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

Moshe A. Barach[*]
University of Minnesota

John J. Horton
MIT Sloan & NBER

Tuesday 22nd October, 2019; 19:35

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 5% lower past average wages and hired workers with 13% lower past average wages. Conditional upon bargaining, workers hired by treated employers struck better wage bargains for themselves.

JEL: J01, J30, M50, M51

---

1 Introduction

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 problem challenging from the firm's perspective is that a worker is 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). As such, 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). Knowing the applicants' past wage could also usefully inform the firm's bargaining strategy, affecting what initial offer to make and how to respond to the worker's proposals (Nash Jr., 1950; Binmore et al., 1986).

Perhaps as a consequence of the usefulness it has to employers, half of workers in the US report that their current employer learned their wage from their previous job (Hall and Krueger, 2012). Over 80% of workers in the US report that if their employer learned their past wage, they learned it before extending a job offer.[1] Despite the apparent usefulness of compensation history information to employers, there is a public policy concern: access to past wage information can potentially create a path dependence in wages, 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 compensation histories of their applicants, whereas control employers could. The compensation 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.[2]

---

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

[2]We use the terms "employer" and "employee" or "worker" for consistency with the

Our empirical focus is on how the absence of compensation history information 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, (4) whether employers made a hire, and (5) the match quality, if a hire was made.

Much of the empirical literature on information in hiring comes from audit studies. The information about particular candidates in manipulated, and the key outcome is typically the characteristics of workers that are called-back or hired (Goldin and Rouse, 2000; Bertrand and Mullainathan, 2004; Autor and Scarborough, 2008; Dobbie et al., 2019). Although the intent of some information-removing policies is to counteract discrimination, a potential down-side is that they may encourage a reliance on other signals, which in turn might harm the very workers the policies are designed to help (Agan and Starr, 2017; Doleac and Hansen, 2018; Craigie, 2017; Shoag and Veuger, 2016). Furthermore, removing information might potentially cause a reduction in hiring altogether if the process becomes too costly relative to the expected benefits.

In our experiment, 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. In short, treated employers responded to their information deficit primarily by acquiring more of their own information. 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, though our estimates are generally imprecise.

In addition to causing more extensive and intensive evaluation, the treatment also changed what *kind* of workers were evaluated and eventually hired. Although treated employers evaluated workers with only slightly lower past wages, workers hired by treated employers had about 13% lower past aver-

_____

economics literature, not as a commentary on the legal nature of the relationships created on the platform.

age 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). The treatment made low experience/low wage bid workers seem like better "deals" relative to high experience/high wage bid workers, as firms with less information infer the workers have productivities closer to the mean. This of course raises questions about what would happen in equilibrium—a consideration we will discuss at length.

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, 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. Treated employers might also "drop out" earlier, deciding not evaluate anyone given their lack of compensation history information. This concern was not borne out in the experiment, as treated employers were *more* likely to make a hire. Furthermore, there is no evidence that treated employers had worse contractual outcomes, though our estimates are imprecise.

We present all of our results using the full sample, but we do find evidence of heterogeneous effects. In particular, we find that all of our 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 compensation history information. This heterogeneity is important, as it suggests different kinds of employers might be differently impacted by policies restricting the use of compensation history information.

A natural question is whether employers would continue to acquire information as in the experiment when compensation history is available again. Creating a panel of employers whose activity spans the experimental period and the return to the pre-experiment *status quo*, we compare experimental behavior to post-experimental behavior. We find that when treated employers regained access to compensation history information, they revert to their previous degree of information acquisition. This implies that employers do not perceive the increased screening costs borne during the experiment as worth the cost savings. However, removing wage history information 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).

A limitation of experiment is that workers did not know that their past wage histories might be hidden from certain employers. In an equilibrium in which the lack of employer access to compensation history was common knowledge, workers could respond to employer behaviors in many ways, such as altering which jobs they apply to or how they present themselves to employers. In our setting, in which workers are the first to bid, relatively low wage workers might adjust their wage bids up, and for relatively high wage workers to adjust their bids down. This would reduce the experimental "bargain hunting" effects we found and instead turn it into a price effect. We explore this possibility with a structural model of hiring, finding that bid adjustments would be exceedingly modest. However, future work looking at entire markets would help alleviate this "partial equilibrium" concern, and indeed, some is already emerging—see McNichols and Hanson (2019).

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 directly explores the role of compensation history in hiring. 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 (as in audit studies). A feature of this paper that differentiates it from the larger literature on information in hiring

is that we document the importance of endogenous information acquisition as an employer response to an information deficit.[3] Endogenous information acquisition could be an important margin of adjustment in other contexts but would typically be overlooked, as it is often hard to measure. This finding highlights an advantage of our online setting.

The experiment is timely, in the sense that it approximates policies that are being implemented—or are under active consideration—in several conventional labor markets. For example, both New York City 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). Early evidence is consistent with our results—using a synthetic control approach, McNichols and Hanson (2019) finds substantial increase in the female to male wage ratio following salary history bans in both California and Massachusetts.

The rest of the paper is organized as follows: Section 2 describes the empirical context. Section 3 presents the experimental design. Section 4 presents the results. Section 5 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 com-

---

[3]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.

[4]"Philadelphia Is About to Ban Employers From Asking Potential Hires About Their Salary History," Fortune Magazine, January 20th, 2017. Accessed online on Tuesday 22nd October, 2019, at: `http://fortune.com/2017/01/20/philadelphia-wage-history-employee-salaries/`.

mon 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 (2019b) 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.[5] 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.[6]

Firms presumably tailor their hiring to attract the employees that generate

---

[5]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.

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

the most match specific surplus for the task at hand. Drawing on the assortative matching in labor markets literature (Rosen, 1982; Sattinger, 1993), we take firms' ex-ante vertical preferences for price and quality as a signal of the project's importance. Firms that indicate they are looking for "entry level" labor are revealing that the project is not essential and they are willing to accept some risk that the task is not completed. 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 worker "profiles." A profile page shows details about the worker's work history on the platform, skills, education, availability to take on more work, and other information that he or she wants to share.[7] 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

---

[7]See Horton (2019a) on the importance of worker capacity information on probability of match formation.

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 in the control group are hired at a wage lower than the proposed wage due to back-and-forth wage bargaining.

Once hired, hours-worked are recorded using platform-provided 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. 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.[8] The change

---

[8]As a robustness check, we run our analysis on private jobs which are composed only of

in the interface available to employers was not explained to treated employers; the interface simply changed.[9] 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 compensation 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.3. 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.

3.1 Employer's view of the applicant pool and measures of screening

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.	$6.15/hour	4.5	123	UK
Paul H.	$8.27/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 as-

---

applicants expressly invited by the employer with whom the employer previously worked. There are no treatment effects on private jobs.

[9]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.

sociated 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.

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, 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.[10]

3.2 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

---

[10]In Appendix A.5 we conduct further analyses which investigate the nature of the messages (i.e., simply coordinating logistics or asking probing questions.

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 row labeled "Applicants" reports statistics on the total number of applicants that applied to the job. "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.

