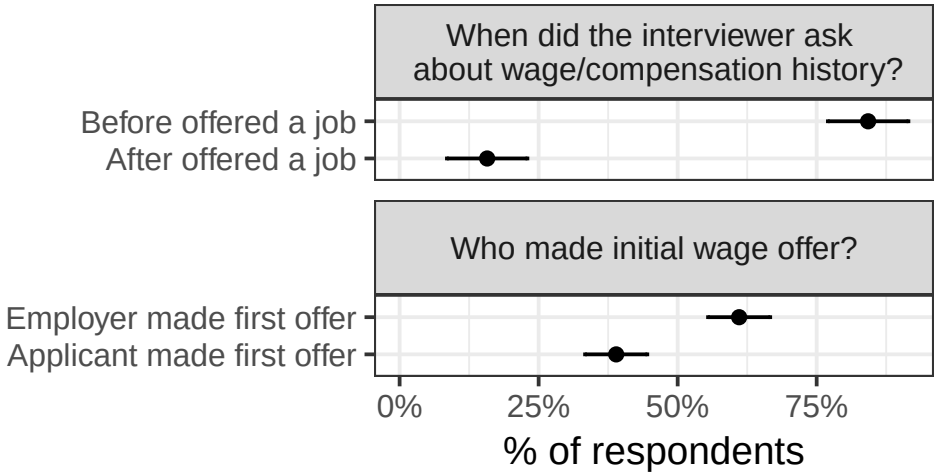

Note: This figure reports population-weighted means from two surveys conducted on Google Surveys. In the top panel, the sample is restricted to respondents reporting that in the last job they interviewed for the employer asked for their compensation history. Respondents answered whether the employer elicited this information before an offer was made or only afterwards. In the bottom panel, the sample is restricted to subjects who reported they bargained over wages in their last job. Respondents answered whether they made the initial offer or the employer made the initial offer.

report the firm asked about wage/compensation before extending a job offer. The population-weighted fraction is shown in Figure 6, in the top panel, with weighted standard errors, which is very close to the unweighted fraction. This first survey shows that asking about wage history is fairly common. Additionally, when employers ask about past compensation, it is much more likely to be asked upfront, presumably because it is used in evaluation, or in forming expectations about the likely outcome of bargaining.

In our second survey, we asked the question:

> *In your current job, did you bargain with your employer over compensation/benefits, and if so, who made the first offer?*

with answer options of:

- *No - wage was known when I applied*

- *Yes - bargained, and I made first offer*
- *Yes - bargained, and firm made first offer*

The results of this survey are publicly available.[19] For this survey, we received responses from 1,509 subjects with demographic weights. Of those, 316 reported they bargained over wages, or 29.4%. Among those asked, 39.2% report they were the first to propose a wage. The population-weighted fraction is shown in Figure 6, in the bottom panel, with weighted standard errors, which is very close to the unweighted fraction. Clearly, among workers bargaining, it is more common for the firm to make the first offer. However, a non-trivial fraction of conventional market bargained outcomes have a bargaining structure similar to our empirical context.

A.2 Randomziation

Table 5 reports the means for a collection of pre-randomization attributes with respect to job opening characteristics, employer characteristics, and the composition of the applicant pool. The groups were well-balanced, with only one pre-treatment characteristic difference in means being conventionally statistically significant—and only marginally so—despite almost 40 covariates used to assess balance.

As with any experiment conducted in a true market, there is a concern about violations of the stable unit treatment value assumption (Blake and Coey, 2014). However, given the short duration of the experiment, we view market-moving equilibrium effects as highly unlikely. The balance in applicant pool composition shown in Table 5 supports this contention.

A.3 Message text analysis

One limitation of "question based" measures of information acquisition is that they might also be capturing coordination-related questioning. For example, our measures would regard "You're hired—when can you start?" as a measure

[19] https://surveys.google.com/reporting/survey?survey=z5eldvvypuvrco4zvo2fkirreq

Table 5: Employer, job opening, and applicant characteristics by experimental group

