

MIT Open Access Articles

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

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Barach, Moshe A. and John J. Horton. "How Do Employers Use Compensation History? Evidence from a Field Experiment." *Journal of Labor Economics*, 39, 1 (October 2020) © 2020 The Author(s)

As Published: 10.1086/709277

Publisher: University of Chicago Press

Persistent URL: <https://hdl.handle.net/1721.1/129809>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



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

Moshe A. Barach, *University of Minnesota Carlson School of Management*

John J. Horton, *Massachusetts Institute of Technology Sloan School of Management and National Bureau of Economic Research*

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 on bargaining, workers hired by treated employers struck better wage bargains for themselves.

I. 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 and Schaefer 2011). What

For helpful comments and advice, thanks to Amanda Agan, Ned Augenblick, Alan Benson, Amanda Pallais, Joe Golden, Liz Lyons, John Morgan, Chris Rider, Steve Tadelis, Stan Veuger, and Noam Yuchtman as well as seminar participants at the University of California, Berkeley, Haas School of Business, Georgetown University McDonough School of Business, and the University of Minnesota Carlson School of Management. Contact the corresponding author, John J. Horton, at jjhorton@mit.edu. Information concerning access to the data used in this paper is available as supplemental material online.

[*Journal of Labor Economics*, 2021, vol. 39, no. 1]
© 2020 by The University of Chicago. All rights reserved. 0734-306X/2021/3901-0006\$10.00
Submitted October 24, 2017; Accepted April 2, 2020; Electronically published October 30, 2020

makes this problem challenging from the firm's perspective is that a worker is an "experience good," so the employer has to make an inference about productivity, relying on whatever signals are 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; Altonji and Pierret 2001; Lange 2007; Oyer and Schaefer 2011; Kahn and Lange 2014). Knowing an applicant's 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 1950; Binmore, Rubinstein, and Wolinsky 1986).

Perhaps as a consequence of the usefulness it has to employers, half of workers in the United States report that their current employer learned their wage from their previous job (Hall and Krueger 2012). More than 80% of workers in the United States report that if their employer learned their past wage, they learned it before extending a job offer. 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.¹

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 is manipulated, and the key outcome is typically the characteristics of workers who are called back or hired (Goldin and Rouse 2000; Bertrand and Mullainathan 2004; Autor and Scarborough 2008; Dobbie et al. 2016). Although the intent of some information-removing policies is to counteract discrimination, a potential downside is that they may encourage a reliance on other signals,

¹ We use the terms "employer" and "employee" or "worker" for consistency with the economics literature, not as a commentary on the legal nature of the relationships created on the platform.

which in turn might harm the very workers the policies are designed to help (Shoag and Veuger 2016; Agan and Starr 2017; Doleac and Hansen 2018; Craigie 2020). 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, although 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 average wages. The observed preference for lower-wage workers—which we refer to as “bargain hunting”—is consistent with the framing of the hiring problem found in the personnel and labor economics literatures (Oyer and Schaefer 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 that 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 on. We find no evidence that the treatment affected the probability that bargaining occurred, but we find evidence that when bargaining did occur, workers hired by treated employers struck a more favorable wage bargain—they were offered and accepted wages that were 9% more of their initial bid compared with 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 to 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, although 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 that 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 preexperimental status quo, we compare experimental behavior to postexperimental 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 being 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 our 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 relatively high-wage workers might 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 Hansen and McNichols 2020).

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 (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.² Endogenous

² 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.

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.³ To the extent that 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 bargaining 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, Hansen and McNichols (2020) find a substantial increase in the female-to-male wage ratio following a salary history ban in California.

The rest of the paper is organized as follows. Section II describes the empirical context. Section III presents the experimental design. Section IV presents the results. Section V concludes.

II. Empirical Context

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

There has been some research that uses online labor markets as an empirical context. Pallais (2014) conducted a field experiment to demonstrate how much value employers place on past on-platform work experience. Gilchrist, Luca, and Malhotra (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. Barach, Golden, and Horton (2020) demonstrate how platform incentives influence the information employers use in making selections, and Barach et al. (2019) investigate how employers use machine learning algorithms in hiring. Agrawal, Lacetera, and Lyons (2016) investigate which factors matter to employers in making selections from an applicant pool and present some evidence of statistical discrimination; their paper also supports the view of

³ “Philadelphia Is About to Ban Employers from Asking Potential Hires about Their Salary History,” *Fortune*, January 20, 2017 (<http://fortune.com/2017/01/20/philadelphia-wage-history-employee-salaries/>).