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—only 44% of a sample of conventional jobs were filled within 30 days, and it is suspected that many of these openings are never filled.[11]

4 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 it is useful, we will also use

---

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

regression.

4.1 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 "*Measures 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

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

where y_j is some outcome of interest, WAGEHISTHID_j is the treatment indicator, and $\mathbf{X}_j$ is a collection of pre-randomization job opening and employer characteristics.[12] We plot the coefficients on WAGEHISTHID_j in Figure 1. For each point estimate, a 95% CI is shown. For each outcome, we plot the coeffi-

---

[12]We control for the category of the job opening, prior jobs billed by the employer, the amount of money the employer has previously spent 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.

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

	Control	Treatment	Difference	% Change
Measures of employer interest (number of applicants per job)				
	N = 2,948	N = 2,974		
Viewed	6.67 (0.17)	7.12 (0.18)	0.45 (0.24)*	6.76
Measures of employer evaluation/elicitation (number of applicants per job)				
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				
	N = 6,530	N = 6,725		
Bid amount	12.85 (0.33)	12.40 (0.29)	-0.46 (0.44)	-3.55
Profile wage rate	12.86 (0.30)	12.36 (0.26)	-0.50 (0.40)	-3.86
Avg 6-month wage	11.09 (0.29)	10.56 (0.25)	-0.53 (0.38)	-4.78
Min 6-month wage	8.93 (0.24)	8.53 (0.21)	-0.40 (0.32)	-4.51
Max 6-month wage	14.11 (0.40)	13.36 (0.31)	-0.76 (0.50)	-5.37
Previous hours worked	1107.33 (46.57)	1072.50 (36.68)	-34.84 (59.28)	-3.15
Prior billed jobs	27.62 (0.93)	26.73 (0.79)	-0.90 (1.22)	-3.24
Avg Feedback	4.70 (0.01)	4.71 (0.01)	0.00 (0.01)	0.08
Job opening outcomes				
Hire made?	0.40 (0.01)	0.43 (0.01)	0.03 (0.01)**	7.23
Hourly Wage Rate	12.29 (0.44)	10.93 (0.36)	-1.35 (0.56)**	-11.01
Feedback (1-10)	8.79 (0.12)	8.91 (0.10)	0.12 (0.15)	1.38
Contract rated a success?	0.59 (0.02)	0.61 (0.02)	0.02 (0.03)	3.65
Characteristics of hired applicants				
	N = 1,520	N = 1,744		
Bid amount	11.77 (0.35)	10.53 (0.47)	-1.25 (0.58)**	-10.59
Profile wage rate	12.05 (0.34)	11.03 (0.48)	-1.02 (0.59)*	-8.47
Avg 6-month wage	10.33 (0.38)	9.00 (0.43)	-1.33 (0.57)**	-12.87
Min 6-month wage	8.36 (0.34)	7.10 (0.38)	-1.26 (0.51)**	-15.08
Max 6-month wage	13.11 (0.49)	11.71 (0.55)	-1.39 (0.74)*	-10.64
Previous hours worked	1140.67 (57.99)	1252.57 (85.72)	111.90 (103.49)	9.81
Prior billed jobs	35.90 (1.75)	34.33 (1.50)	-1.58 (2.30)	-4.39
Avg Feedback	4.72 (0.01)	4.71 (0.02)	-0.01 (0.02)	-0.20
*Wage Bargaining	hire made*				
Any bargaining?	0.13 (0.01)	0.12 (0.01)	-0.01 (0.02)	-8.04
Wage-to-bid	bargaining	0.88 (0.03)	0.96 (0.03)	0.08 (0.04)*	9.22

Notes: This table reports means errors across experimental groups. Sample sizes are reported inline for most outcome categories. When the sample size differs within category of outcome they are reported below. The sample size for evaluating if there was a hire is 2,974 treatment jobs and 2,948 control jobs. The sample sizes for evaluating wage rate and if there is any bargaining is 907 treatment jobs and 842 control jobs. The sample sizes for evaluating feedback is 604 treatment jobs and 547 control jobs. The sample sizes for evaluating the change in wage-to-bid ratio is 105 treatment jobs and 105 control jobs. 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 $p \leq .01 : ***$.

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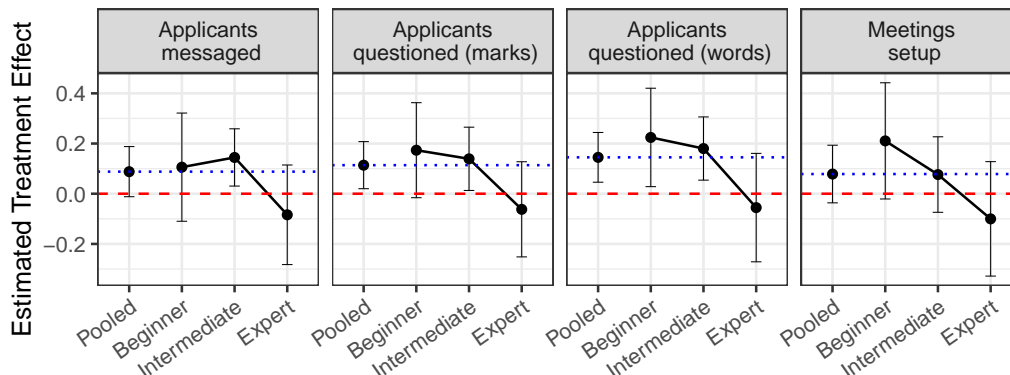
Note: This figure shows the relationship between measures of information acquisition and the treatment separated by requested expertise of the worker. The level of observation is the job opening. The model and controls are the same as used in the estimation of Equation 1, except that each employer vertical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

cient 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 9.2% 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. While economically meaningful, the effect on face-to-face meetings setup is borderline insignificant in the "Pooled" sample.

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]

---

[13]Another margin which employers could react would be by directly asking applicants about their past earnings. Additional analysis, which is available from the authors, shows that treatment and control employers do not differentially ask about "wage[s]", or "earn[ing—s]" or "rate[s]" in messages to applicants.

In fact, when restricting the sample to employers interested in hiring "Beginner" applicants, the effect on face-to-face meetings setup is significant at the 10% level.

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:

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

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 exponentiated to be interpretable as odds ratios. The left panel of the figure report estimates for the question marks, the center panel reports odds ratios for question words, and the right panel reports odds ratios of 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

---

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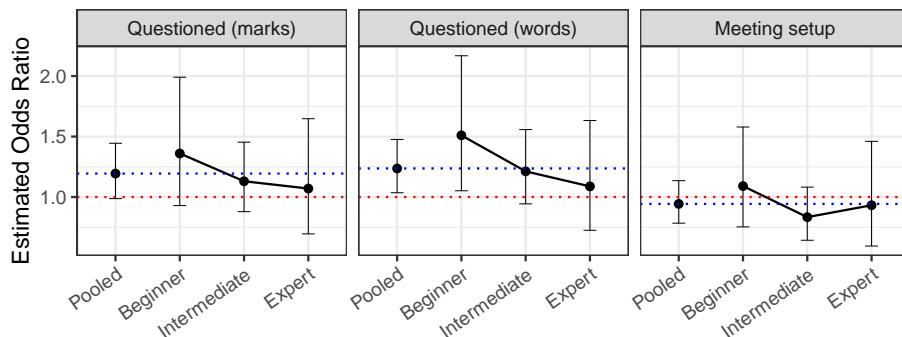
Note: This figure shows the relationship between intensity of information acquisition and the treatment, by the employer's reported vertical preference. The model and controls are the same as used in the estimation of Equation 2, except that each employer vertical preference group is fit separately.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

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 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.

