

Sponsored Advertising in Labor Markets: Evidence from a Field Experiment*

Apostolos Filippas

John J. Horton

Prasanna Parasurama

Diego Urraca

April 24, 2025

Abstract

Digitization of labor markets and the recent proliferation of AI tools have dramatically lowered the search and application costs for jobseekers. Although lower costs are welcome, there is growing concern that this cost reduction has led to over-application, where jobseekers indiscriminately apply to a large number of jobs even when they are not a good fit. Jobseekers who are highly motivated and know they are an excellent fit have little ability to signal that fact and stand out from the crowd. A potential solution is a market-based mechanism that allows jobseekers to credibly signal their fit/interest to employers. We study the introduction of such a costly signaling mechanism, sponsored advertising, in a large online labor market. With sponsored advertising, workers could bid using an on-platform currency to sponsor their applications and receive one of three top positions in the employer’s application list. The platform randomly varied whether employers were exposed to sponsored applications, and within the exposure group, whether employers could see a disclosure that the application was sponsored. We find that sponsored applications are positively selected and that sponsoring an application increases the likelihood of a worker being hired by 41%. The experimental design allows us to estimate that 80% of this increase was due to the ranking effect—sponsored applications ranking higher—and 20% was due to the signaling effect—the disclosure that the application is sponsored. We find no difference in post-hire outcomes between the treated and control employers. We discuss the implications of our results for policy-makers and for designing sponsored advertising in online marketplaces.

*The latest version of this paper can be found at <https://www.apostolos-filippas.com>. “Diego Urraca” is a pen name for an employee of the marketplace where this experiment was conducted and is used to preserve the anonymity of the platform. This work was not subject to prior review and reflects only the opinions of the authors.

1 Introduction

Labor markets are increasingly computer-mediated (Autor, 2001; Varian, 2010). Technological advances have dramatically reduced jobseekers’ search and application costs, but there is a growing concern that this has also led to over-application.¹ Generative AI arguably exacerbates this problem by making it nearly costless for jobseekers to write cover letters, tailor or embellish their resumes, and apply to jobs indiscriminately—even when they are not a good fit.² As a result, jobseekers are competing for an interview against hundreds or thousands of other applicants, who are all seemingly well-qualified. At the same time, jobseekers who are highly motivated and know they are an excellent fit often have little ability to signal that fact and stand out from the broader applicant pool.

One possible solution is a paid mechanism for jobseekers to signal fit/interest to employers. Such costly signaling mechanisms already exist in other matching markets. For example, dating platforms allow users to send a scarce signal (e.g., “SuperLikes” on Tinder or “Roses” on Hinge) to potential matches to convey high interest and stand out.³ However, adverse selection is a first-order concern with this type of costly signaling: if low-quality jobseekers are more likely to pay to send the signal, or if employers perceive signaling jobseekers as desperate, then the signaling mechanism will not work in equilibrium.

In this paper, we study whether sponsored advertising, a type of costly signaling mechanism, can work in a labor market setting. We report the results of a field experiment conducted on a large online labor market. In this labor market, employers post jobs that can be done remotely, and workers search for and apply to these jobs. Employers view applications that are sorted algorithmically, and then shortlist, interview, and hire candidates. During the experiment, all workers became eligible to bid to “boost” (i.e., advertise) their applications when applying for a job. Boosted applications were displayed at the top of the employer’s job-specific application list, and in some cases, with a disclosure that the application was boosted. The platform randomly varied whether employers were exposed to boosted applications, and within the exposure group, whether employers could see a disclosure that the application was boosted.

The type of sponsored advertising described here has been adopted by numerous online platforms. Examples include sellers promoting their products on Amazon, real estate agents promoting their listings on Zillow, course creators promoting their courses on Udemy, and

¹For example, LinkedIn’s “Easy Apply” feature allows jobseekers to apply to jobs with a click of a button. See <https://www.linkedin.com/help/linkedin/answer/a512388>.

²For example, a WSJ article reports on the deluge of AI-crafted applications that recruiters have received since the advent of generative AI tools. See <https://www.wsj.com/tech/ai/automation-tools-ai-resumes-human-vetting-65a7100d?st>

³On Hinge, a popular dating app, users can send “Roses” to potential matches to signal high interest. Users are allotted one free Rose each week and have the ability to purchase more. See <https://help.hinge.co/hc/en-us/articles/36311177115027-Roses>. A similar mechanism is used on Tinder using “SuperLikes”.

daters promoting themselves to their potential matches on dating platforms. Sponsored advertising “works” in part because it places listings in a location more likely to be noticed—e.g., at the top of the search results. But as economists have long noted, there is also a potential signaling explanation for why advertising can be helpful at attracting buyers, where the act of advertising itself sends an informative signal to the buyer (Nelson, 1974; Kihlstrom and Riordan, 1984; Milgrom and Roberts, 1986). In an advertising or signaling equilibrium, high-quality sellers will have more to gain from advertising compared to low-quality sellers and therefore advertise at a higher level—i.e., positive selection into advertising. Buyers will in turn infer that the act of advertising itself is a signal of high quality. On the other hand, in a pooling equilibrium, if low-quality sellers also advertise, buyers may see advertising as a sign of desperation, and therefore as a signal of low quality.

Disentangling these distinct but concomitant effects of sponsored advertising is an empirical challenge. If advertisers have better outcomes than non-advertisers, it could be because they are positively selected, ranked higher, and/or because advertising sends a signal of high quality. The unique design of our experiment allows us to disentangle these effects of sponsored advertising into selection, ranking, and signaling effects, respectively.

During the experiment, all workers became eligible to bid to advertise through “boosted applications” when applying for a job. But what the employer saw depended on their treatment assignment. In the control group (PLACEBO), employers saw the organic results—they experienced no change to their application list. Employers in the active treatment cells saw the highest bidding applicants at the top of their application list. Alongside the ranking change, the platform further varied the information it displayed to treated employers. In the first treatment group (ADON), employers saw a disclosure that the application was boosted. In the second treatment group (ADNODISCLOSURE), employers did not see this disclosure. In the third treatment group (ADNOREC), employers saw this disclosure, but did not see an additional algorithmically determined recommendation label—which employers in all other groups saw.⁴ We summarize our main results below.

First, we find strong evidence of positive self-selection into advertising. Workers who advertised were more likely to be sought out by employers and more likely to be hired both before and during the experiment. To estimate the self-selection effect, we compare the outcomes of workers who chose to advertise to those who did not within the PLACEBO cell. This comparison allows us to observe the outcomes of boosted applications as if they had not been boosted, because boosting had no effect on the PLACEBO employer’s application list. We find that boosted applications in the PLACEBO cell were 101% (1.5 pp) more likely to result in a hire compared to non-boosted applications in the same cell. These effects remain present when we add worker fixed effects, suggesting that workers selectively boost when they

⁴Algorithmic recommendation labels are common in many platforms. For example, Amazon displays “Overall Pick” and “Best Seller” labels next to some listings in addition to sponsored listings.

are a particularly good match with a job, and/or they put more effort into applications that they choose to boost.

Second, we find that boosting an application increases the likelihood of a worker getting hired. To estimate the effect of boosting, we compare the difference in outcomes between boosted and non-boosted applications in the ADON cell to the difference in outcomes in the PLACEBO cell. This comparison differences out the self-selection effect, and isolates the boosting effect. Boosting an application increases the likelihood of a worker getting hired by 40.8% (1.2 pp).

Third, the experimental design allows us to decompose the effect of boosting into its signaling and ranking effects. To do so, we compare the difference in outcomes between the ADON cell—where boosting had both a ranking and a signaling effect—and the ADNODISCLOSURE cell—where boosting only had a ranking effect. We find that boosted applications being ranked higher increases the likelihood of a worker getting hired by 32.5%, and the disclosure that the application was boosted increases the likelihood of a worker getting hired by 8.3%. Percentage-wise, the ranking effect accounts for about 80% of the total effect of boosting, and the signaling effect accounts for about 20% of the total effect.

Fourth, we find that the effects of boosting are similar in the presence and in the absence of the platform’s algorithmically determined recommendation label. We compare the difference in outcomes between the ADON cell—where employers see the platform’s algorithmically determined recommendation label—and the ADNOREC cell—where employers do not see that label. In the ADNOREC cell, the effect of boosting on the likelihood of a worker getting hired is 41.5%. This is not a statistically significant difference compared to the effect of boosting in the ADON cell, 40.8%. We interpret this finding as evidence against the common concern that the effectiveness of sponsored advertising is lower when it competes with other algorithmic recommendation labels.

One might worry that employers, encountering advertising for the first time on the platform, might have been eager to try hiring advertising workers only to be disappointed with their performance later. However, this was not the case: on all available metrics for post-hire outcomes (amount spent, performance ratings, etc.), we find no detectable difference in employer outcomes between treated and control groups. We also do not find any detectable differences in other employer outcomes, including the number of interviews extended, the time to hire, and the total number of hires made. Employers simply substituted their hires in favor of advertising workers, and away from non-advertising workers. This type of crowd-out has also been documented in other labor market contexts (Crépon et al., 2013). We discuss some possible reasons why we do not observe a net increase in hires in Section 9.