	Control	Treatment	Difference	% Change
Employer attributes				
Prior job openings	23.49 (0.90)	23.98 (0.95)	0.49 (1.31)	2.10
Prior billed jobs	10.71 (0.43)	11.29 (0.46)	0.58 (0.63)	5.45
Prior spend by employers	5643.10 (312.83)	6053.35 (328.90)	410.25 (453.99)	7.27
Num prior contractors	10.84 (0.46)	11.73 (0.65)	0.89 (0.80)	8.18
Avg feedback score of employer	4.81 (0.01)	4.79 (0.01)	-0.03 (0.02)*	-0.55
Num of reviews of employer	8.05 (0.40)	8.84 (0.59)	0.79 (0.71)	9.82
Job opening attributes				
Number non-invited applicants	33.62 (0.80)	33.44 (0.74)	-0.18 (1.09)	-0.53
Avg best match score	0.36 (0.00)	0.36 (0.00)	0.00 (0.00)	0.83
Avg bid	12.76 (0.17)	12.60 (0.17)	-0.16 (0.24)	-1.23
Prefered experiance in hours	33.69 (2.43)	34.25 (2.37)	0.56 (3.40)	1.65
Estimated job duration in weeks	17.19 (0.39)	16.93 (0.39)	-0.26 (0.55)	-1.53
Applicant attributes				
Tenure in days	868.81 (1.98)	866.77 (2.74)	-2.03 (3.38)	-0.23
Hours worked to date	1212.73 (8.39)	1207.21 (11.56)	-5.53 (14.28)	-0.46
Num past jobs worked	33.59 (0.19)	33.81 (0.28)	0.22 (0.34)	0.67
Past hourly earnings	9812.41 (87.48)	9706.79 (121.18)	-105.62 (149.45)	-1.08
Past fixed wage earnings	2035.25 (17.94)	2027.06 (25.70)	-8.19 (31.34)	-0.40
Num prior employers	25.80 (0.13)	26.01 (0.20)	0.21 (0.24)	0.80
Wage bid	10.99 (0.07)	10.94 (0.10)	-0.05 (0.12)	-0.48
Profile wage	10.69 (0.06)	10.64 (0.09)	-0.05 (0.10)	-0.50
Min hr. wage (6 months)	6.93 (0.04)	6.89 (0.06)	-0.04 (0.08)	-0.63
Avg hr. wage (6 months)	8.46 (0.05)	8.42 (0.07)	-0.04 (0.09)	-0.48
Max hr. wage (6 months)	10.56 (0.06)	10.54 (0.09)	-0.02 (0.11)	-0.20

Notes: This table reports means for a number of pre-randomization characteristics for the employer, job opening, and applicant pool, by experimental group. Standard errors are reported next to the estimate, in parentheses. The far right column also reports the percentage change in the treatment group, relative to the mean in the control group. For the wage measures in "Applicant attributes," the 99th and 1st percentile are removed as outliers, as s very small number of wage bids are not bona fide hourly wage proposals, but instead place-holders with "wages" of \$0.01/hour or \$999/hour. In the bottom panel, standard errors are clustered at the job opening level, as applicants are nested within job openings. The associated significance stars on the difference are for a two-sided t-tests of the null hypothesis of no difference in means across groups. Significance indicators: $p \leq 0.10 : *$, $p \leq 0.05 : **$, and $p \leq .01 : ***$.

of information acquisition, even though the employer is not assessing the applicant's productivity. To address this shortcoming, we also look for words in messages that would more likely indicate continued probing.

We first look for evidence of employers directly asking for past compensation history information. Although treated employers asked more questions, surprisingly, we find no evidence they tried to ask about past wage history. There is no increase in the terms "wage," "wage history," "hourly rate" and so on in the messages of the treatment group.[20]

We also analyze the message text by counting how frequently words appear that are found in the Linguistic Inquiry and Word Count 2015 (LIWC) dictionary (Pennebaker et al., 2015). The LIWC is a word dictionary that classifies words into categories often studied in social, health, and personality psychology such as cognitive processes, perceptual process, and biological processes. Each of these categories also includes sub-categories. We focus on the categories of "interrogation," and "cognitive processes," and two "cognitive processes" sub-categories: "insight" and "causation." For example, the sub-category of "insight" contains words such as: "think" and "know," while the sub-category of "causation" contains words such as: "because" and "effect." The idea is that coordination messages would be less likely to have these kinds of terms than interviewing messages. To analyze this textual data, we use a message-level logit model weighted by the number of words in each message thread, with standard errors clustered at the job opening level (Papke and Wooldridge, 1996). If the employer sent more than one message to an applicant, we combine all messages sent by an employer into a single message. On average there are about 3.5 messages sent by an employer to an applicant, per message thread.