4.2 Characteristics of called-back applicants

In the absence of compensation 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, though these differences were not, taken one by one, statistically significant at the 10% level. However, the effects are all directionally the same—called-back workers in the treatment bid about 4% less, had 4% lower profile rates, and 5% lower average past wages. All effects are consistent with treated employers evaluating applicants who have a greater potential to be a bargain.

To study how these effects differ by employer's vertical hiring preferences, we estimate an applicant-level selection model

$$
\begin{aligned}
\text{CALLEDBACK}_{i,j} = {} & \beta_0 + \beta_1 \log(\text{PROFILERATE}_{i,j}) \\
& + \beta_2 \text{WAGEHISTHID}_j + \beta_3 \left(\log(\text{PROFILERATE}_{i,j}) \times \text{WAGEHISTHID}_j \right) \\
& \hspace{4cm} + \epsilon_{i,j} \Big| \text{VIEWED}_{i,j} = 1, \quad (3)
\end{aligned}
$$

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, with separate panels for 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.

4.3   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

Figure 3: Employer probability of calling back a worker conditional on the applicant log profile rate, the treatment assignment, and the employer vertical preference
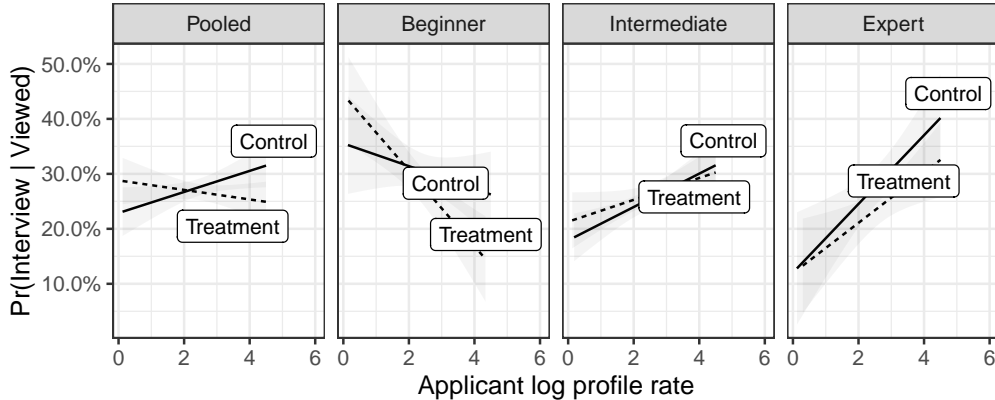
Note: This figure plots, for each employer requested expertise of the worker, predicted values from a linear model which regresses if the employer called-back an applicant, conditional on viewing the applicant, on the log of the applicant's profile wage rate by treatment status of the employer. Standard errors are clustered at the job opening level.

rate in the control group of 40%.

As we saw in Section 4.2, called-back applicants in the treatment have slightly lower past wages, albeit not significantly so. However, this bargain hunting is much more evident in 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 11% lower and average past wages were about 13% lower.[15] These differences are substantially larger than the differences in mean attributes of the called-back applicants.

Although the bargain hunting effect is sizeable, as we noted earlier, the treated group did have about a 7% higher fill rate, and so some of the decrease in the average past wages could be a composition effect due to these additional filled job openings. However, it is impossible for the change to be purely due to

---

[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.

the additional filled job openings as marginally filled jobs must have applicants with *negative* past wages to achieve a 13% reduction overall in past wages.[16]

We might suspect that, in the absence of compensation history information, employers would put more weight on other available signals, such as total past earnings, tenure on the platform, past feedback, and so on. We fit a model of employer hiring on the basis of observable characteristics and find no evidence that the weight put on various attributes differed by employer treatment status. However, many of the point estimates are quite imprecise. This analysis is in Appendix A.2.

4.4 Prevalence and outcome of wage bargaining

Although workers make the first wage offer, employers are free to counter-offer, which the worker can counter, and so on. We can measure the extent of this bargaining, and see to what extent bargaining explains the decline in hired worker wages. We consider two aspects of wage bargaining: (1) whether any bargaining 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 means for both measures in the treatment and control. There is no strong evidence that the treatment caused a change in the fraction of worker/employer pairs that negotiate, but there is some evidence that hired workers in the treated group strike better wage bargains, conditional upon bargaining.

Given that we know the treatment encouraged the hiring of lower wage workers, this could simply be a selection effect if lower-wage workers typically have better bargaining outcomes. In a regression framework, we can potentially control for these composition changes, as well as see whether bargaining effects are concentrated among relatively lower-wage workers.

First, we re-capitulate the means comparison results in Column (1) of Ta-

---

[16]The marginally filled jobs would need average past wages of $-8.68 per hour for the 13% reduction in past wages to be entirely driven by the increased fill rate, which is impossible.

ble 3, estimating

$$\text{AnyBargaining}_j = \beta_0 + \beta_1\text{WageHistHid}_j + \epsilon_j \big| \text{Hired}_j = 1. \qquad (4)$$

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 $p \leq .01 : ***$..

As expected, 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. Column (2) reports an

estimate of

$$\textsc{WageToBidRatio}_j = \beta_0 + \beta_1 \textsc{WageHistHid}_j + \epsilon_j \Big| \textsc{AnyBargaining}_j = 1. \tag{5}$$

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, we might suspect that 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 selection concern 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).

4.5 Contract outcomes

The treatment induced employers to hire workers with lower past average wages. If past wages simply reflected productivity, then we would expect more hours-worked if project sizes remained the same, but an unchanged wage bill. If these hired workers were "worse," this might have reduced the quality of the work the employer received. This in turn might create less employer surplus, though it could be offset by the lower wage bill. We do not directly measure employer surplus, but we do have some reasonble proxies.

In the panel labeled "*Job opening outcomes*" in Table 2, we report means for 1) the "private" feedback score, which was a rating on a scale from 1-10 and made only to the platform and (2) whether the employer rated the completed contract a success, which is a binary indicator. These two feedback

measures perhaps most closely map to surplus, as they are less subject to strategic misreporting (Filippas et al., 2019).

It is important to note that for these outcome measures, we are conditioning on filled jobs. This introduces both selection concern and it means the sample is smaller. We only have 5,922 job postings, and given that more than half of job postings are not filled we are left with 2,715 jobs eligible to provide feedback. Furthermore the baseline percentage of employers that actually leave feedback upon completion of a job is not1 100%. These caveats aside, both outcome measures in Table 2 are higher in the treatment, though the effects are far from conventionally significant.