This paper adds to the emerging empirical literature studying the effects of sponsored advertising in online marketplaces. Prior literature has studied sponsored advertising in online search engines (Blake et al., 2015; Coviello et al., 2017), e-commerce platforms (Moshary,

2024; Abhishek et al., 2022; Joo et al., 2024), and restaurant platforms (Sahni and Nair, 2020a; Dai et al., 2023), documenting both positive and negative effects on sellers. Although our labor market setting shares the basic economic problem of advertising as other platforms, there are several distinct characteristics about labor markets that arguably make it a stronger test bed for studying the economics of sponsored advertising. First, hiring is a high-stakes context, and engaging with all applicants can be beneficial to the employer for wage negotiation reasons. Whether advertising could meaningfully affect hiring decisions in a high-stakes environment was unclear. Second, there is arguably more information asymmetry in labor markets compared to e-commerce platforms. Workers have private information about their interest in a job that neither the employer nor the platform has access to. Because advertising is highly targeted in our setting—workers choose exactly which jobs to advertise to—boosting an application can send a positive signal of interest to the employer. At the same time, employers may expect high-quality workers to already have enough work to not have to advertise (Horton, 2019). This concern is particularly exacerbated in the labor market setting because workers have tight capacity constraints. Because a worker can only take on so many jobs at a time, and the marginal benefit of advertising is highest when the worker has little to no work, boosting an application could send a negative signal of desperation to the employer (Kroft et al., 2013). Despite these contrasting factors, we see positive selection into advertising and a positive effect of advertising on hiring outcomes.

Although the experiment shows that advertising “works” in practice, we also consider why it might work in theory. Specifically, we explore why the labor market setting might allow for positive selection into advertising and thus an advertising equilibrium to exist. Classical signaling models of advertising use repeat purchases as the linkage that prevents low-quality sellers from advertising at the same level as high-quality sellers (Nelson, 1974; Milgrom and Roberts, 1986). In a labor market setting, the multi-stage nature of the hiring process naturally provides the “repeat interaction” similar to repeat purchases in signaling models. If boosted applications increase the likelihood of a worker being interviewed, but not the likelihood of being hired *conditional* on being interviewed, then high-quality workers will have more to gain from advertising. This is because even if low-quality workers advertise to increase their chances of getting an interview, the employer will likely find out their true quality during the interview stage and will be less likely to hire them. This discourages low-quality workers from advertising at the same level as high-quality workers in the first place, leading to positive self-selection into advertising. We provide empirical evidence in support of this mechanism in our experiment: boosted applications increase the likelihood of a worker being interviewed, but not the likelihood of being hired *conditional* on being interviewed.

This paper makes several contributions to the economics of advertising. We disentangle the effects of advertising into selection, ranking, and signaling effects. We provide direct, experimental evidence of positive selection into advertising and the resulting signaling effect

of advertising as hypothesized by [Nelson \(1974\)](#); [Kihlstrom and Riordan \(1984\)](#); [Milgrom and Roberts \(1986\)](#). This is despite a search advertising context where one might presume only ranking effects matter. Second, because we can observe buyer outcomes, we can directly look at consumer welfare, finding that advertising in our context does not lead to an aggregate welfare loss. Lastly, to the best of our knowledge, this is the first paper to study sponsored advertising in a labor market setting.

The rest of the paper is organized as follows. Section 2 reviews relevant work on sponsored advertising. Section 3 describes our empirical context. Section 4 describes the design of the experiment and the estimation strategy. Section 5 reports the effects of boosted applications on workers. Section 6 examines how workers use boosted applications. Section 7 reports the effects of boosted applications on employers. Section 8 provides evidence for the mechanism that rationalizes our results. Section 9 discusses the implications of our results and concludes.

2 Related works

This paper is related to two broad streams of literature: the literature on the effects of advertising in sponsored search settings and the literature on algorithmic hiring.

First, this paper is related to the experimental literature on the effects of advertising in sponsored search settings. [Blake et al. \(2015\)](#) and [Coviello et al. \(2017\)](#) provide early evidence on the effects of advertising in search engines on website traffic, finding drastically different results (see also [Golden and Horton \(2021\)](#)). [Blake et al. \(2015\)](#) find that turning off paid advertising for eBay.com had no effect on website traffic to the site, while [Coviello et al. \(2017\)](#) find that turning off paid advertising for Edmunds.com led to more than a half-fold decrease in traffic. Both studies randomize on the seller side by varying the level of advertising. A drawback of this design is that it is difficult to scale the experiment to a large number of sellers. Instead, these studies only measure the effects of advertising for a single seller.

Recent platform-based studies have overcome this single-seller issue by randomly varying access to advertising for sellers or varying exposure of ads to buyers within a platform where there are a large number of both buyers and sellers. This literature is still nascent, and the empirical evidence on the effects of advertising on sellers is mixed. [Dai et al. \(2023\)](#) study advertising on a restaurant search engine (Yelp) and find that restaurants that randomly received access to free advertising saw an increase in purchase intention outcomes. [Moshary \(2024\)](#) finds that advertising increased the likelihood of a product being purchased in an e-commerce platform, but this came at the cost of reduced purchases for organic listings, leading to a *net* decrease in purchases. [Abhishek et al. \(2022\)](#) also study advertising in an e-commerce platform and find that additional advertising increased purchases for some categories but not others.

A number of studies aim to disentangle the ad disclosure effect from the ranking effect of sponsored advertising. [Sahni and Nair \(2020a\)](#) study advertising in a restaurant search engine (Zomato), where they experimentally manipulate the disclosure of ads to users. They find that disclosure that a listing is an advertisement increased calls to the restaurant (see also [Sahni \(2015\)](#); [Sahni and Nair \(2020b\)](#)). [Joo et al. \(2024\)](#) also manipulate ad disclosure in an e-commerce platform and find that ad disclosure in an e-commerce platform *decreased* click-through and conversion rates. A key design difference between our study and the above studies is that in both [Sahni and Nair \(2020a\)](#) and [Joo et al. \(2024\)](#), the ads can occupy any position in the search results, whereas in our design, the ads are always at the top of the search results. Our design therefore alleviates concerns about confounding disclosure effects with visibility or salience effects of advertising.

Our study is also related to the algorithmic hiring literature that studies the role of digital technologies in matching workers to jobs—e.g., ranking and recommendation algorithms on hiring platforms. Much of this work tends to be technical in nature, focusing on the design of the algorithm (see e.g., [Ramanath et al. \(2018\)](#); [Geyik et al. \(2019\)](#); [Kokkodis and Ipeirotis \(2023\)](#)). A few experimental studies have examined the effects of algorithmic hiring on a number of labor market outcomes ([Horton, 2017](#); [Cowgill, 2019](#); [Li et al., 2020](#)). These studies show how algorithmic recommendations can lead to higher fill rates and better matches. When making hiring recommendations, these algorithms take into account observable job and worker characteristics that the platforms have visibility into. Workers have private information about their interest and own fit with a job that neither the employer nor the platform has. This information could be useful for the platform when ranking and recommending workers but is missing for economic reasons, and not for technical reasons. In this paper, we study how incorporating this private information (through a market-based mechanism) into the ranking algorithm affects worker and employer outcomes; and whether advertising competes with other algorithmic recommendation labels.

3 Empirical context

Our study is conducted in a large online labor market ([Horton, 2010](#); [Agrawal et al., 2015](#); [Horton et al., 2017](#)). In online labor markets, employers hire workers to perform tasks that can be done remotely, such as computer programming, graphic design, and writing. Each market differs in its scope and focus, but platforms commonly provide ancillary services that include maintaining job listings, hosting user profile pages, arbitrating disputes, certifying worker skills, and maintaining feedback systems ([Filippas et al., 2020](#)).

The most important features of conventional labor markets also exist in our context. Employers and workers are free to enter and exit the market at any time. Employers post job descriptions, and workers search for and apply to jobs. Employers can assess promising

candidates through interviews. Employers and workers can negotiate over wages, which take the form of either hourly salaries or fixed amounts, and form contracts. More generally, employers and workers face substantial search frictions, barriers to entry, and information asymmetries (Pallais, 2013; Stanton and Thomas, 2016; Horton, 2017, 2019; Benson et al., 2019; Filippas et al., 2021).

3.1 The status-quo job application process

Workers search for job openings and can apply to any job by using an in-platform currency called “coins.”⁵ The number of coins required to apply to a job—the cost of an application—is determined by the platform using a proprietary formula that takes into account only job-specific attributes, such as the anticipated job duration and earnings. Employers may also invite workers to apply to jobs; a job application following an employer invite uses up no coins (Filippas et al., 2024).

Each employer has a job-specific application tracking system (ATS). The ATS keeps track of all applications sent to a given job posting. In the ATS, the tile of each application conveys information including the name and profile picture of the worker, the amount of money the worker has earned on the platform, and a snippet of the cover letter the worker sent along with their application. Applications are displayed in a ranked order determined by the platform’s proprietary algorithm. If an application is determined algorithmically to be a good match for the job posting, a “Best Match” label is displayed on the tile of the application. Appendix A.2 provides an illustration of the employer’s ATS interface.

Table 1 provides summary statistics on the status quo job search and matching, for job postings created in the pre-experimental period. Job postings received 20.1 applications on average, with some postings receiving as many as 2,000 applications. The average job fill rate was 0.53, meaning that employers made a hire in about half of the jobs. Although the median number of hires per job was 1, some jobs had multiple hires. On the worker side, workers submitted 5.07 applications on average, and were hired for 0.17 jobs. The probability of an application leading to a hire was about 0.03.

4 Experiment

4.1 Experimental design

During the experiment, all workers became eligible to bid in a sealed auction to “boost” their job applications. There were three boosted application slots per job. Employers were allocated randomly to one of four treatment groups upon posting a job. The treatment

⁵Coins are sold through the platform and cost \$0.15 each, are placed in a non-interest-bearing account, cannot be converted back to cash, and expire one year after the purchase.