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.

A. 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.⁴ 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.⁵

Firms presumably tailor their hiring to attract the employees who generate the most match-specific surplus for the task at hand. Drawing on the assortative matching in the 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 that 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.⁶ A worker also lists his or her “profile rate,” which is an hourly wage. Although it is

⁴ 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.

⁵ See Horton and Johari (2015) on the effects of this feature on applicant sorting.

⁶ See Horton (2019a) on the importance of worker capacity information on the probability of match formation.

self-reported, it is usually close to the typical wage that the worker earns, and employers consider it when deciding who to invite to their openings. The profile rate is a useful measure for our purposes, as it is recorded even if the worker has no work history on the platform. Furthermore, unlike average past wages, the profile rate is not dragged down by wages from jobs completed far in the past.

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

Employers can screen 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 because of 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.

III. Experimental Design

The experiment was conducted in late 2014 by the platform. All employers who 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 ($n = 2,948$). If an assigned employer posted an additional job opening, this job opening also received the treatment assignment of the initial job opening. However, we use only 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.⁷ The

⁷ As a robustness check, we run our analysis on private jobs, which are composed only of applicants expressly invited by the employer with whom the employer previously worked. There are no treatment effects on private jobs.

change in the interface available to employers was not explained to treated employers; the interface simply changed.⁸ 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 prerandomization attributes with respect to job opening characteristics, employer characteristics, and the composition of the applicant pool are shown in section A.3 of the appendix (available online). The experimental groups are well balanced, which is unsurprising given that the software used to allocate employers to treatment cells has been used many times and has proven reliable.

A. 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 associated with each past job held by the worker, but in the treatment group they could not. For example, a work history item for an applicant would be presented to a treated employer as follows:

Job title:	Lead data scientist
Contract type:	Hourly
Total:	\$451.34
Time:	December 2014 to present

However, a control employer viewing the same applicant would see the following:

⁸ 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.

Job title:	Lead data scientist
Contract type:	Hourly
Hourly wage:	\$17.00
Total:	\$451.34
Time:	December 2014 to present

Note that the control employer could see that the worker worked at \$17.00 a 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 six 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 searched these message threads for a number of measures of employer evaluation and constructed 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, let us say that the employer’s message to Julia was the following:

Hi Julia—Nice application, looking forward to working together!

This message does 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:

Table 1
Per-Opening Summary Statistics for the Control Group ($n = 2,948$)

Statistic	Mean	SD	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 who 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
Organic applicants “questioned”	1.121	2.050	0	0	36
Organic applicants interviewed face-to-face	.890	1.890	0	0	23
Applicants hired	.580	1.039	0	0	26

NOTE.—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 who applied to the job. “Invited” workers are those who the employer sought out and asked to apply for the opening. “Organic” applicants are workers who applied without being invited. An application is “viewed” if the employer clicked on a worker’s application to learn more about the applicant.

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.⁹

B. 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 seven of the applications submitted to the job by organic applicants and only message (i.e., call back) about two of these applicants.

Turning to the interviewing phase, employers specifically ask at least one question to about 62% of the applicants they message. About half of applicants who are messaged are asked to conduct a face-to-face interview, at least as measured by appearance of the Skype keyword. On average, this hiring process leads to about 40% of job openings posted being filled within 6 months of being posted (on average 0.58 applicants are hired, as some jobs hire more than one applicant). This is not dissimilar to the traditional labor market—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.¹⁰

⁹ In sec. A.5 of the appendix we conduct further analyses that investigate the nature of the messages (i.e., simply coordinating logistics or asking probing questions).

¹⁰ Report by Centre for Economics and Business Research (<http://press.indeed.com/wp-content/uploads/2015/01/Time-to-fill-jobs-in-the-US.pdf>).

IV. Results

We will present experimental results chronologically with respect to the hiring process, going from initial screening to posthire contractual outcomes. As we have a true experiment, we will always present results as simple means comparisons at the job opening level, although when it is useful, we will also use regression.

A. 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 panel 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 seven applications per opening, or about 7% more applicants than in the control.