We expand our collecetion of contractual outcomes in Table 4, estimating effects using linear regression and adding several pre-randomization controls.[17] The outcomes are: (1) the number of hours the applicant billed on the job, (2) the numerical public feedback the employer left for the worker (1 to 5 scale), (3) the numerical "private" feedback on the worker (reported only to the platform) (1 to 10 scale), also standardized (4) the dollar amount of any bonus left for the worker, and (5) if the worker was rehired by the employer after completing the job.

---

[17]We add controls for the number of prior contracts the employer billed, the number of times the employer has worked with this specific contractor and the category of the job.

Table 4: Effect of the treatment on employer's subjective measures of contract outcomes

| | *Dependent variable:* | | | | | |
| | log(Hours Billed) | Feedback Left | Public Feedback | Private Feedback | Bonus Given | Rehired |
	(1)	(2)	(3)	(4)	(5)	(6)
Wage history hidden	−0.010	−0.002	0.001	0.002	0.935	−0.003
	(0.021)	(0.077)	(0.056)	(0.115)	(4.623)	(0.015)
Constant	2.603***	0.746**	4.691***	9.935***	3.550	0.110
	(0.105)	(0.377)	(0.160)	(0.095)	(3.675)	(0.076)
Employer-level covariates	Yes	Yes	Yes	Yes	Yes	Yes
Assignment-level covariates	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,851	1,851	1,002	1,095	1,672	1,848

Notes: This table reports regressions where the outcome is a measure of the employer's subjective evaluation of the contract. The sample is restricted to fill job openings where feedback was left. In Column (1) the outcome is an indicator if the employer left any public feedback. In Column (2) the outcome is the log of the number of hours billed on the job. In Column (3) the outcome is the z-score of public feedback score left by the employer on the employee's performance. In Column (4) the outcome is the z-score of the private feedback score left by the employer on the employee's preformace. In Column (5) the outcome is an indicator if the employer left the worker a bonus. In Column (6) the outcome is an indicator if the employer rehired the worker for a later job. Covariates included are category indicators, the number of prior jobs the employer filled, and the number of prior jobs the employer and worker completed together, and total job billings. Significance indicators: $p \leq 0.10 : *$, $p \leq 0.05 : **$, and $p \leq .01 : * * *$.

Each of these measures is not without some complications in interpretation. Changes in hours-worked could reflect hiring a more productive worker (which would tend to lower hours-worked) but if these same workers are hired a wage lower than their marginal product, the employer might ask for more work to be completed. As it is, in Column (1) we see there appears to be no evidence of a change in hours-worked in the treatment, though the estimate is imprecise.

Before presenting the feedback measures, we first report a regression whether the outcome is whether any feedback was left, in Column (2). Importantly, the treatment does not seem to have any effect whatsoever on the feedback response rate. In Columns (3) and (4), the outcome is the public and private feedback scores. There is no evidence of a difference by treatment assignment—both point estimates are very close to zero, with the private feedback measure now being closer to zero compared to what we observed in Table 2. In Column (5) and (6), the outcomes are bonuses and re-hires. As with our other results, there is no evidence of a difference by treatment assignment. The totality of the evidence is that matches were not detectably better or worse in the treatment group.

4.6   Is limiting information efficient from the employer's perspective?

If treated employers enjoy lower wages and no worse contractual outcomes, might they change how they screen going forward, post-experiment? The answer presumably depends on the cost of the added screening. We answer this question using a difference-in-differences design by combining our experimental data with data from a post period in which the platform reverted to the pre-experiment *status quo*.

We have two time periods, $t = 0$, corresponding to the period when the experiment was active, and $t = 1$, which is after the experiment ended. We can

then estimate

$$y_{jt} = \beta_0 + \beta_1 \text{WAGEHISTHID}_j + \beta_2 t + \beta_3 \left(\text{WAGEHISTHID}_j \times t \right) + \mathbf{X}_{jt}\gamma + \epsilon_j, \tag{6}$$

where y_{jt} is one of the measures of intensive search used previously on job j, in time period t, WAGEHISTHID_j is the treatment assignment the employer posting job opening j, and $\mathbf{X}_j$ are pre-randomization job characteristics.[18] Our sample consists of employers that also posted a job with one month after the experiment ended. This sample is necessarily smaller, and to the extent the treatment affected an employer's probability of posting a subsequent job opening in the post-period, the sample could be selected, though we have no evidence that this was the case.

The coefficient on WAGEHISTHID_j should be similar to the experimental estimates, subject to the caveat that the sample is only those employers posting jobs in both periods. If the change in employer screening behavior persisted post-experiment, then the coefficient on $\left(\text{WAGEHISTHID}_j \times t \right)$ should be zero. If instead the employers revert, the point estimate on the intrreaction term should have the same magnitude but opposite sign as the coefficient on WAGEHISTHID_j.

Table 5 reports estimates of Equation 6 using a quasi-maximum likelihood Poisson regressions with heteroskedasticity-robust standard errors. In general, the point estimate on the treatment indicator is similar to the experimental estimates, though less precise (as expected). In Columns (1) and (2), we see some evidence that treated employers reverted to evauating smaller pools in the post period, though the estimates are imprecise and smaller in magnitude than the treatment effects.

---

[18]Controls include the category of the job opening, employer's prior jobs, employer's prior spending, number of applications to job, number of recommended applications to job, skill requested indicator, and the average applicant's bid.

Table 5: Effect of treatment on information elicitation by experimental period

	Dependent variable:				
	Applicants viewed, NUMVIEWED	Applicants messaged, NUMMESSAGED	Applicants questioned (marks), NUMQUESTIONEDMARKS	Applicants questioned (words), NUMQUESTIONEDWORDS	Meetings setup, NUMFACETOFACE
	(1)	(2)	(3)	(4)	(5)
WAGEHISTHID	0.054	0.093	0.149**	0.153**	0.072
	(0.045)	(0.063)	(0.061)	(0.066)	(0.084)
Period, t	0.018	0.048	0.092	0.093	0.036
	(0.048)	(0.068)	(0.069)	(0.074)	(0.079)
t × WAGEHISTHID	−0.017	−0.032	−0.220**	−0.246**	0.022
	(0.066)	(0.101)	(0.096)	(0.105)	(0.124)
Constant	1.675***	0.536***	0.034	0.165	−0.522*
	(0.133)	(0.186)	(0.206)	(0.224)	(0.274)
Opening-level covariates	Yes	Yes	Yes	Yes	Yes
Employer-level covariates	Yes	Yes	Yes	Yes	Yes
Observations	5,975	5,975	5,975	5,975	5,975

Notes: This table reports the results of Poisson regressions where the outcomes are various measures of employer information elicitation. The level of observation is the job opening. The sample is limited to employers who posted a job within 1 month of the experiment ending. Estimates are from quasi-maximum likelihood Poisson models. Heteroskedasticity-robust standard errors are reported.The outcome in Column (1) is a count of the number of applicants viewed. The outcome in Column (2) is a count of the number of applicants messaged. The outcome in Column (3) is a count of the number of applications that exchanged messages including a question mark with the employer. The outcome in Column (4) is a count of the number of applications that included at least one of the following question words: who, what, where, when, why, or how. The outcome in Column (5) is a count of the number of applications that included the word "Skype" in a message with the employer. All models include as controls indicators for the category of work, the number of prior jobs billed by the employer, the employers prior spending on the platform, the number of applications to the job opening, the number of recommended applications to the job opening, the average applicant's bid, and an indicator if the employer requested a specific skill. These other regression coefficients are not shown. Significance indicators: $p \leq 0.10 : *$, $p \leq 0.05 : **$, and $p \leq .01 : ***$.