Table 1: Pre-experimental summary statistics

	Mean	Median	SD	Min	Max
<i>Employer/Job statistics</i>					
number of posts per employer	4.49	2	10.10	1	1,255
number of apps per job	20.10	13	28.67	0	2,243
number of invited apps per job	5.19	1	58.83	0	12,588
number of hires per job	0.75	1	2.01	0	228
job filled indicator	0.53	1	0.50	0	1
amount spent per job	256.63	0	1360.36	0	251,551
<i>Worker statistics</i>					
number of applications	5.07	2	17.58	1	3,401
number of invites received	0.47	0	2.27	0	183
number of contracts formed	0.17	0	0.75	0	80
number of contracts per application	0.03	0	0.13	0	1
average hourly asking wage	22.15	15	34.75	0	999
average fixed asking wage	785.97	75	11940.27	5	1,000,000

Notes: This table reports the summary statistics on the status-quo job search and matching on the platform. The sample includes jobs posted in the pre-experimental period, between June 8, 2021 and September 8, 2021. For employer/job-level statistics, we report (i) the number of job postings per employer, (ii) the number of applications per job posting, (iii) number of invited applications per job posting, (iv) the number of workers hired per job posting, (v) whether the job filled, that is, at least one worker was hired for a given job, and (vi) the total amount of money spent per job posting in the 60-day period after being posted on the platform. On the worker side, we report (i) the number of applications, (ii) the number of invites received, (iii) the number of contracts formed (iv) the number of contracts formed per application sent, (v) the average asking wage for hourly jobs, and (vi) the average asking wage for fixed-price jobs.

changed the version of the ATS that each employer saw, but workers did not know which treatment group each job belonged to. We describe the employer treatment arms below, and we also summarize them in Table 2.

- **ADON:** Boosted applications were displayed at the top of the employer’s ATS, and a “Highly Interested” label was displayed on each application’s tile. Hovering over the label revealed that the worker paid to get noticed.
- **ADNoDISCLOSURE:** Boosted applications were displayed at the top of the employer’s ATS, but did not include a “Highly Interested” label—i.e., there was no disclosure that the application was boosted.
- **ADNoREC:** Boosted applications were displayed at the top of the employer’s ATS, and included a “Highly Interested” label, but the algorithmically determined “Best Match” label was not displayed in any application.
- **PLACEBO:** Boosted applications had no effect on the employer’s ATS—there was neither

a ranking change nor a “Highly Interested” label. Employers were shown the organic results.

Table 2: Comparison of feature changes in ATS across treatment groups

	ADON	ADNoDISCLOSURE	ADNoREC	PLACEBO
Boosted apps pinned on top	✓	✓	✓	
“Highly Interested” label	✓		✓	
“Best Match” label	✓	✓		✓

The design of the experiment allows us to answer a rich set of questions about the effects of boosted applications. First, we can examine whether there is positive selection into advertising, by comparing the outcomes of workers who chose to boost their applications to those who did not within the PLACEBO cell. Second, we can estimate the causal effect of the boosted applications on employer and worker outcomes, by comparing the ADON and the PLACEBO cells. Third, we can separate the ranking effect (applications ranking higher) from the signaling effect (displaying the “Highly Interested” label) of a boosted application, by comparing the ADON and the ADNoDISCLOSURE cells against the PLACEBO cell. Fourth, we can examine the difference in the relative efficiency of the algorithmically determined and the sponsored ad-determined labels, by comparing the ADON and the ADNoREC cells.

4.2 The boosted application auction format

Workers could bid using coins to compete for the boosted application slots, and a sealed bid auction was used to determine the winners. The auction worked as follows: for a job posted by a treated employer, interested workers could set the maximum number of coins they were willing to spend for a boosted application slot. The top 3 bidders at any given point in time were pinned to the boosted application slots. The winners paid the lowest winning bid if they were in a boosted application slot at the end of the auction (7 days after the job posting date), or if they had an interaction with the employer while they were in a boosted application slot. A worker would get a full reimbursement either if her application was outbid and not interacted with, or if the employer who posted the job was allocated to the PLACEBO cell.⁶ In a special case when the number of bidders was less than three, the lowest winning bid was considered to be zero; in such cases, all bidders, de facto, would have their bids reimbursed. Appendix A.2 provides an illustration of the worker bidding interface.

⁶For workers that did not get the disclosure, the platform refunded the coins they spent on the boosted application after the experiment.

4.3 Treatment administration

The experiment began on September 8, 2021 and ended on October 13, 2021. A total of 106,788 employers were part of the experiment, of whom 37,616 (35.22%) were allocated to ADON, 37,417 (35.04%) were allocated to PLACEBO, 15,838 (14.83%) were allocated to ADNoREC, and 15,917 (14.91%) were allocated to ADNoDISCLOSURE. A total of 510,975 workers were part of the experiment, and they submitted 3,665,555 applications to 167,322 jobs during the experiment. These sample sizes were selected based on a power analysis to detect a 2% change in the probability that a treated employer would make a hire within 7 days with 80% power.

All employers in our data received the “correct” treatment and remained in the same experimental group throughout the experiment. The platform did not inform the employers that they received different treatments. The experimental groups were well-balanced. In Appendix A.1, we report two-sided t-tests for various employer-level attributes, and plot allocations over time.

5 The effects of boosted applications on workers

5.1 Estimation strategy

We estimate the treatment effect of boosting on workers’ outcomes by comparing the differences in outcomes between the active treatment groups and the PLACEBO group. Recall that the randomization was at the employer level, but we are interested in effects at the worker/application level. This requires us to use the following specification to estimate the causal effect of boosted applications on workers’ outcomes:

$$\begin{aligned} y_{i,j} = & \beta_0 \\ & + \beta_1 \text{TRTADON}_j + \beta_2 \text{TRTADNoDISCLOSURE}_j + \beta_3 \text{TRTADNoREC}_j \\ & + \beta_4 \text{BOOST}_{i,j} \\ & + \beta_5 (\text{TRTADON}_j \times \text{BOOST}_{i,j}) \\ & + \beta_6 (\text{TRTADNoDISCLOSURE}_j \times \text{BOOST}_{i,j}) \\ & + \beta_7 (\text{TRTADNoREC}_j \times \text{BOOST}_{i,j}) \\ & + \epsilon_{i,j}, \end{aligned} \tag{1}$$

where $y_{i,j}$ is the outcome of interest for an application submitted by worker i for job posting j , TRTADON_j , $\text{TRTADNoDISCLOSURE}_j$ and TRTADNoREC_j indicate the treatment assignment for the employer who posted job j , $\text{BOOST}_{i,j}$ is an indicator for whether the application submitted by worker i for job j was boosted, and $\epsilon_{i,j}$ is the error term.

It is worth examining what each of the coefficients of Equation 1 captures. The coefficient β_0 is the average outcome of non-boosted applications in the PLACEBO group. The coefficients β_1, β_2 , and β_3 are the differences in the outcomes of non-boosted applications in the three active treatment groups, compared to the PLACEBO group. These coefficients capture any crowd-out effects of boosted applications on non-boosted applications (if the coefficients are negative) or any positive spillover effects of boosted applications on non-boosted applications (if the coefficients are positive). The coefficient β_4 is the difference in the outcome of boosted and non-boosted applications in the PLACEBO group. Because workers self-selected into submitting a boosted application, and boosted applications had no effect on employers in the PLACEBO group, β_4 is an estimate of the self-selection effects of boosting. A positive β_4 would indicate that boosted applications are positively selected due to higher-quality applicants self-selecting into boosting their applications and/or due to the greater effort workers put into applications they boost. A negative β_4 would indicate that boosted applications are adversely selected due to, for example, lower-quality applicants boosting their application out of desperation. Finally, the coefficients β_5, β_6 , and β_7 measure the difference in outcomes between boosted and non-boosted applications in the three active treatment groups, compared to the PLACEBO group. These estimates capture the causal effect of boosted applications on workers' outcomes in each of the three active treatment groups.

We estimate Equation 1 for two outcomes: (a) whether worker i was interviewed for job j , and (b) whether worker i was hired for job j . We report the regression estimates for interview and hire outcomes in Table 3. In each table, Columns (1-2) reports regression estimates for the specification of Equation 1. Columns (3-4) add job posting fixed effects to control for heterogeneity in boosted applications and outcomes across jobs. This is useful for scenarios such as when workers boost their applications at a higher rate for jobs that receive more applications and thus have a lower baseline probability of being hired, which could bias the treatment effect estimates. Columns (5-6) add worker fixed effects to use within worker variation in boosted applications. Columns (7-8) add both worker and job posting fixed effects. The worker + job posting fixed effects is our preferred specification because it most closely answers our core question: for a given job, what is the effect of a worker boosting their application?

In Figure 1, we plot the estimated effects of boosting on being hired as a percent increase compared to the baseline hire rate. This allows for an easier, visual comparison of the relative magnitude of each effect.