In our experimental context, the equivalent of an interview callback 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, although 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 the presence of a 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, so we estimate a Poisson count regression

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

where y_j is some outcome of interest, WAGEHISTHID_j is the treatment indicator, and \mathbf{X}_j is a collection of prerandomization job opening and employer characteristics.¹¹ We plot the coefficients on WAGEHISTHID_j in figure 1. For each point estimate, a 95% confidence interval is shown. For each outcome, we plot the coefficient using the full sample (labeled “Pooled”) and for each of the three employer vertical preference levels (labeled “Beginner,” “Intermediate,” and “Expert”).

Reassuringly, the regression coefficients shown in figure 1 for the pooled sample give marginal effects similar to those found in the means comparison:

¹¹ 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 for whether the employer requested specific skills.

Table 2
Effects of Hiding Applicant Compensation History on Various Outcomes

	Control	Treatment	Difference	% Change
Measures of Employer Interest (No. of Applicants per Job)				
	<i>n</i> = 2,948	<i>n</i> = 2,974		
Viewed	6.67 (.17)	7.12 (.18)	.45 (.24)*	6.76
Measures of Employer Evaluation/Elicitation (No. of Applicants per Job)				
Messaged	1.80 (.07)	1.93 (.07)	.13 (.10)	7.09
Questioned (word)	1.12 (.04)	1.27 (.05)	.15 (.06)**	13.40
Questioned (mark)	1.19 (.04)	1.31 (.05)	.12 (.06)**	10.14
Face-to-face scheduled	.89 (.03)	.95 (.04)	.06 (.06)	6.30
Characteristics of Called-Back (i.e., Messaged) Applicants				
	<i>n</i> = 6,530	<i>n</i> = 6,725		
Bid amount	12.85 (.33)	12.40 (.29)	-.46 (.44)	-3.55
Profile wage rate	12.86 (.30)	12.36 (.26)	-.50 (.40)	-3.86
Average 6-month wage	11.09 (.29)	10.56 (.25)	-.53 (.38)	-4.78
Minimum 6-month wage	8.93 (.24)	8.53 (.21)	-.40 (.32)	-4.51
Maximum 6-month wage	14.11 (.40)	13.36 (.31)	-.76 (.50)	-5.37
Previous hours worked	1,107.33 (46.57)	1,072.50 (36.68)	-34.84 (59.28)	-3.15
Prior billed jobs	27.62 (.93)	26.73 (.79)	-.90 (1.22)	-3.24
Average feedback	4.70 (.01)	4.71 (.01)	.00 (.01)	.08
Job Opening Outcomes				
Hire made?	.40 (.01)	.43 (.01)	.03 (.01)**	7.23
Hourly wage rate	12.29 (.44)	10.93 (.36)	-1.35 (.56)**	-11.01
Feedback (1-10)	8.79 (.12)	8.91 (.10)	.12 (.15)	1.38
Contract rated a success?	.59 (.02)	.61 (.02)	.02 (.03)	3.65
Characteristics of Hired Applicants				
	<i>n</i> = 1,520	<i>n</i> = 1,744		
Bid amount	11.77 (.35)	10.53 (.47)	-1.25 (.58)**	-10.59
Profile wage rate	12.05 (.34)	11.03 (.48)	-1.02 (.59)*	-8.47
Average 6-month wage	10.33 (.38)	9.00 (.43)	-1.33 (.57)**	-12.87
Minimum 6-month wage	8.36 (.34)	7.10 (.38)	-1.26 (.51)**	-15.08
Maximum 6-month wage	13.11 (.49)	11.71 (.55)	-1.39 (.74)*	-10.64
Previous hours worked	1,140.67 (57.99)	1,252.57 (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
Average feedback	4.72 (.01)	4.71 (.02)	-.01 (.02)	-.20
Wage Bargaining Hire Made				
Any bargaining?	.13 (.01)	.12 (.01)	-.01 (.02)	-8.04
Wage-to-bid bargaining	.88 (.03)	.96 (.03)	.08 (.04)*	9.22

NOTE.—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 whether there was a hire is 2,974 treatment jobs and 2,948 control jobs. The sample sizes for evaluating wage rate and whether 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. Asterisks indicate *p*-values for a two-sided *t*-test of the null hypothesis of no difference in means across groups.

* *p* ≤ .10.
 ** *p* ≤ .05.




FIG. 1.—Effects of the treatment on extensive margin measure of employer information acquisition. 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. A color version of this figure is available online.

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 did control employers. While economically meaningful, the effect on the number of face-to-face meetings set up is borderline insignificant in the pooled sample.