The coefficients on the interaction, WAGEHISTHID $\times t$, in Column (3) and Column (4) are negative and significant, indicating that treatment employers revert in the post-experiment period. The effect sizes are large enough that the experimental effect is completely undone, though the effects are imprecise enough to make strong claims unwarranted. It appears that the upfront costs of more screening do not seem to be worth the potentially better match and lower wage. In short, in the absence of the experimental intervention, we cannot rule out that employers return to their *status quo* screening strategies.

4.7   Applicant outcomes

Although our experiment was conducted at the level of employer, we can also analyze the experiment at the application level. We might expect, given the shift towards hiring workers with lower past wages, that the applicants with comparatively higher past wages are disadvantaged by the treatment. However, this may not necessarily be the case given that removing compensation history information also increased hiring overall.[19]

We compare the monetary benefit to an applicant stemming from that applicants submitting a bids to both treatment and control job postings. The outcome is TOTALVALUE$_{i,j}$, the total earnings that originated from application from worker i to opening j. The outcome is \$0.00 if the application does not result in employment and a positive value if the applicant is hired. We take advantage of the fact that applicants could apply to both jobs posted by treated employers as well as control employers, allowing for a within-worker analysis. We estimate the following model:

$$\log(1 + \text{TOTALVALUE}_{i,j}) = \beta_0 \text{WAGEHISTHID}_j + \mathbf{X}_j \gamma + \eta_i + \epsilon_{i,j}, \quad (7)$$

where η_i is an applicant-fixed effect and $\mathbf{X}_j$ is a collection of pre-randomization job-opening and employer characteristics.[20] Figure 4 plots the coefficients on

---

[19]In this section we examine which individuals benefit differentially from bids submitted in treatment and control jobs. In Appendix A.6 we investigate longer-run effects on applicants, although our experimental design makes detecting these differences difficult.

[20]Covariates include: job category indicators, an indicator if the applicant is from the

Figure 4: Effects of the treatment on applicant's expected value of bidding
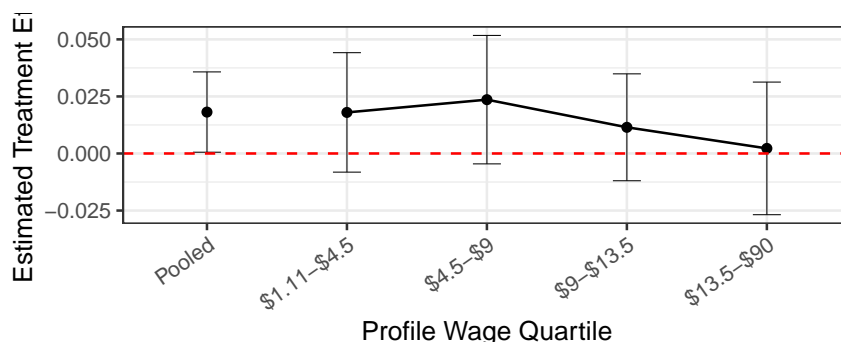
Note: This figure shows the relationship between an applicants expected value from submitting a bid, and the treatment by the applicant's profile wage. The level of observation is the bid. The first quartile of profile wages range from \$1.11-\$4.5, the second quartile ranges from \$4.5-\$9, the third quartile ranges from \$9-\$13.5, and the fourth quartile ranges from \$13.5-\$90. Models all include an applicant level fixed effect. Observations are weighted by 1/(number of total applications sent by the worker). Heteroskedasticity-robust standard errors clustered by job opening are used for 95% confidence intervals.

WAGEHISTHID$_j$ from each regression by quartile of the profile wage rate. For each point estimate a 95% CI is shown. The full sample estimate are labeled "Pooled."

The regression coefficients shown in Figure 4 for the "Pooled" sample indicate that on average, applicants' expected value of each bid submitted increases by about 1.8%. The quartile regression results indicate that while applicants whose profile wage rate is between \$4.50 and \$9.00 benefit the most from hiding past compensation history. There is no strong evidence that workers in the other quartiles were harmed (or helped), though these estimates are all fairly imprecise.

5 Discussion and conclusion

This paper demonstrates how employers react when compensation history is removed. They respond primarily by acquiring more information, expanding their

---

same country as the employer, the applicant's tenure, the algorithm score assigned to the applicant, and a indicator for the order the application was submitted to the job.)

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. There is no evidence they ultimatley form worse matches. The strongest finding of the paper is a marked shift towards evaluating and hiring lower wage workers.

Endogenous information acquisition by firms has been almost entirely overlooked in extant empirical work. 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)[21]. 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 is some literature that touches on it. 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).

Our findings suggest that policies that limit employer access to compensation 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 "entry level" vertical preferences were the employers who responded by more information acquisition and a shift in hiring.

---

[21]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.

References

Agan, Amanda and Sonja Starr, "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment," *The Quarterly Journal of Economics*, 08 2017, *133* (1), 191–235.

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.

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.

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.

Craigie, Terry-Ann, "Ban the box, convictions, and public sector employment," *Convictions, and Public Sector Employment (January 27, 2017)*, 2017.

Dobbie, Will, Paul Goldsmith-Pinkham, Neale Mahoney, and Jae Song, "Bad Credit, No Problem? Credit and Labor Market Consequences of Bad Credit Reports," Technical Report, Becker Friedman Institute for Research in Economics Working Paper 2019.

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," *Journal of Labor Economics*, 2018.

Filippas, Apostolos, John J Horton, and Joseph M Golden, "Reputation Inflation," *Working Paper*, 2019.

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.

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," *Management Science (forthcoming)*, 2019.

_ , "Price floors and employer preferences: Evidence from a minimum wage experiment," *Working Paper*, 2019.

_ **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.

McNichols, Drew and Benjamin Hanson, "Information and the Persistence of the Gender Gap: Early Evidence from California's Salary History Ban," *Working Paper*, 2019.

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, James W., Martha E. Francis, and Roger J. Booth, "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 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.