5.2 Boosting applicants were positively selected

In Table 3, the estimate for the variable BOOST captures the difference in outcomes between boosted and non-boosted applications within the PLACEBO group. Looking at Columns (1-

Table 3: Treatment effect estimates of boosting on the likelihood of being interviewed and hired

Dependent Variables: Specification:	Interviewed OLS	Hired OLS	Interviewed Job FE	Hired Job FE	Interviewed Worker FE	Hired Worker FE	Interviewed Worker + Job FE	Hired Worker + Job FE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PLACEBO	0.0741*** (0.0022)	0.0150*** (0.0008)						
AdON	-0.0047‡ (0.0024)	-0.0021* (0.0009)			-0.0044* (0.0017)	-0.0020** (0.0006)		
AdNoREC	-0.0052‡ (0.0028)	-0.0025** (0.0009)			-0.0045* (0.0021)	-0.0022*** (0.0007)		
AdNoDISCLOSURE	-0.0054‡ (0.0028)	-0.0019* (0.0009)			-0.0050* (0.0021)	-0.0020** (0.0007)		
BOOST	0.0597*** (0.0017)	0.0152*** (0.0007)	0.0247*** (0.0009)	0.0089*** (0.0005)	0.0507*** (0.0015)	0.0103*** (0.0007)	0.0108*** (0.0010)	0.0035*** (0.0006)
AdON × BOOST	0.0352*** (0.0021)	0.0112*** (0.0009)	0.0376*** (0.0014)	0.0119*** (0.0008)	0.0350*** (0.0019)	0.0114*** (0.0009)	0.0384*** (0.0015)	0.0123*** (0.0008)
AdNoREC × BOOST	0.0334*** (0.0026)	0.0115*** (0.0011)	0.0365*** (0.0019)	0.0124*** (0.0010)	0.0330*** (0.0025)	0.0112*** (0.0011)	0.0372*** (0.0020)	0.0125*** (0.0011)
AdNoDISCLOSURE × BOOST	0.0250*** (0.0027)	0.0089*** (0.0011)	0.0259*** (0.0019)	0.0094*** (0.0010)	0.0246*** (0.0025)	0.0088*** (0.0011)	0.0272*** (0.0020)	0.0098*** (0.0011)
<i>Fixed-effects</i>								
Job posting			✓	✓			✓	✓
Worker					✓	✓	✓	✓
<i>Fit statistics</i>								
Observations	3,403,094	3,403,094	3,403,094	3,403,094	3,403,094	3,403,094	3,403,094	3,403,094

Clustered (employer) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1*


Notes: This table reports the estimates of the effect of boosted application on the likelihood of being interviewed and hired. The independent variables are: (i) a binary indicator for whether an application is a boosted application, (ii) treatment indicators for the experimental assignment of the employer, and (iii) interactions between the boosted application and treatment indicators. Columns (1-2) of the table reports the OLS estimates, Columns (3-4) include only job posting fixed effects, Columns (5-6) include only worker fixed effects, and Columns (7-8) include both worker and job posting fixed effects.

2), boosted applications were 80.5% (5.97 pp) more likely to be interviewed and 101% (1.5 pp) more likely to be hired compared to non-boosted applications. Because boosting had no effect on the employer's application list in the PLACEBO cell, these results indicate that boosted applications are positively selected.

Looking at Columns (5-8), we see that the estimate for the variable BOOST remains positive when adding worker fixed effects. This suggests that workers selectively boost when the match quality is high for a particular job, and/or put more effort into applications they choose to boost. In other words, workers' boosting is specific and targeted.

We provide further evidence of positive selection into boosting in Section 6 using observational data. We find that boosters were more likely to be sought out by employers even in the pre-experimental period.

Figure 1: Selection, ranking, and signaling effects of boosting



Notes: This figure plots the estimated effects of boosting on a worker being hired as a percentage increase from the baseline probability of being hired. The estimates are based on the specification that includes worker and job fixed effects (Column (8) of Table 3). In the PLACEBO cell, workers could self-select into boosting their applications, but employers did not see the boosted applications at the top of the ATS, so this PLACEBO estimate captures the self-selection effects of boosting. In the ADNODISCLOSURE cell, employers could see the boosted applications at the top of the ATS, but did not see the disclosure that the application was boosted, so the difference between the ADNODISCLOSURE and the PLACEBO estimates captures the ranking effect of boosting. In the ADON cell, employers could see both the boosted applications and the disclosure that the application was boosted, so the difference between the ADON and the ADNODISCLOSURE estimates captures the signaling effect of boosting.

5.3 The aggregate effect of boosting

The interaction terms in Table 3 capture the causal effect of boosting on the likelihood of being interviewed and hired. Looking at Columns (7-8), we see that boosting an application increases the likelihood of a worker being interviewed by 28.7% (3.8 pp) and the likelihood of being hired by 40.8% (1.2 pp)—compared to the corresponding outcomes in the PLACEBO cell. The estimates for the ADON, ADNODISCLOSURE, and ADNOREC are all negative, suggesting that non-boosted applications are less likely to result in an interview or a hire in the three active treatment cells compared to the PLACEBO cell. This indicates a crowd-out effect of boosted applications on non-boosted applications.

In Appendix B, we estimate the same specification using as our outcome variable an

indicator for whether an application led both to a hire and to the hired worker earning a positive amount for the job and including invited workers in the sample. This is a more restrictive robustness check for our results because we are moving further downstream of the treatment. The results remain highly similar to those of Table 3.

5.4 The ranking effect of boosting

In Table 3, the estimates for the interaction term $\text{AdNoDISCLOSURE}_j \times \text{BOOST}_{i,j}$ captures the ranking effect of boosting applications. The reason is that in the AdNoDISCLOSURE treatment arm, employers could not see that applicants had boosted their applications, but boosted applications were displayed at the top of the employers’ application tracking system.

Ranking higher increased the likelihood of a worker being interviewed by 20.3% (2.7 pp), and being hired by 32.5% (1 pp). This is in accordance with results from previous research, which show that the allocation of the top “real estate” influences buyer outcomes substantially.

5.5 The signaling effect of boosting

Boosted applications ranked at the top of employers’ ATS in both the AdON and the AdNoDISCLOSURE cells, but employers could see the disclosure that an application was boosted only in the AdON cell. This allows us to estimate the signaling effect of boosting by comparing application outcomes between these two cells. Specifically, the difference in the treatment effects between these cells provides the causal effect of the disclosure on the likelihood of a worker being interviewed and hired.

In the specification of Equation 1, β_5 captures the total effect of boosted application, β_6 captures the ranking change effect of boosted application, and hence the difference $\beta_5 - \beta_6$ captures the signaling effect. Disclosing that an application is boosted increases the likelihood of being interviewed by 8.4% (1.1 pp; $p\text{-value} = 7.4\text{e-}08$), and being hired by 8.3% (0.2 pp; $p\text{-value} = 0.027$). Conducting linear hypothesis tests against the hypothesis that $\beta_5 - \beta_6 = 0$ shows that the estimates are statistically significant.

One concern with our interpretation of this estimate is that the disclosure may have an attention-grabbing or visibility effect, rather than a signaling effect. However, because boosted applications only appear in the top 3 positions of the application list, it is unlikely that the disclosure increased the visibility of the top-ranked applicants. To build more confidence in our findings, we test the relationship between the disclosure effect and the selection effect. If the effect we detect is explained by signaling, we should expect the disclosure effect to be stronger in jobs where workers are more positively selected; if the effect can be explained by attention-grabbing effect, we would expect to find no relationship between the two.

Our approach to test for the attention-grabbing effect is as follows. First, we split our sample based on the category to which a job belongs.⁷ Then, for each subsample, we estimate the specification of Equation 1 including both worker and job posting fixed effects, and we obtain estimates for the selection effect β_4 and the disclosure effect $(\beta_5 - \beta_6)$. Finally, we regress the disclosure effect on the selection effect. We report the results of our empirical exercise in Table 4.

In Column (1) of Table 4, we report the weighted least squares estimates, where the weights are the number of jobs in each category. In Column (2) we report estimates using an errors-in-variables (EiV) model; this allows us to account for the fact that the selection and disclosure effects are estimated with error. Across both estimators, we find a positive relationship between the selection and disclosure effects. This corroborates our interpretation that the disclosure effect that we detect is, at least in part, a signaling effect, rather than a pure attention-grabbing effect.

Table 4: Disclosure vs. Selection Effect

Estimator:	WLS	EiV
(Intercept)	-0.022 (0.075)	-0.081 (0.066)
Selection Effect	0.237* (0.097)	0.32** (0.098)
<i>Fit statistics</i>		
Observations	170	170

Standard-errors in parentheses

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, †: 0.1*

Notes: This table reports regression estimates of the disclosure effect vs. selection effect. The dependent variable is the disclosure effect in job category k , and the independent variable is the selection effect in job category k . The first column reports the weighted least squares estimates, where the weights are the number of jobs in each category. The second column reports the errors-in-variables estimates, where the errors are the standard errors of the selection and disclosure effects.

5.6 The interplay of boosting and algorithmic recommendations

The presence of algorithmic recommendations in the employer’s ATS could interact with the effects of boosting. In our context, recall that algorithmic recommendations take the form of a “Best Match” label that is displayed on the tile of applications that the platform’s algorithm determines to be a good match for the job posting.

⁷Note that we have to conduct this analysis at the job sub-category level, rather than at the job level, because there is no variation in treatment within a job. There are 176 such job sub-categories, such as Article & Blog Writing, 3D Animation, Back-End Development, and so on.

We compare the relative effect of boosting in the presence and absence of the algorithmically determined “Best Match” label by comparing the ADON and ADNOREC cells. Specifically, we test against the hypothesis that the treatment effect of boosted application in the ADNOREC cell is equal to the treatment effect of boosted application in the ADON cell, i.e., $\beta_5 - \beta_7 = 0$. We find that we cannot reject this hypothesis, that is, the difference in the two estimates is not statistically significant. This suggests that the algorithmic recommendations of the platform do not affect the effects of boosting.

5.7 Heterogeneity analysis

The effects of boosted applications could vary by job category. In our context, each job belongs to one of 12 broad job categories, such as Writing, Data Science & Analytics, and Web, Mobile & Software Development. We split our data by job category, estimate the specification of Equation 1 including worker and job posting fixed effects specification for each subsample, and we report the ADON versus the PLACEBO treatment effect estimates in Figure 2a. The left panel shows the nominal treatment effect estimates. Because the baseline hire rate varies across job categories, in the right panel we report the relative treatment effect estimates as a percentage increase from the baseline hire rate. There is some heterogeneity across job types, but the estimates are positive across all job types.