Column (1) of Table 6 reports a regression where the outcome of interest is equal to the proportion of words in a message which are associated with "interrogation" sent by the employer on a job opening. We find that the proportion of words in a message that are associated with "interrogation" is 1.09 times the proportion of words in a message that are associated with "interrogation" for control employers.

[20] Analysis not shown, but available upon request.

Table 6: Effect of hiding past wage history on the kinds of words used in messages from employers to applicants

	Dependent variable:			
	Interrogative term usage	Cognitive process term usage	Insight term usage	Causation term usage
	(1)	(2)	(3)	(4)
Wage history hidden	0.091^{*}	0.035^{*}	0.089^{**}	0.047
	(0.051)	(0.021)	(0.039)	(0.047)
Constant	-4.493^{***}	-2.252^{***}	-3.957^{***}	-4.104^{***}
	(0.041)	(0.016)	(0.031)	(0.036)
Observations	36,814	36,814	36,814	36,814

Notes: This table reports regressions where the outcomes are the fractions of words in a message that come from specialized list of thematically related words. The unit of observation is messages sent by the employer to applicants. On average there are about 3.5 messages sent by an employer to an applicant, per message thread. Estimates are from logit models weighted by the number of words in each message. The dependent variable in Column (1) is the percentage of total words associated with interrogation and includes words like: how, when, what. The dependent variable in Column (2) is the percentage of total words associated with cognitive processes and contains words like: cause, know, ought. The dependent variable in Column (3) is the percentage of total words associated with insight and contains words like: think, know. The dependent variable in Column (4) is the percentage of total words associated with causation and contains words like: because, effect. Heteroskedasticity-robust standard errors clustered at the job posting level are reported. The sample is restricted to hourly first job posts by an employer. Significance indicators: $p \leq 0.10 : *$, $p \leq 0.05 : **$, and $p \leq .01 : ***$.

Columns (2), (3) and (4) report regression results in which the outcome is the proportion of words belonging to different LIWC categories. The results indicate that treated employers use 1.04 times as many "cognitive process" words (from Column (2)), and 1.09 as many "insight" words (from Column (3)). The Column (4) regression shows "causation" words are used more frequently, but the effect is not conventionally significant.

A.4 Estimate of demand curve

In Column (1) of Table 7, we report a regression of a hiring indicator for each application on the log wage bid of the applicant interacted with the treatment assignment. This regression is run without a worker-specific effect. The positive coefficient on wage bid reflects the fact that higher wages are positively correlated with signals of worker productivity and more productive workers are more likely to be hired. The treatment indicator is positive and highly significant, reflecting the increase in probability of employers hiring when they cannot observe past wage history. The negative coefficient on the interaction term implies that this increase in hiring probability is focused on applicants with lower wage bids, reflecting the earlier documented "bargain hunting."

In Column (2), we include a worker fixed effect. After including a worker-specific fixed effect, which controls for the worker specific general quality, a higher wage bid decreases the probability that worker is hired, as we would expect.

In Column (3), we include both a worker-specific fixed effect and a job-posting-specific fixed effect. The identifying variation comes from workers applying at different wages to different job postings over a short period of time. This controls both for employer and job specific heterogeneity in overall hiring probability as well as worker-specific heterogeneity. We using the coefficient on the log(Wage bid) term as the change in probability of hiring due to change in wage bid. Thus, we treat $p'(w) \approx -0.011$. We will use this estimate for all workers when calculating equilibrium wage adjustments.

Table 7: Wage bid and probability of hiring

	Dependent variable:		
	Applicant hired		
	(1)	(2)	(3)
Log wage bid	0.003^{***}	-0.011^{***}	-0.011^{***}
	(0.001)	(0.002)	(0.002)
WageHistHid	0.006^{***}	0.003	
	(0.002)	(0.002)	
Log wage bid × WageHistHid	-0.002^{**}	-0.001	-0.001
	(0.001)	(0.001)	(0.001)
Constant	0.011^{***}		
	(0.001)		
Outcome mean	0.018	0.018	0.018
Worker FE	*N*	*Y*	*Y*
Job Opening FE	*N*	*N*	*Y*
Observations	188,833	188,833	188,833

Notes: This table reports regressions where the outcome is an indicator for whether the worker was hired. The unit of analysis is the individual application. Standard errors are clustered at the level of the individual application. Significance indicators: $p \leq 0.10 : *$, $p \leq 0.05 : **$, and $p \leq .01 : ***$.