Turning to the employer vertical preference subsamples, 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.¹² In fact, when restricting the sample to employers interested in hiring beginner applicants, the effect on the number of face-to-face meetings set up is significant at the 10% level.

The effects presented so far are all extensive margin effects, but we are also interested in intensive margin effects—that is, did treated employers evaluate the applicants they called back 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:

$$y_{ij} = \text{logit}^{-1}(\beta_0 + \beta_1 \text{WAGEHISTHID}_j + \mathbf{X}_{ij}\gamma + \epsilon_j) | \text{CALLED_BACK}_{ij} = 1, \quad (2)$$

where y_{ij} is some outcome for applicant i to job opening j , such as whether he or she was asked a question, and \mathbf{X}_{ij} is a collection of prerandomization job opening and applicant characteristics.¹³ 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 (again labeled “Pooled”) and each of the vertical preference tiers. Coefficients are exponentiated to be interpretable as odds ratios. The left panel of the figure reports estimates for the question marks, the center panel reports odds ratios for question words, and the right panel reports odds ratios for face-to-face meetings set up.

As with the extensive margin estimates, treatment effects appear to be concentrated among employers with low and medium vertical preferences. For example, employers with low vertical preferences are 1.51 times more likely to question an applicant they message (using the “question words”

¹² Another margin that 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],” “earn[ing—s],” or “rate[s]” in messages to applicants.

¹³ 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 for whether 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.




FIG. 2.—Effects of the treatment on intensive margin information acquisition measures, by employer vertical preferences. 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. A color version of this figure is available online.

measure). In contrast, employers with the highest vertical preferences show no treatment effects. Across groups, there is no evidence of an intensive margin difference in face-to-face interviewing.

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

B. Characteristics of Called-Back Applicants

In the absence of compensation history information, treated employers might change the kind of applicants they evaluate. The panel in table 2 labeled “Characteristics of Called-Back (i.e., Messaged) 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, although 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 had 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,

$$\begin{aligned} \text{CALLED_BACK}_{ij} = & \beta_0 + \beta_1 \log(\text{PROFILE_RATE}_{ij}) + \beta_2 \text{WAGE_HIST_HID}_j \\ & + \beta_3 (\log(\text{PROFILE_RATE}_{ij}) \times \text{WAGE_HIST_HID}_j) \quad (3) \\ & + \epsilon_{ij} | \text{VIEWED}_{ij} = 1, \end{aligned}$$

for the whole sample as well as for the three employer vertical preference tier subsamples.

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 callback probability toward workers with relatively lower profile rates. In the low (“Beginner”) and medium (“Intermediate”) vertical preference panels, we mostly see a shift in evaluation toward workers with relatively lower profile rates, whereas in the highest tier we simply see slightly less hiring of workers with the highest profile rates.




FIG. 3.—Employer probability of calling back a worker conditional on the applicant log profile rate, the treatment assignment, and the employer vertical preference. This figure plots, for each employer-requested expertise of the worker, predicted values from a linear model that regresses whether 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.

C. Probability of Hiring and the Characteristics of Hired Workers

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

As we saw in section IV.B, 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 toward 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.¹⁴ 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, 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 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.¹⁵

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 section A.2 of the appendix.

D. Prevalence and Outcome of Wage Bargaining

Although workers make the first wage offer, employers are free to make a counteroffer, which the worker can then 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:

¹⁴ 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 him or her directly with a private job opening, and we do not include private job openings in our sample.

¹⁵ The marginally filled jobs would need to average past wages of negative \$8.68 per hour for the 13% reduction in past wages to be entirely driven by the increased fill rate, which is impossible.

(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 who negotiate, but there is some evidence that hired workers in the treated group strike better wage bargains, conditional on bargaining.

Given that we know that 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 recapitulate the means comparison results in column 1 of table 3, estimating

$$\text{ANYBARGAINING}_j = \beta_0 + \beta_1 \text{WAGEHISTHID}_j + \epsilon_j | \text{HIRED}_j = 1. \quad (4)$$

As expected, the coefficient on the treatment indicator is essentially a precisely estimated zero—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,

Table 3
Effect of the Treatment on the Existence and Outcomes
of Hired Worker Wage Bargaining

	Dependent Variable		
	ANYBARGAINING (1)	WAGETOBIDRATIO (2) (3)	
Wage history hidden, WAGEHISTHID	-.005 (.016)	.089** (.043)	.281** (.122)
LPR			.008 (.034)
WAGEHISTHID × LPR			-.093* (.054)
Constant	.108*** (.012)	.836*** (.029)	.818*** (.083)
Observations	1,424	150	150

NOTE.—This table reports regressions where the outcomes are measures of bargaining. In col. 1, the outcome is whether any wage bargaining occurred. The sample is restricted to employers who made a single hire. In cols. 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 occurred. Heteroskedasticity-robust standard errors are reported in parentheses. The top and bottom 0.5% of wage-to-bid ratios are dropped. LPR = applicant profile rate in logs.