The effects of boosted applications could also vary by worker experience on the platform. We replicate the analysis of Figure 2a, but now splitting our data by worker experience on the platform. We measure worker experience by counting the total number of jobs held by the worker in the pre-experiment period and bucketing them into four groups: 0 jobs (i.e., no jobs held in the pre-experiment period), 1 job, 2 jobs, and 3 or more jobs. We report the results of this empirical exercise in Figure 2b. The left panel shows the nominal treatment effect and the right panel shows the relative treatment effect as a percentage increase from the baseline hire rate within that category. We find that the effect of boosted application on the probability of being hired is positive across all groups. The nominal treatment effect is smaller for inexperienced workers, i.e., workers who held 0 jobs in the pre-experiment period. However, this is because the baseline hire rate for inexperienced workers is lower compared to other groups. Looking at the treatment effects as a percent increase from the baseline hire rate, the relative treatment effect is slightly higher for less experienced workers, although this difference is not statistically significant.

5.8 Post-hire outcomes for boosted vs non-boosted applications


We estimate the effects of boosted applications on post-hire outcomes for workers conditional on being hired. Although this analysis is observational because we condition on a post-treatment variable, it provides us with an informative comparison of the quality of boosting

Figure 2: Heterogeneity analysis of the effects boosting on hiring

(a) By job type



(b) By worker experience on the platform



Notes: This figure shows the treatment effect estimates of boosted application on the probability of being hired. In the top panel, we split our sample by job type, estimate the specification of Equation 1 including worker and job posting fixed effects FE model, and report the ADON vs PLACEBO treatment effects. In the bottom panel, we split our sample by worker experience on the platform, estimate the same model for each experience group, and report the ADON vs PLACEBO treatment effects. In both panels, the left side shows the nominal treatment effect estimates and the right side shows the relative treatment effect estimates as a percentage increase from the baseline hire rate of each respective group.

and non-boosting applicants. We estimate the following regression specification:

$$\begin{aligned}
y_{i,j} | \text{hired} = & \beta_0 \\
& + \beta_1 \text{BOOST}_{i,j} \\
& + \beta_2 (\text{TRTAdON}_j \times \text{BOOST}_{i,j}) \\
& + \beta_3 (\text{TRTAdNoDISCLOSURE}_j \times \text{BOOST}_{i,j}) \\
& + \beta_4 (\text{TRTAdNoREC}_j \times \text{BOOST}_{i,j}) \\
& + \gamma_j + \epsilon_{i,j}.
\end{aligned} \tag{2}$$

We consider three post-hire outcomes: (1) earnings from the job, (2) private feedback from the employer to the worker, and (3) public feedback from the employer to the worker. We report the results in Table 5. We do not find statistically significant effects of boosted applications on post-hire outcomes conditional on being hired. This suggests that boosting workers achieved comparable outcomes to non-boosting workers after being hired for a job.

Table 5: OLS estimates of post-hire outcomes

Dependent Variables:	log(Earnings) (1)	log(Private feedback) (2)	log(Public feedback) (3)
BOOST	0.0879* (0.0433)	-0.0029 (0.0198)	-0.0167 (0.0155)
BOOST \times AdON	-0.0589 (0.0613)	0.0019 (0.0268)	-0.0274 (0.0212)
BOOST \times AdNoREC	-0.0659 (0.0764)	0.0006 (0.0362)	0.0181 (0.0225)
BOOST \times AdNoDISCLOSURE	-0.2187‡ (0.1177)	0.0348 (0.0346)	-0.0063 (0.0285)
<i>Fixed-effects</i>			
Job posting	✓	✓	✓
<i>Fit statistics</i>			
Observations	46,233	33,001	35,257

Clustered (employer) standard-errors in parentheses

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1*

Notes: This table reports the OLS estimates of the differences in post-hire outcomes between boosted applications and non-boosted applications conditional on being hired. Column (1) reports earnings from the job, (2) reports private feedback from the employer, and (3) reports public feedback from the employer.

6 Workers use of boosted applications

In this section, we examine how often workers use boosted applications, what type of jobs they use boosted applications for, and what type of workers use boosted applications.

6.1 Intensity of using boosted applications

22.9% workers boosted at least once during the experiment period. Among these workers, the median worker boosted 33.3% of their applications. However, this result is skewed by the fact that workers who submitted only one or two applications boosted all their applications. Further restricting the sample to those workers who submitted at least 5 applications during the experiment period, we find that the median worker boosted 17.7% of their applications. This suggests that workers use boosted applications judiciously. Figure 3 shows the distribution of these numbers across workers.

6.2 When workers boost their applications

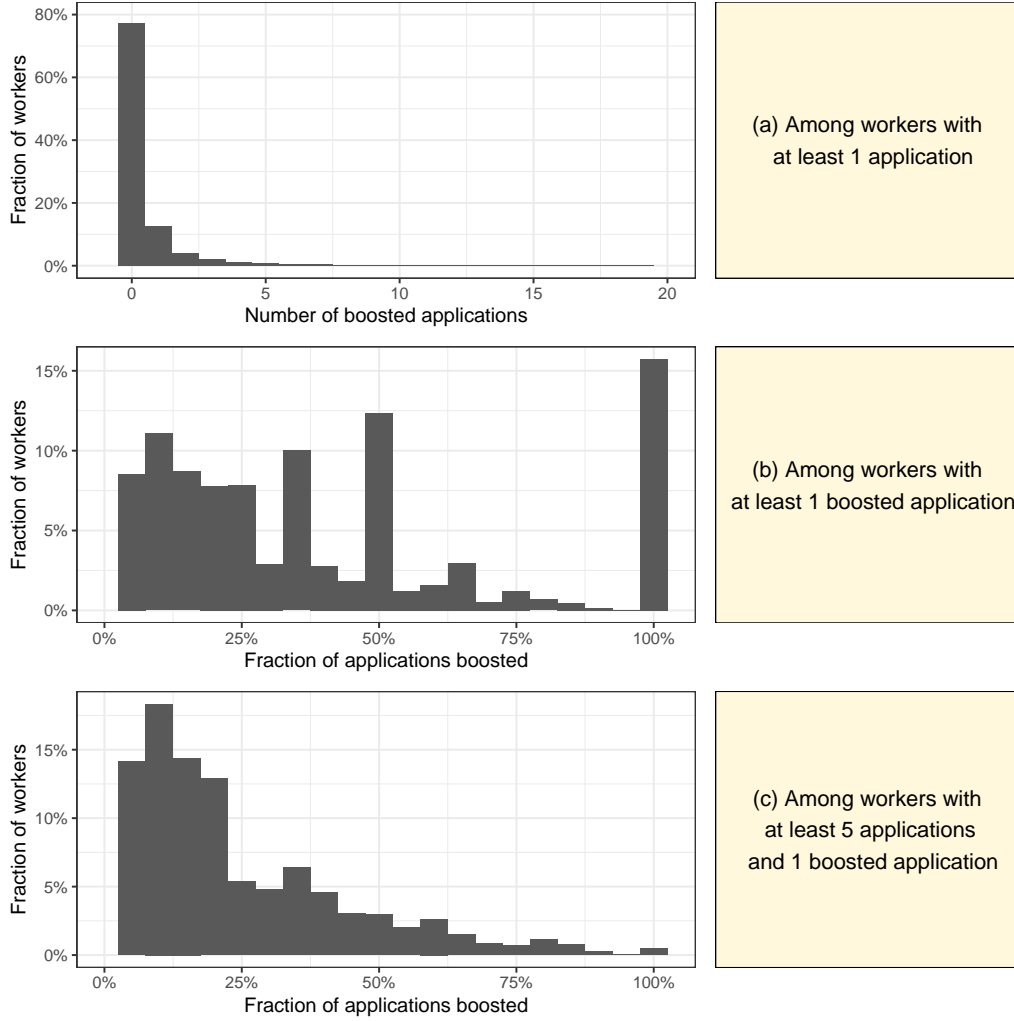
We next examine what type of jobs workers use boosted applications for. Specifically, we consider two main characteristics of the job that is salient to the worker when deciding whether to place a boosted application bid: job congestion and job value. We measure job congestion using the number of applications on the job posting, and job value using the average wage bid for the job. We estimate the following regression specification:

$$y_{i,j} = \beta_0 + \beta_1 NApps_j + \beta_2 NApps_j^2 + \beta_3 AvgWageBid_j + \beta_4 AvgWageBid_j^2 + \epsilon_{i,j} \quad (3)$$

where $y_{i,j}$ is an indicator for (a) whether the worker i placed a bid for a boosted application slot and (b) whether the worker won a boosted application slot for job j , $NApps_j$ is the number of applications received for job j , and $AvgWageBid_j$ is the average wage bid received for job j . We estimate this specification using logit regression and plot the predicted probabilities in Figure 4.

For job congestion, we see an inverted U-shaped relationship in the probability of placing a boosted application bid. Workers are most likely to place a boosted application bid when the job is moderately congested. However, the shape of the relationship is inverted for the probability of winning a boosted application slot. As more workers compete for a boosted application slot for these jobs, the probability of winning a boosted application slot decreases. For job value, the probability of boosted application bid increases with job value up to \$100/hr, and then decreases. The probability of winning a boosted application slot, increases with job value up to \$125/hr, after which the probability estimates become too imprecise.