35

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 compensation 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.[22] For this survey, we received responses from 391 subjects with demographic weights. Of those, 115 reported they were asked about their compensation history, or 29.4%. Among those

---

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

Figure 5: Survey evidence on when compensation history is asked about and the order of wage bargaining
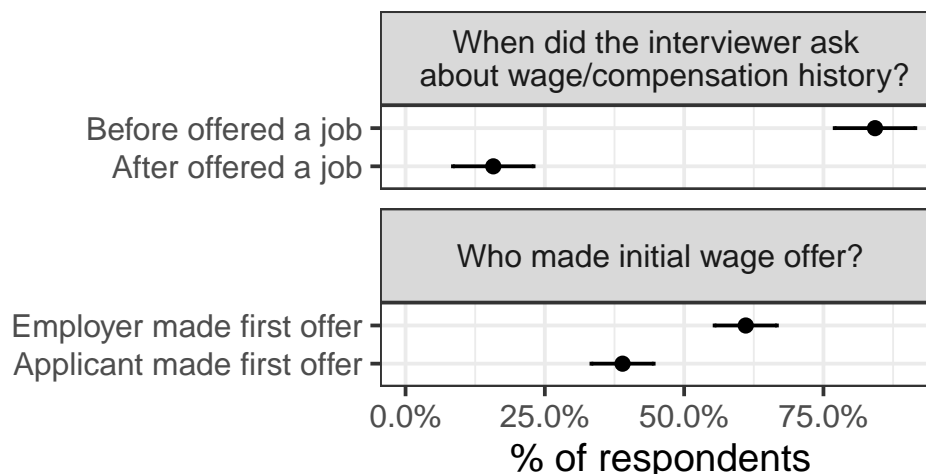
Note: This figure reports population-weighted means from two surveys conducted on Google Surveys. In the top panel, the sample is restricted to respondents reporting that in the last job they interviewed for the employer asked for their compensation history. Respondents answered whether the employer 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.

asked, 82.6% report the firm asked about wage/compensation before extending a job offer. The population-weighted fraction is shown in Figure 5, in the top panel, with weighted standard errors, which is very close to the unweighted fraction. This first survey shows that asking about compensation 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.[23] 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 5, 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   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 6 reports the regression coefficients on each signal from regressions run separately for treatment and control employers, and by employer vertical

---

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

preference, from a regression of the form:

$$\text{CALLEDBACK}_{i,j} = \sum_{x=0}^{5} \beta_x \text{SIGNAL}_x + \gamma_j + \epsilon_{i,j} \Big| \text{VIEWED}_{i,j} = 1, \qquad (8)$$

where x indexes x-score measures of five readily available signals of productivity of each applicant i to job j including: Wage bid, Profile wage rate, Prior feedback score, Prior number of jobs completed, Prior earnings, and Platform tenure in days. 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 6 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 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.

Figure 6: Applicant characteristics on the probability call-back, by vertical preference tier
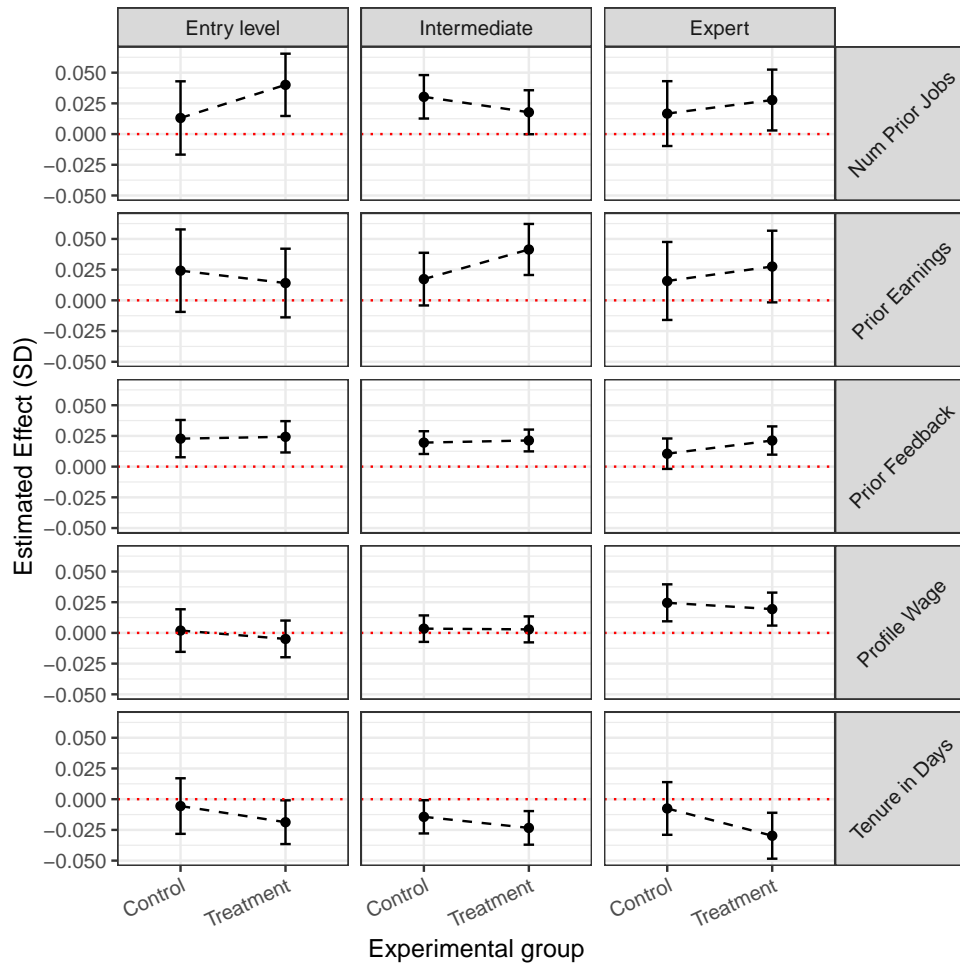
Note: This figure shows the coefficients on the non-wage bid parameters in Equation 8. The sample consists of viewed applications. Independent variables are transformed into z-scores normalized within job. Estimates represent change in standard deviations in probability of interviewing an applicant with 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.

A.3 Randomization

Table 6 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 6 supports this contention.

A.4 Expectations of available information

Employers' expectations concerning available information could influence our interpretation of observed treatment effects. It is possible that the observed increase in number of applicants viewed could be driven by employers experiencing surprise or confusion over the lack of usually available information on applicants' compensation histories. Especially employers who have had previous experience using the platform and, therefore, might react to not being able to view the past compensation history. It is possible that after viewing one application and noticing this information is missing, the employer might view another application to investigate Furthermore, an increase in the number of applicants viewed could also influence downstream measures of information acquisition and potentially even probability of hiring. Thus, we must investigate if expectations of available information causes a treatment effect which is unrelated to exogenous efforts by employers to obtain more information about a wider set of applicants.

To investigate this alternative explanation for our results we compare employers who have never previously posted a job on the platform and thus have no expectations of which information may or may not be provided in this mar-

Table 6: 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 a 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 $p \leq .01 : ***$.