* $p \leq .10$.
 ** $p \leq .05$.
 *** $p \leq .01$.

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

$$\begin{aligned} \text{WAGE_TO_BID_RATIO}_j &= \beta_0 + \beta_1 \text{WAGE_HIST_HID}_j \\ &+ \epsilon_j | \text{ANY_BARGAINING}_j = 1. \end{aligned} \quad (5)$$

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 `WAGE_HIST_HIDj` 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).

E. 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, although it could be offset by the lower wage bill. We do not directly measure employer surplus, but we do have some reasonable 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 to 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, Horton, and Golden 2019).

It is important to note that for these outcome measures we are conditioning on filled jobs. This both introduces selection concern and means the sample is smaller. We have only 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 who actually leave feedback on completion of a job is not 100%. These caveats aside, both outcome measures in table 2 are higher in the treatment, although the effects are far from conventionally significant.

We expand our collection of contractual outcomes in table 4, estimating effects using linear regression and adding several prerandomization

Table 4
Effect of the Treatment on the Employer’s Subjective Measures of Contract Outcomes

	Dependent Variable					
	Log(Hours Billed) (1)	Feedback Left (2)	Public Feedback (3)	Private Feedback (4)	Bonus Given (5)	Rehired (6)
Wage history hidden	-.010 (.021)	-.002 (.077)	.001 (.056)	.002 (.115)	.935 (4.623)	-.003 (.015)
Constant	2.603*** (.105)	.746** (.377)	4.691*** (.160)	9.935*** (.095)	3.550 (3.675)	.110 (.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

NOTE.—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 col. 1 the outcome is an indicator for whether the employer left any public feedback. In col. 2 the outcome is the log of the number of hours billed on the job. In col. 3 the outcome is the z-score for the public feedback score left by the employer on the employee’s performance. In col. 4 the outcome is the z-score for the private feedback score left by the employer on the employee’s performance. In col. 5 the outcome is an indicator for whether the employer left the worker a bonus. In col. 6 the outcome is an indicator for whether the employer rehired the worker for a later job. Covariates included are category indicators, the number of prior jobs the employer filled, the number of prior jobs the employer and worker completed together, and total job billings.

** $p \leq .05$.
 *** $p \leq .01$.

controls.¹⁶ 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 (scale of 1 to 5); (3) the numerical private feedback on the worker (reported only to the platform; scale of 1 to 10), also standardized; (4) the dollar amount of any bonus left for the worker; and (5) whether the worker was rehired by the employer after completing the job.

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 at 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 that there appears to be no evidence of a change in hours worked in the treatment, although the estimate is imprecise.

Before presenting the feedback measures, in column 2 we first report a regression where the outcome is whether any feedback was left. 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

¹⁶ 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.

now being closer to zero compared with what we observed in table 2. In column 5 and 6, the outcomes are bonuses and rehires. 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.

F. 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, after the 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 postperiod in which the platform reverted to the preexperimental status quo.

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

$$y_{jt} = \beta_0 + \beta_1 \text{WAGEHISTHID}_j + \beta_2 t + \beta_3 (\text{WAGEHISTHID}_j \times t) + X_{jt} \gamma + \epsilon_j, \quad (6)$$

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 of the employer posting job opening j , and \mathbf{X}_j are prerandomization job characteristics.¹⁷ Our sample consists of employers who also posted a job with 1 month after the experiment ended. This sample is necessarily smaller, and to the extent that the treatment affected an employer's probability of posting a subsequent job opening in the postperiod, the sample could be selected, although 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 after the experiment, then the coefficient on $(\text{WAGEHISTHID}_j \times t)$ should be zero. If instead the employers revert, the point estimate on the interaction term should have the same magnitude but opposite sign as the coefficient on WAGEHISTHID_j .