Figure 3: Distribution of boosted applications across workers

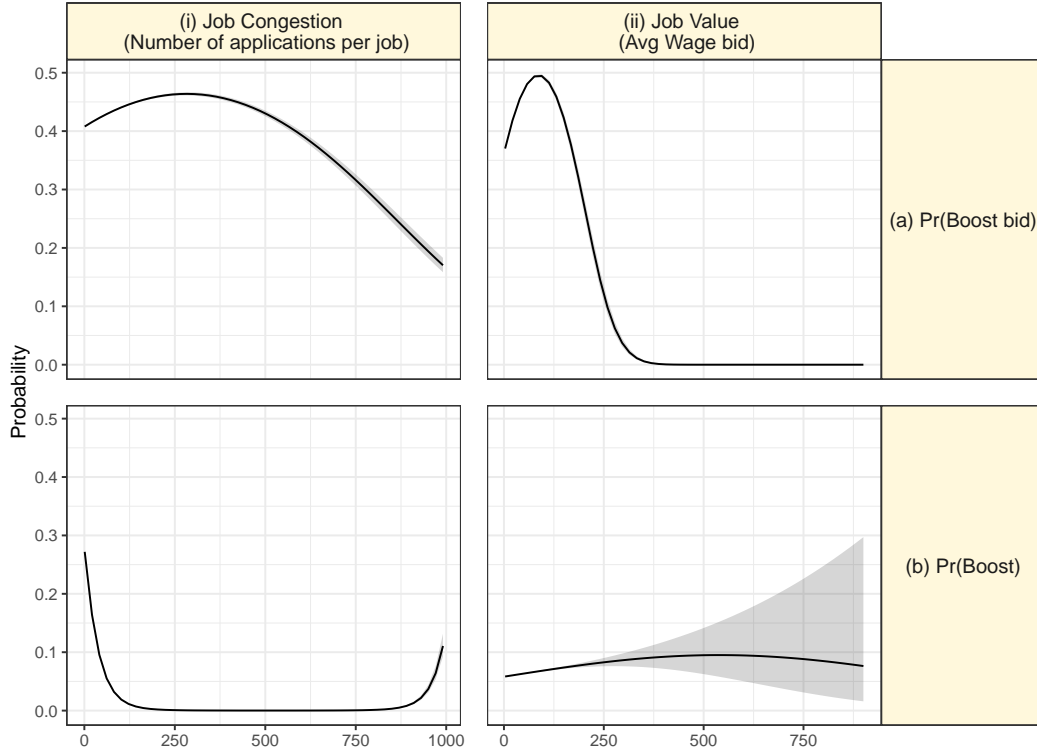


Notes: This figure shows the distribution of boosted applications across workers. Panel (a) plots the distribution of the number of boosted applications submitted by workers. Panel (b) restricts the sample to workers who submitted at least one boosted application, and plots the distribution of the fraction of a worker's applications worker that were boosted. Panel (c) further restricts the sample to workers that submitted at least five applications and one boosted application.

6.3 Differences between boosters and non-boosters

We next examine whether there are differences in the attributes of workers who use boosted applications compared to those who do not. Recall that in Section 5 we showed boosted applications are positively selected. There are two possible explanations for this result: (i) workers who boost their applications are of inherently higher quality, or (ii) workers who boost their applications put more effort into their applications. To understand whether there

Figure 4: Predicted probability of placing a boosted application bid



Notes: This figure shows the predicted probabilities of boosting as a function of job congestion and job value. Job congestion is shown in column (i), and is measured using the number of applications received for a job. Job value is shown in column (ii), and is measured using the average wage bid for a job. Row (a) shows the probability of a worker placing a bid for a boosted application slot. Row (b) shows the overall probability of a worker winning a boosted application slot.

are quality differences at the worker-level, we next compare the attributes of workers who had at least one boosted application (boosters) to workers who had no boosted applications.

Table 6 reports the mean differences in attributes between boosters and non-boosters, both in the pre-experimental and experimental periods. On average, boosters received more invitations from employers, applied to more jobs, asked for higher wages, and were hired more often than non-boosters. This was the case both in the pre-experimental and experimental periods. This suggests that boosters are inherently different than non-boosters, and the positive selection that we see cannot be explained by job-specific reasons.

Table 6: Mean differences in attributes between boosters and non-boosters

	Non-Boosters	Boosters	Diff. in Means	<i>p</i> -value
<i>Pre-experimental Period</i>				
num applications	7.6	18.3	10.7	<0.001
num invited applications	0.7	1.5	0.8	<0.001
num contracts formed	0.3	0.6	0.4	<0.001
avg hourly asking wage	20.7	24.6	4.0	<0.001
avg fixed asking wage	311.2	395.9	84.7	<0.001
<i>Experimental Period</i>				
num applications	4.8	15.1	10.3	<0.001
num invited applications	0.4	0.9	0.4	<0.001
num contracts formed	0.1	0.4	0.3	<0.001
avg hourly asking wage	23.0	25.4	2.5	<0.001
avg fixed asking wage	326.7	398.9	72.2	<0.001
Observations	362,327	117,128		

Notes: This table reports the mean pre-experimental and experimental attributes for workers who applied to at least one job posting during the experimental period. We define a worker as a “Booster” if they had at least one boosted application during the experimental period and as a “Non-Booster” otherwise. Invited workers are excluded from the analysis since invitees were not eligible to submit a boosted application. For each attribute, we report the mean value for each group, the difference in means, and the *p*-value of a two-sided t-test for the difference in means.

7 The effects of boosted applications on employers

We now turn our attention to the effects of boosted applications on employers. Because randomization took place at the employer-level, it is straightforward to obtain causal estimates of the effects of boosting on employers. We estimate the following specification:

$$y_k = \beta_0 + \beta_1 \text{ADON}_k + \beta_2 \text{ADNoDISCLOSURE}_k + \beta_3 \text{ADNoREC}_k + \epsilon_k, \quad (4)$$

where y_k is the outcome for employer k , each treatment indicator captures the treatment assignment for employer k (we set PLACEBO=0), and ϵ_k is the error term. The coefficient β_0 captures the average outcome for employers in the PLACEBO cell, and coefficients $\beta_1, \beta_2, \beta_3$ capture the treatment effect of boosted application on employer outcomes in each of the active treatment cells.

We consider the following employer outcomes: (1) the total number of job postings an employer makes during the experiment period, (2) the number of applications received per job posting, (3) the number of invited applications per job posting, (4) the average applicant rating per job posting, (5) the number of hires per job posting, (6) the average time to hire (in days) per job posting (7) the total expenditure per job posting, and (8) the average feedback from employer to worker per job posting. We report the estimates for all job postings during

the experiment period in Table 7a, and the estimates for the first job posting an employer makes once they are allocated to a treatment cell in Table 7b. The latter sample ensures that the estimates are not biased by any long-run indirect effects, e.g., employers changing their hiring behavior after being exposed to the treatment.

For all outcomes, we do not find any statistically significant effects of boosted applications on employers in either sample. This suggests that boosted applications do not affect employers outcomes in our setting—even though they affect worker outcomes. These results allay the concern that employers who encounter advertising for the first time on the platform might be eager to try hiring advertising workers, only to be later disappointed with their performance. Lastly, the fact that there was no net increase in hires corroborates the findings that employers substituted away from non-boosted applications in favor of boosted applications.

Table 7: Treatment effect estimates of boosted application on employer outcomes

(a) Outcomes using all job posts during the experiment period

Dependent Variables:	num posts (1)	num apps (2)	num invites (3)	avg app rtg (4)	num hires (5)	avg time to hire (6)	amt spent (7)	avg feedback (8)
PLACEBO (Intercept)	1.557*** (0.0091)	22.09*** (0.1440)	3.844*** (0.1516)	0.9329*** (0.0003)	0.6051*** (0.0180)	6.836*** (0.1457)	216.4*** (4.761)	8.641*** (0.0252)
AdON	0.0157 (0.0159)	-0.0207 (0.2599)	0.3140 (0.3634)	0.0007 (0.0005)	-0.0122 (0.0223)	0.1633 (0.2245)	0.6042 (6.723)	0.0239 (0.0407)
AdNoREC	0.0266 (0.0182)	0.4331 (0.2790)	0.6758 (0.4855)	-0.0005 (0.0005)	-0.0168 (0.0221)	-0.4185 [‡] (0.2380)	3.740 (8.233)	-0.0249 (0.0483)
AdNoDISCLOSURE	0.0010 (0.0181)	0.1266 (0.2991)	0.1277 (0.4577)	0.0008 (0.0005)	-0.0276 (0.0202)	-0.0758 (0.2677)	-9.685 (8.305)	0.0411 (0.0471)
<i>Fit statistics</i>								
Observations	106,788	167,322	167,322	155,943	167,322	69,046	167,322	45,017

Clustered (employer) standard-errors in parentheses
*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1*

(b) Outcomes using first job post after treatment allocation

Dependent Variables:	num posts (1)	num apps (2)	num invites (3)	avg app rtg (4)	num hires (5)	avg time to hire (6)	amt spent (7)	avg feedback (8)
PLACEBO (Intercept)	1.557*** (0.0091)	22.25*** (0.1455)	3.736*** (0.2661)	0.9334*** (0.0003)	0.5466*** (0.0076)	7.333*** (0.1630)	223.0*** (5.231)	8.605*** (0.0277)
AdON	0.0157 (0.0159)	-0.0826 (0.2055)	0.1108 (0.3759)	0.0005 (0.0004)	-0.0168 (0.0107)	-0.0085 (0.2312)	1.780 (7.388)	0.0006 (0.0391)
AdNoREC	0.0266 (0.0182)	0.1771 (0.2669)	0.4310 (0.4880)	-0.0003 (0.0006)	0.0162 (0.0139)	-0.4163 (0.2994)	0.4973 (9.593)	0.0350 (0.0505)
AdNoDISCLOSURE	0.0010 (0.0181)	0.1723 (0.2664)	0.5112 (0.4872)	0.0002 (0.0006)	-0.0049 (0.0139)	-0.3188 (0.3001)	-5.873 (9.576)	0.0653 (0.0507)
<i>Fit statistics</i>								
Observations	106,788	106,788	106,788	100,075	106,788	42,697	106,788	26,537

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1*

Notes: This table reports the OLS estimates of the effect on boosted application on employer outcomes. The independent variables are treatment indicators. The reported outcomes are (1) the total number of posts an employer makes during the experiment period, (2) the number of applications received per post, (3) the number of invited applications per post, (4) the average applicant rating per post, (5) the number of hires per post, (6) the average time to hire (in days) per post, (7) the total expenditure per post, and (8) the average feedback from employer to worker per post. In panel (a), Post-level outcomes are estimated using all posts the employer made during the experiment period. In panel (b), they are estimated using the first post the employer made after being allocated to a treatment cell.