Figure 7: Effects of the treatment by employer information expectations
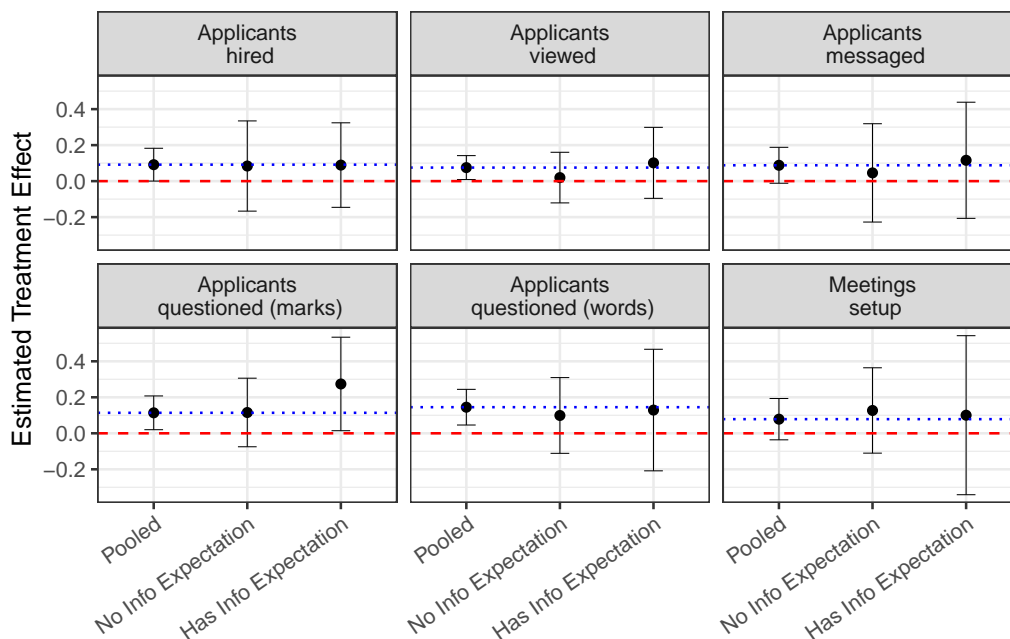
Note: This figure shows the relationship between measures of information acquisition and the treatment separated by the employer's expectation of available information. The level of observation is the job opening. The sample is limited to employers with either zero or one previous on-platform experience. Heteroskedasticity-robust standard errors are used for 95% confidence intervals.

ket to employers who have only posted one previous job. We limit the analysis to those with only one previous job as much more experienced applicants use different screening practices, which limits comparability. We re-estimate the same Poisson count regression as in Model 1 separately for employers with zero previous job postings and for employers with exactly one previous job posting.[24] We plot the coefficients of WAGEHISTHID$_j$ in Figure 7. For each point estimate, a 95% CI is shown.

Reassuringly, the regression coefficients shown in Figure 1 for the "Pooled" sample, which consists of only employers with either zero or one previous plat-

---

[24]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.

form job postings, give marginal effects similar to those for the complete sample used in our main results. The results of Figure 1 also allow us to conclude that the effect of removing compensation history information is not statistically different for employers who have expectations about platform-provided information ("Has Info Expectation") and employers with no expectation over what information is provided by the platform ("No Info Expectation").

We also investigate expectations affecting treatment effects by looking for differences in the treatment effect over the duration of experimental period for employers who posted multiple jobs over the experimental period. By analyzing all of an employer's job postings after being assigned to treatment or control, we can look for decay in the experimental results. If the treatment effect is driven primarily by a surprise reaction to the lack of available information, once employers adapt to not expecting to observe applicant's past compensation history, the treatment effect should decrease.

The median number of job postings by an employer during the experimental treatment is 1, and the 75th percentile is only 2 job postings. For this reason, we group job postings into four groups: Employers' first job posted during the experiment, employers' 2nd job posted during the experiment, employers' third through fifth jobs posted during the experiment, and job postings for which the employer has posted greater than 5 previous postings during the experimental period.

We estimate the same Poisson count regression as in Model 1 however, we now add an employer random effect to help control for individual level heterogeneity. We plot the coefficients of $\textsc{WageHistHid}_j$ in Figure 8. For each point estimate, a 95% CI is shown. The treatment effects are stable over the entire experimental period. There is no statistical difference in estimated treatment effect based on how many prior jobs the employer has posted over the experimental period for any of our information acquisition measures. Together, these investigations allow us to conclude that the treatment effect is not driven by expectations over the available information in the market, but by employers seeking to obtain additional available information in the market.

Figure 8: Effects of the treatment by number of previous experiment job postings
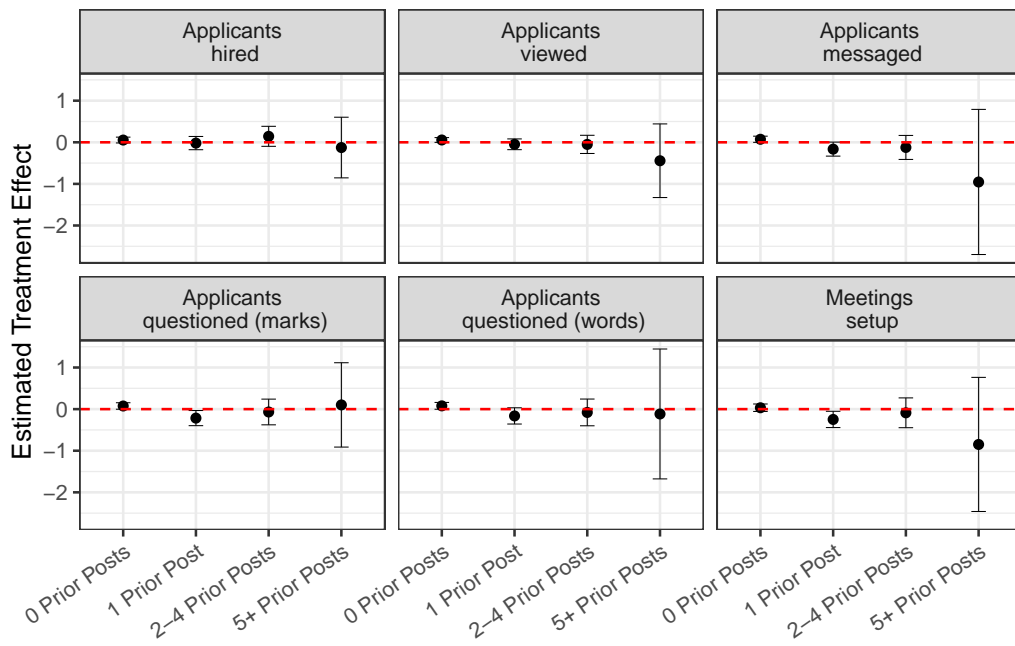
Note: This figure shows the relationship between measures of information acquisition and the treatment by the number of previous job postings by an employer during the experiment. The level of observation is the job opening. Models all include an employer-level random effect term. Heteroskedasticity-robust standard errors clustered by employer are used for 95% confidence intervals.

A.5 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 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 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 logic behind this analysis 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 7 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.

Columns (2), (3) and (4) report regression results in which the outcome is

Table 7: 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 $p \leq .01 : ***$.