Table 5 reports estimates of equation (6) using quasi maximum likelihood Poisson regressions with heteroskedasticity-robust standard errors. In general, the point estimate on the treatment indicator is similar to the experimental

¹⁷ Controls include the category of the job opening, employer's prior jobs, employer's prior spending, the number of applications to the job, the number of recommended applications to the job, a 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 (1)	Applicants Messaged, NUMMESSAGED (2)	Applicants Questioned (Marks), NUMQUESTIONEDMARKS (3)	Applicants Questioned (Words), NUMQUESTIONEDWORDS (4)	Meetings Set Up, NUMFACETOFACE (5)
WAGEHISTHID	.054 (.045)	.093 (.063)	.149** (.061)	.153** (.066)	.072 (.084)
Period, t	.018 (.048)	.048 (.068)	.092 (.069)	.093 (.074)	.036 (.079)
$t \times$ WAGEHISTHID	-.017 (.066)	-.032 (.101)	-.220** (.096)	-.246** (.105)	.022 (.124)
Constant	1.675*** (.133)	.536*** (.186)	.034 (.206)	.165 (.224)	-.522* (.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

NOTE.—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 in parentheses. The outcome in col. 1 is a count of the number of applicants viewed. The outcome in col. 2 is a count of the number of applicants messaged. The outcome in col. 3 is a count of the number of applications in which messages including a question mark were exchanged with the employer. The outcome in col. 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 col. 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 employer’s 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 for whether the employer requested a specific skill. The other regression coefficients are not shown.

* $p \leq .10$.
 ** $p \leq .05$.
 *** $p \leq .01$.

estimates, although it is less precise (as expected). In columns 1 and 2, we see some evidence that treated employers reverted to evaluating smaller pools in the postperiod, although the estimates are imprecise and smaller in magnitude than the treatment effects.

The coefficients on the interaction $WAGEHISTHID \times t$ in columns 3 and 4 are negative and significant, indicating that treatment employers revert in the postexperiment period. The effect sizes are large enough that the experimental effect is completely undone, although the effects are imprecise enough to make strong claims unwarranted. It appears that the up-front 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.

V. Discussion and Conclusion

This paper demonstrates how employers react when compensation history is removed. They respond primarily by acquiring more information, expanding their evaluation on both the extensive margin and the intensive margin. We find no evidence that they simply abandon their search or rely more heavily on other signals. There is no evidence they ultimately form worse matches. The strongest finding of the paper is a marked shift toward evaluating and hiring lower-wage workers.

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 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 with more information acquisition and a shift in hiring.

References

- Agan, Amanda, and Sonja Starr. 2017. Ban the box, criminal records, and racial discrimination: A field experiment. *Quarterly Journal of Economics* 133, no. 1:191–235.
- Agrawal, Ajay K., Nicola Lacetera, and Elizabeth Lyons. 2016. Does information help or hinder job applicants from less developed countries in online markets? *Journal of International Economics* 103:1–12.
- Altonji, Joseph G., and Charles R. Pierret. 2001. Employer learning and statistical discrimination. *Quarterly Journal of Economics* 116, no. 1:313–50.
- Autor, David H., and David Scarborough. 2008. Does job testing harm minority workers? Evidence from retail establishments. *Quarterly Journal of Economics* 123:219–77.