8 Why does an advertising equilibrium exist in our setting?

We showed empirically that boosting—i.e., advertising—sends a positive signal to employers. For this signaling equilibrium to exist, there must be a separating equilibrium where high-quality workers find it more advantageous to advertise than low-quality workers.

Classical signaling models of advertising rely on “repeat purchases” as the mechanism that discourages low-quality sellers from advertising (Nelson, 1974; Milgrom and Roberts, 1986). In a labor market, the multi-stage nature of the hiring process plays the role of repeat interaction. The key logic that discourages low-quality workers from advertising at the same level as high-quality workers is that advertising affects the chances of getting an interview, but conditional on getting an interview, advertising does not affect the chances of getting hired. Intuitively, this means that even if low-quality workers advertise to increase their chances of getting an interview, the employer will likely find out their true quality during the interview stage and will be less likely to hire them. This discourages low-quality workers from advertising at the same level as high-quality workers, resulting in a signaling equilibrium.

Our data allows us to test directly this mechanism—whether advertising affects the chances of getting an interview but not the chances of getting hired conditional on an interview. Toward that end, we estimate the effect of boosted applications on these two outcomes. In Table 8, Columns (1) and (2) report estimates on the effect of boosting an application on the likelihood of getting an interview, and columns (3) and (4) on the likelihood of being hired conditional on getting an interview.

We find that boosted applications increase the likelihood of getting an interview, but do not affect the likelihood of a worker being hired conditional on getting an interview. These findings rationalize why advertising workers are positively selected, and why advertising is a positive signal in the hiring equilibrium.

9 Discussion and Conclusion

Online labor platforms have dramatically lowered the search and application costs for workers. Although these reduced costs are welcome, an unintended consequence is increased congestion—i.e., workers applying to jobs even when they are not a good fit. This congestion makes it more difficult for highly motivated workers to stand out from the crowd, and employers to find these workers. Algorithmic matching technologies mitigate some of this congestion by recommending matches that are more likely to be successful. However, there are limits to these matching technologies. Workers have private information about their fit and interest in jobs that neither the employers nor the platform can observe, which can be useful in the matching process. A costly signaling mechanism, such as advertising, can extract this private information.

Table 8: Treatment effect estimates of boosting on being interviewed and hired conditional on getting an interview

Dependent Variables:	Interviewed		Hired Interview	
	(1)	(2)	(3)	(4)
BOOST	0.0247*** (0.0009)	0.0108*** (0.0010)	0.0193*** (0.0035)	0.0114 [‡] (0.0059)
ADON \times BOOST	0.0376*** (0.0014)	0.0384*** (0.0015)	-0.0028 (0.0048)	-0.0043 (0.0076)
ADNoREC \times BOOST	0.0365*** (0.0019)	0.0372*** (0.0020)	-0.0002 (0.0060)	-0.0156 (0.0096)
ADNoDISCLOSURE \times BOOST	0.0259*** (0.0019)	0.0272*** (0.0020)	0.0031 (0.0062)	-0.0028 (0.0096)
<i>Fixed-effects</i>				
Job posting	✓	✓	✓	✓
Worker		✓		✓
<i>Fit statistics</i>				
Observations	3,403,094	3,403,094	270,608	270,608
<i>Clustered (employer) standard-errors in parentheses</i>				
<i>Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1</i>				

Notes: This table reports the OLS estimates of the effect of boosted application on: (1) the likelihood of getting an interview with job posting fixed effects; (2) the likelihood of getting an interview with job posting and worker fixed effects; (3) the likelihood of being hired conditional on getting an interview with job posting fixed effects; and (4) the likelihood of being hired conditional on getting an interview with job posting and worker fixed effects.

We reported on the results of a large-scale field experiment in an online labor market that can help us understand whether advertising can help workers get hired and whether the platform can use this information to improve the matching process. We showed that boosted applications are positively selected, and increase the likelihood of a worker being hired. This effect is driven by both the ranking and signaling effects of boosted applications. Interestingly, we found no significant effects of boosted applications on a wide range of employer outcomes.

Our findings have several implications for the design of online labor markets and labor market intermediaries such as LinkedIn. First, our results show that boosted applications can be a useful tool for workers to signal their interest and fit in a job to employers, increasing their likelihood of being hired. This mitigates concerns that advertising in a hiring setting may send a negative signal of desperation to employers. Rather, we find that employers take the signal at face value and view boosted applications as a net-positive signal. This can be helpful for highly-motivated workers to stand out from the crowd, and for employers to find these workers. It can also be especially helpful to new workers, who struggle to find jobs

without having built a reputation on the platform: we show that boosted applications are just as effective for less experienced workers as they are for experienced workers. As such, advertising offers an alternative mechanism to mitigate the “cold start” problem in labor markets.

The costly signaling mechanism in our setting took the form of sponsored advertising. However, labor platforms can use other costly signaling mechanisms similar to the ones widely used on dating platforms such as “Roses” on Hinge or “Superlikes” on Tinder. For example, LinkedIn could provide its users a limited number of “Premium” applications each month. And users could choose which job applications to use these “Premium” applications on.

The ranking of workers in employers’ application tracking system continues to play a crucial role in the matching process, even in the presence of advertising. This result is well established in e-commerce settings but less understood in labor markets. In contrast to an e-commerce market, engaging with all sellers (workers) is beneficial to the buyer (employer) as it provides an opportunity to negotiate and obtain a better price. Yet, we find that employers still rely on the “organic” ranking of workers. This suggests that even in the absence of advertising, labor platforms could use information about sellers’ propensity to advertise or past advertising behavior to improve their ranking algorithms as suggested by [Long et al. \(2022\)](#).

We did not find any significant differences in employer outcomes due to boosting. By providing the employers and the platform with more information about the fit/interest of workers to jobs, boosted applications should, in theory, increase matching efficiency. One reason why we do not see significant increases in the final match success between employers and workers might be due to where in the hiring process the increase in matching efficiency is happening. Because hiring is a multi-stage process and advertising only affects the first stage (screening), all the information there is to be extracted from advertising is likely being extracted in the first stage. Consequently, advertising is likely increasing the matching efficiency in the first stage, either by increasing the quality of the employers’ shortlists or by reducing the amount of time employers spend reviewing and screening workers, neither of which are directly observable in our data. Future research could explore the mechanisms through which advertising affects the matching process in more detail, either with more detailed data or through additional field experiments.

References

- Abhishek, V., Jerath, K., and Sharma, S. (2022). The impact of retail media on online marketplaces: Insights from a field experiment.
- Agrawal, A., Horton, J. J., Lacetera, N., and Lyons, E. (2015). Digitization and the contract labor market: A research agenda. In *Economic analysis of the digital economy*, pages 219–250. University of Chicago Press.
- Autor, D. H. (2001). Wiring the labor market. *The Journal of Economic Perspectives*, 15(1):25–40.
- Benson, A., Sojourner, A., and Umyarov, A. (2019). Can reputation discipline the gig economy? experimental evidence from an online labor market. *Management Science*.
- Blake, T., Nosko, C., and Tadelis, S. (2015). Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica*, 83(1):155–174.
- Coviello, L., Gneezy, U., and Goette, L. (2017). A large-scale field experiment to evaluate the effectiveness of paid search advertising.
- Cowgill, B. (2019). Bias and productivity in humans and machines.
- Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., and Zamora, P. (2013). Do labor market policies have displacement effects? evidence from a clustered randomized experiment. *The Quarterly Journal of Economics*, 128(2):531–580.
- Dai, W., Kim, H., and Luca, M. (2023). Frontiers: Which firms gain from digital advertising? evidence from a field experiment. *Marketing Science*, 42(3):429–439.
- Filippas, A., Horton, J., Parasurama, P., and Urraca, D. (2024). Costly capacity signaling increases matching efficiency: Evidence from a field experiment. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 414–415.
- Filippas, A., Horton, J. J., and Golden, J. (2021). Reputation inflation. *Marketing Science*.
- Filippas, A., Horton, J. J., and Zeckhauser, R. J. (2020). Owning, using, and renting: Some simple economics of the “sharing economy”. *Management Science*, 66(9):4152–4172.
- Geyik, S. C., Ambler, S., and Kenthapadi, K. (2019). Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. *arXiv:1905.01989 [cs]*.
- Golden, J. and Horton, J. J. (2021). The effects of search advertising on competitors: An experiment before a merger. *Management Science*, 67(1):342–362.