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.
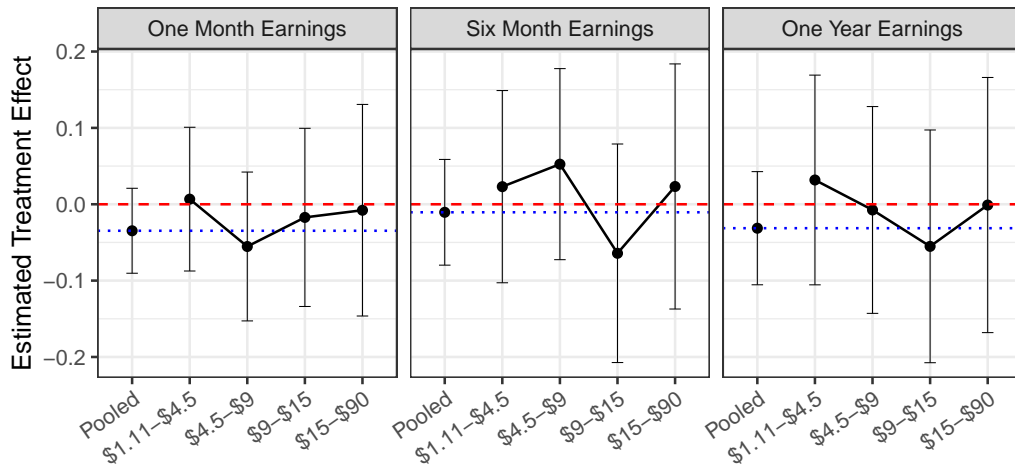
A.6 Longer-term Effects

Due to the short duration of the experiment, our ability to identify longer-term effects of removing past compensation history on applicants and employers is limited. However, it is possible that the short-term advantage preference for lower wage applicants caused by the removal of past wage information may lead to greater future employment opportunities, and greater future wages for these advantaged applicants. To evaluate these longer-term effects, we need to look at outcomes which occurred long after the experiment ended. This creates a problem of control-treatment interference. To help eliminate treatment vs. control overlap, we remove all applicants who experienced both treatment and control jobs. We compare applicants who experienced only treatment jobs to applicants who experienced only control jobs. It is important to note that the sample in this analysis is composed of applicants who only applied to either treatment or control jobs and thus will contain applicants who applied to fewer jobs on average. The more jobs an applicant applies to the greater the probability they will experience both a treatment and control job.

Figure 9 plots the the coefficients from log-linear regressions where the output of interest is the log of future wages. As such, the coefficients can be interpreted as percent difference in treatment applicants' future wage from the future wages of control applicants.

Across all of the models in Figure 9, the large standard errors make it impossible for us to conclude whether the short-run experiment had long-run effects on applicants' earnings. Furthermore, there is no detectable time trend either increasing or decreasing among any of the profile wage groups. Thus, we can not draw any conclusions concerning longer-term wage effects, which is not surprising given the constrains of our experimental design.

Figure 9: Effects of the treatment on future applicant earnings

A.7 Equilibrium considerations

In the experiment, workers did not know that employers would lack access to their compensation history. With an actual market-wide policy change, workers would know that their compensation history is hidden, and could respond. Workers could respond in many ways, such as altering which 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:

$$\Pr(\text{HIRED}_i = 1) = \frac{e^{x_i'\beta}}{\sum_a e^{x_a'\beta}}, \tag{9}$$

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.

A.7.1 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 9 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.[25] We estimate the β coefficients separately for treatment job openings and for the control job openings. Table 8 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.

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 that worker's wage bid using Equation 9. A worker's predicted change in hiring probability for each job opening they apply to is thus:

$$\Delta p_i = \Pr(\text{HIRED}|w_i, \hat{\beta}_{WHH=1}) - \Pr(\text{HIRED}|w_i, \hat{\beta}_{WHH=0}). \tag{10}$$

---

[25]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.

Table 8: Estimated Coeffcient 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 co-effcient vector for treatment and control job postings. Estimates are generated from Equation 9 using maximum likelihood estimation.

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 compensation history is hidden and $\hat{\beta}_{\text{WHH}=0}$ are the estimated coefficients for the control. In the top panel of Figure 10, 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.

A.7.2 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

$$m(\hat{y}_i) = e^{\gamma_0 + \gamma_1 \log \hat{y}_i}, \tag{11}$$

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 $\gamma_1 < 0$.

Figure 10: Estimated effects on individual worker's probability of hire
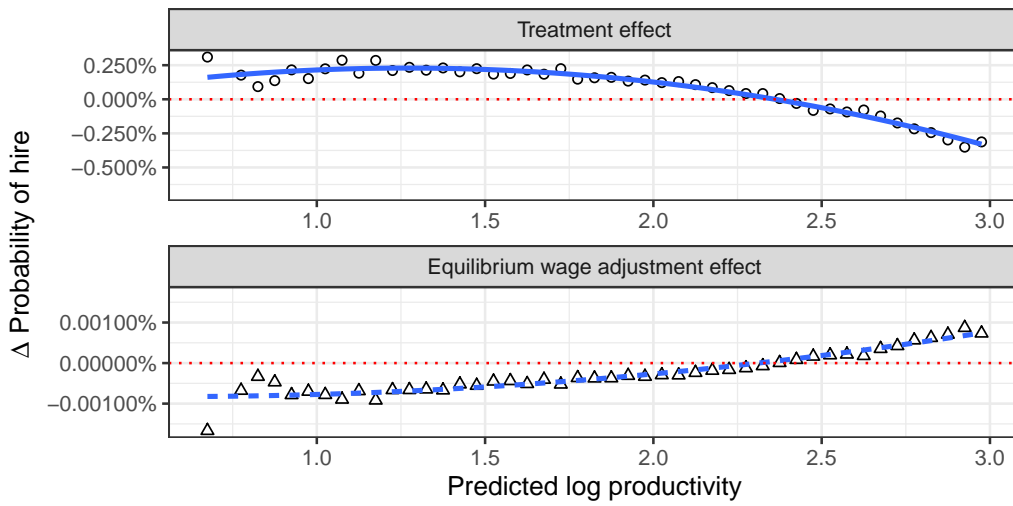
Note: This figure shows the relationship between worker productivity and the change in probability of being hired due to (1) the employer's ability or lack thereof to observe past wages, in the top panel, and (2) the worker's equilibrium wage adjustment, in the bottom panel. Plotted points are averages 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.

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, 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' inability 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:

$$dw^* = -\frac{dp(w)}{p'(w)}.$$

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

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

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

worker:

$$\mathbb{E}\left[\frac{\mathrm{d}p}{p - c_i p'(w)} - \frac{\mathrm{d}w}{w}\right] = 0. \tag{12}$$

To use Equation 12, 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.7.3, 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 $c_i = 0.9 w_i$, but we try a range of values.[26]

We estimate the the γ vector from Equation 11 by minimizing the square error in the sample analog of Equation 12. 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 12. 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 10 shows how little of an effect wage bidding adjustments have on hire probability. As expected, workers' equilibrium wage adjustment decreases the probability of being hired for relatively low productivity 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.

A.7.3 Estimate of demand curve

In Column (1) of Table 9, 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 compensation 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-

Table 9: 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 $p \leq .01 : ***$.

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.