- Barach, Moshe A., Joseph M. Golden, and John J. Horton. 2020. Steering in online markets: The role of platform incentives and credibility. *Management Science* 66, no. 9:4047–70.
- Barach, Moshe A., Aseem Kaul, Ming D. Leung, and Sibio Lu. 2019. Strategic redundancy in the use of big data: Evidence from a two-sided labor market. *Strategy Science* 4, no. 4:298–322.
- Bertrand, Marianne, and Sendhil Mullainathan. 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review* 94, no. 4:991–1013.
- Binmore, Ken, Ariel Rubinstein, and Asher Wolinsky. 1986. The Nash bargaining solution in economic modelling. *RAND Journal of Economics* 17:176–88.
- Clemens, Jeffrey, and Michael Wither. 2014. The minimum wage and the Great Recession: Evidence of effects on the employment and income trajectories of low-skilled workers. NBER Working Paper no. 20724, National Bureau of Economic Research, Cambridge, MA. <http://www.nber.org/papers/w20724>.
- Craigie, Terry-Ann. 2020. Ban the box, convictions, and public employment. *Economic Inquiry* 58, no. 1:425–45.
- Dobbie, Will, Paul Goldsmith-Pinkham, Neale Mahoney, and Jae Song. 2016. Bad credit, no problem? Credit and labor market consequences of bad credit reports. NBER Working Paper no. 22711, National Bureau of Economic Research, Cambridge, MA. <http://www.nber.org/papers/w22711>.
- Doleac, Jennifer L., and Benjamin Hansen. 2020. The unintended consequences of “ban the box”: Statistical discrimination and employment outcomes when criminal histories are hidden. *Journal of Labor Economics* 38, no. 2:321–74.
- Filippas, Apostolos, John J. Horton, and Joseph M. Golden. 2019. Reputation inflation. NBER Working Paper no. 25857, National Bureau of Economic Research, Cambridge, MA. <http://www.nber.org/papers/w25857>.
- Gilchrist, Duncan S., Michael Luca, and Deepak Malhotra. 2016. When $3 + 1 > 4$: Gift structure and reciprocity in the field. *Management Science* 62, no. 9:2639–50.
- Goldin, Claudia, and Cecilia Rouse. 2000. Orchestrating impartiality: The impact of “blind” auditions on female musicians. *American Economic Review* 90, no. 4:715–41.
- Hall, Robert E., and Alan B. Krueger. 2012. Evidence on the incidence of wage posting, wage bargaining, and on-the-job search. *American Economic Journal: Macroeconomics* 4, no. 4:56–67.
- Hansen, Benjamin, and Drew McNichols. 2020. Information and the persistence of the gender wage gap: Early evidence from California’s salary history ban. NBER Working Paper no. w27054, National Bureau of Economic Research, Cambridge, MA.

- Holzer, Harry J. 1987. Hiring procedures in the firm: Their economic determinants and outcomes. NBER Working Paper no. 2185, National Bureau of Economic Research, Cambridge, MA. <http://www.nber.org/papers/w2185>.
- Horton, John J. 2010. Online labor markets. In *International workshop on internet and network economics*, 515–22. Berlin: Springer.
- . 2019a. Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution. *Management Science* 65, no. 8:3518–40.
- . 2019b. Price floors and employer preferences: Evidence from a minimum wage experiment. Working paper.
- Horton, John J., and Ramesh Johari. 2015. 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*, 507. New York: Association for Computing Machinery.
- Kahn, Lisa B., and Fabian Lange. 2014. Employer learning, productivity, and the earnings distribution: Evidence from performance measures. *Review of Economic Studies* 81, no. 4:1575–613.
- Kotlikoff, Laurence J., and Jagadeesh Gokhale. 1992. Estimating a firm's age-productivity profile using the present value of workers' earnings. *Quarterly Journal of Economics* 107, no. 4:1215–42.
- Kuhn, Peter, and Kailing Shen. 2013. Gender discrimination in job ads: Evidence from China. *Quarterly Journal of Economics* 128, no. 1:287–336. <https://doi.org/10.1093/qje/qjs046>.
- Lange, Fabian. 2007. The speed of employer learning. *Journal of Labor Economics* 25, no. 1:1–35.
- Nash, John F., Jr. 1950. The bargaining problem. *Econometrica* 18, no. 2:155–62.
- Oyer, Paul, and Scott Schaefer. 2011. Personnel economics: Hiring and incentives. In *Handbook of labor economics*, vol. 4, 1769–823. Amsterdam: Elsevier.
- Pallais, Amanda. 2014. Inefficient hiring in entry-level labor markets. *American Economic Review* 104, no. 11:3565–99.
- Rosen, Sherwin. 1982. Authority, control, and the distribution of earnings. *Bell Journal of Economics* 13, no. 2:311–23.
- Sattinger, Michael. 1993. Assignment models of the distribution of earnings. *Journal of Economic Literature* 31, no. 2:831–80.
- Shoag, Daniel, and Stan Veuger. 2016. No woman no crime: Ban the box, employment, and upskilling. Working Paper no. 16-015, Harvard Kennedy School, Cambridge, MA.
- Spence, Michael. 1973. Job market signaling. *Quarterly Journal of Economics* 87, no. 3:355–74. <https://doi.org/10.2307/1882010>.
- Stanton, Christopher, and Catherine Thomas. 2015. Landing the first job: The value of intermediaries in online hiring. *Review of Economic Studies* 83, no. 2:810–54.
- Terviö, Marko. 2009. Superstars and mediocrities: Market failure in the discovery of talent. *Review of Economic Studies* 76, no. 2:829–50.