- Horton, J. J. (2010). Online labor markets. *Internet and Network Economics*, pages 515–522.
- Horton, J. J. (2017). The effects of algorithmic labor market recommendations: Evidence from a field experiment. *Journal of Labor Economics*, 35(2):345–385.
- Horton, J. J. (2019). Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution. *Management Science*.
- Horton, J. J., Kerr, W. R., and Stanton, C. (2017). Digital labor markets and global talent flows. Technical report, National Bureau of Economic Research.
- Joo, M., Shi, J., and Abhishek, V. (2024). Do sellers benefit from sponsored product listings? evidence from an online marketplace. *Marketing Science*.
- Kihlstrom, R. E. and Riordan, M. H. (1984). Advertising as a signal. *Journal of Political Economy*, 92(3):427–450.
- Kokkodis, M. and Ipeirotis, P. G. (2023). The good, the bad, and the unhirable: Recommending job applicants in online labor markets. *Management Science*, 69(11):6969–6987.
- Kroft, K., Lange, F., and Notowidigdo, M. J. (2013). Duration dependence and labor market conditions: Evidence from a field experiment. *The Quarterly journal of economics*, 128(3):1123–1167.
- Li, D., Raymond, L. R., and Bergman, P. (2020). Hiring as exploration. Technical Report w27736, National Bureau of Economic Research.
- Long, F., Jerath, K., and Sarvary, M. (2022). Designing an online retail marketplace: Leveraging information from sponsored advertising. *Marketing Science*, 41(1):115–138.
- Milgrom, P. and Roberts, J. (1986). Price and advertising signals of product quality. *Journal of Political Economy*, 94(4):796–821.
- Moshary, S. (2024). Does sponsored search advertising augment organic search? evidence from an e-commerce platform. *Management Science*.
- Nelson, P. (1974). Advertising as information. *Journal of Political Economy*, 82(4):729–754.
- Pallais, A. (2013). Inefficient hiring in entry-level labor markets. *American Economic Review*.
- Ramanath, R., Inan, H., Polatkan, G., Hu, B., Guo, Q., Ozcaglar, C., Wu, X., Kenthapadi, K., and Geyik, S. C. (2018). Towards deep and representation learning for talent search at linkedin. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, CIKM ’18, pages 2253–2261, New York, NY, USA. Association for Computing Machinery.

- Sahni, N. S. (2015). Effect of temporal spacing between advertising exposures: Evidence from online field experiments. *Quantitative Marketing and Economics*, 13(3):203–247.
- Sahni, N. S. and Nair, H. S. (2020a). Does advertising serve as a signal? evidence from a field experiment in mobile search. *The Review of Economic Studies*, 87(3):1529–1564.
- Sahni, N. S. and Nair, H. S. (2020b). Sponsorship disclosure and consumer deception: Experimental evidence from native advertising in mobile search. *Marketing Science*, 39(1):5–32.
- Stanton, C. T. and Thomas, C. (2016). Landing the first job: The value of intermediaries in online hiring. *The Review of Economic Studies*, 83(2):810–854.
- Varian, H. R. (2010). Computer mediated transactions. *American Economic Review*, 100(2):1–10.

A More details on the experiment

A.1 Internal validity


One way to assess whether the randomized assignment was performed correctly is to try to detect systematic differences in observable pre-treatment characteristics between employers assigned to the control and the treatment groups. In Table 9, we perform a series of two-sided t-tests for various job attributes. We find no evidence of systematic differences between these job-level characteristics. In addition, Figure 5 plots the allocation of employers to the treatment and control cells over time.

Table 9: Balance test table

	ADON mean $\overline{X}_{\text{ADON}}$	PLACEBO mean $\overline{X}_{\text{PLACEBO}}$	p-value
<i>Post Characteristics</i>			
number of posts	1.79	1.86	0.18
amount spent	69.34	71.66	0.404
invites sent	9.17	9.17	0.589
fill probability	1.4	1.4	0.983
<i>Observation counts</i>	37,417	37,616	0.468

Notes: This table reports averages and p-values of two-sided t-tests for various pre-treatment observables, for workers assigned to the ADON and PLACEBO experimental groups. Each outcome is an employer-level aggregate between June 8, 2021 and September 8, 2021. The reported outcomes are (i) the number of posts the employer created, (i) the amount an employer spent, (ii) the number of invites the employer sent to workers, and (iv) the number of hires that the employer made on the platform. Performing the same tests for other experimental arms yields no evidence of systematic differences.

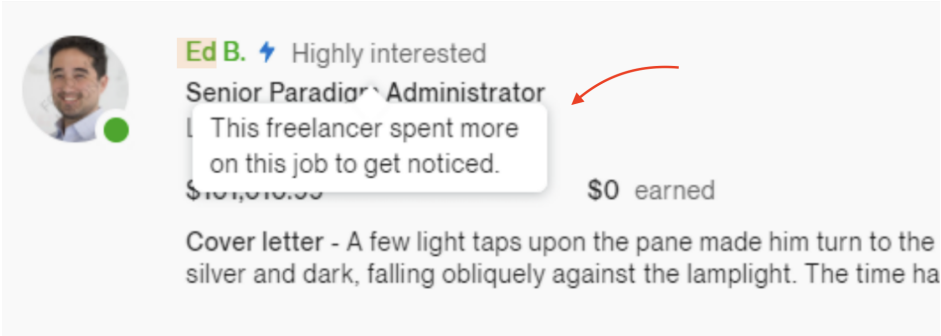
Figure 5: Employers allocated to the experimental groups over time.



Notes: This figure plots the number of employers allocated to the treatment groups each day of the allocation period. The allocation period began on September 8, 2021 and ended on October 13, 2021.

A.2 User Interfaces

Figure 6: Employer view of a boosted application.



Notes: This figure shows the view of a boosted application of an employer that was assigned the ADON treatment. The boosted application is pinned to the top of the list of applications, and is marked with a “Highly interested” label.

Figure 7: Example worker view of the auction interface.

Boost your application (optional)

Bid for one of 3 boosted application spaces at the top of the employer's ATS

How bidding works ▼

Slot	Bid
1st place	20 Coins. 1 hour ago
2nd place	15 Coins. 1 hour ago
3rd place	10 Coins. 30 minutes ago

+ Set a Bid

Notes: This figure shows the view of the auction interface for a worker after they have submitted an application.

B Additional analyses

Table 10: Treatment effect estimates of boosted application on the likelihood of being hired (with > 0 earnings)

	(1)	(2)	(3)	(4)
PLACEBO	0.0138*** (0.0008)			
AdON	-0.0019* (0.0009)	-0.0018** (0.0006)		
AdNoREC	-0.0023** (0.0009)	-0.0020** (0.0007)		
AdNoDISCLOSURE	-0.0017 [‡] (0.0009)	-0.0018** (0.0007)		
BOOST	0.0132*** (0.0007)	0.0087*** (0.0007)	0.0079*** (0.0005)	0.0030*** (0.0006)
AdON \times BOOST	0.0106*** (0.0009)	0.0108*** (0.0008)	0.0111*** (0.0008)	0.0115*** (0.0008)
AdNoREC \times BOOST	0.0111*** (0.0011)	0.0109*** (0.0011)	0.0118*** (0.0010)	0.0120*** (0.0011)
AdNoDISCLOSURE \times BOOST	0.0085*** (0.0011)	0.0084*** (0.0011)	0.0090*** (0.0010)	0.0093*** (0.0010)
<i>Fixed-effects</i>				
Worker		✓		✓
Job posting			✓	✓
<i>Fit statistics</i>				
Observations	3,403,094	3,403,094	3,403,094	3,403,094

Clustered (employer) standard-errors in parentheses

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1*

Notes: This table reports the OLS estimates of the effect of boosted application on the likelihood of being hired. We use a more restrictive outcome variable, where we set the outcome to 1 if the worker was hired and earned more than \$0 from the job within 60 days, and 0 otherwise. The independent variables are: (i) a binary indicator for whether an application is a boosted application, (ii) treatment indicators for the experimental assignment of the employer, and (iii) interactions between the boosted application and treatment indicators. Column (1) of the table reports the OLS estimates, Column (2) includes only worker fixed effects, Column (3) includes only job posting fixed effects, and Column (4) includes both worker and job posting fixed effects. Invited applications are included in this.

Table 11: Treatment effect estimates of boosted application on the likelihood of being hired (includes invited applications)

	(1)	(2)	(3)	(4)
PLACEBO	0.0249*** (0.0008)			
AdON	-0.0019 [‡] (0.0010)	-0.0020** (0.0006)		
AdNoREC	-0.0028** (0.0009)	-0.0024*** (0.0007)		
AdNoDISCLOSURE	-0.0022* (0.0009)	-0.0023*** (0.0007)		
BOOST	0.0053*** (0.0007)	5.66×10^{-5} (0.0007)	-0.0012* (0.0005)	-0.0062*** (0.0006)
AdON \times BOOST	0.0110*** (0.0010)	0.0115*** (0.0009)	0.0119*** (0.0008)	0.0124*** (0.0009)
AdNoREC \times BOOST	0.0119*** (0.0012)	0.0118*** (0.0012)	0.0120*** (0.0011)	0.0126*** (0.0011)
AdNoDISCLOSURE \times BOOST	0.0093*** (0.0012)	0.0091*** (0.0012)	0.0094*** (0.0011)	0.0098*** (0.0011)
<i>Fixed-effects</i>				
Worker		✓		✓
Job posting			✓	✓
<i>Fit statistics</i>				
Observations	3,665,555	3,665,555	3,665,555	3,665,555

Clustered (employer) standard-errors in parentheses

*Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, ‡: 0.1*

Notes: This table reports the OLS estimates of the effect of boosted application on the likelihood of being hired. The independent variables are: (i) a binary indicator for whether an application is a boosted application, (ii) treatment indicators for the experimental assignment of the employer, and (iii) interactions between the boosted application and treatment indicators. Column (1) of the table reports the OLS estimates, Column (2) includes only worker fixed effects, Column (3) includes only job posting fixed effects, and Column (4) includes both worker and job posting fixed effects. Invited workers are excluded from the analysis, since invitees were not eligible to submit a boosted application