

Decomposing the Signaling and Ranking Effects of Sponsored Advertising: Experimental Evidence from an Online Labor Market

Anonymized for review

April 22, 2024

Abstract

We study the introduction of advertising in an online labor market where applicants could bid to “boost” their applications. Applicants with the highest bids received a prominent position at the top of the employer’s application list. The platform randomly varied whether the employer was exposed to boosted applications, and within the exposure group, whether the employer was shown a disclosure that the application was boosted. We find that boosted applications are positively selected and that boosting an application increases the likelihood of an applicant being hired by 40.8%. The experimental design allows us to estimate that 79.8% of this increase was due to the ranking effect—boosted applications ranking higher—and 20.2% was due to the signaling effect—the disclosure that the application is boosted.

1 Introduction

Several online platforms allow sellers to pay to have their listings shown more prominently alongside organic results. This “sponsored advertising” was pioneered by internet search engines and has since become a core feature in online platforms.¹ An emerging experimental literature has studied the effects of sponsored advertising in online search engines (Blake et al., 2015; Coviello et al., 2017), e-commerce platforms (Moshary, 2021; Abhishek et al., 2022; Joo et al., 2024), and restaurant platforms (Sahni and Nair, 2020a; Dai et al., 2023). This paper contributes to this literature by providing the first empirical study of sponsored advertising in an online labor market.

Labor markets have distinct characteristics that make predicting the effects of sponsored advertising difficult. Advertising could make a seller more prominent in the rankings and hence increase click-through and conversion rates (Animesh et al., 2011; Rutz et al., 2012; Narayanan and Kalyanam, 2015; Jeziorski and Moorthy, 2018). However, compared to an e-commerce setting, prominence may matter less in labor markets, as the employer’s consideration set is likely smaller, and engaging with all applicants can be informative in the employer’s wage negotiations.²

Advertising could also function as a signal.³ In our setting, advertising is highly targeted—workers choose exactly which jobs to advertise to. Since workers have private information about their fit and interest in a job, advertising could send a positive signal to the employer. At the same time, employers may expect high-quality workers to already have enough work to not have to advertise, sending a negative signal of desperation. This concern is exacerbated in our setting because workers have tight capacity constraints—a worker could only take on so many jobs at a time. Furthermore, because hiring is typically a high-stakes decision, an employer generally spends a substantial amount of time and resources evaluating a worker. Whether advertising could meaningfully affect hiring decisions at all is unclear.

We study the effects of sponsored advertising through an experiment conducted by a large online labor market. In this labor market, employers post job descriptions for jobs that can be done remotely (e.g., computer programming, graphic design, writing). Workers search for and apply to jobs. For a given job, the applications are shown to the employer in a ranked order determined by the platform’s proprietary algorithm. From this list, the employer can

¹Sponsored advertising is a substantial revenue stream for online platforms. For example, Amazon reported a \$47B revenue from advertising in 2023. See <https://www.statista.com/statistics/259814/amazons-worldwide-advertising-revenue-development>.

²In our context, a job receives about 20 applications on average.

³Signaling models of advertising posit that advertising in and of itself can provide an informative signal about the quality of the seller. In the signaling equilibrium of advertising, high-quality sellers have more to gain from signaling their quality compared to low-quality sellers and are therefore more incentivized to advertise. As a result, high-quality sellers advertise and low-quality sellers do not. This provides useful information to buyers, who can use advertising as a signal to discern high-quality from low-quality sellers (Nelson, 1974; Kihlstrom and Riordan, 1984; Milgrom and Roberts, 1986).

shortlist a subset of the candidates, further assess them via interviews, negotiate over wages, and make hiring decisions.

During the experiment, all workers became eligible to bid to advertise through “boosted” applications when applying for a job. The platform randomly varied whether and how employers were exposed to boosted applications. Employers assigned to the control group (PLACEBO) saw no change to their application list and were shown the organic results. Treated employers 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 did not see an additional algorithmically determined recommendation label—which employers in all other groups saw. The design of the experiment allows us to disentangle the distinct but often concomitant effects of sponsored advertising—the self-selection effect into advertising, the ranking effect of advertising, the signaling effect of advertising, and the interaction of advertising effects with other algorithmic recommendation labels.⁴

Our results show positive self-selection into advertising. Based on observational data, we find that workers who advertised were more likely to be sought out by employers and more likely to be hired both in the experimental and pre-experimental periods. To precisely estimate the self-selection effect into advertising, we compare the outcomes of workers who chose to advertise to those who did not within the PLACEBO. This comparison allows us to observe the outcomes of boosted applications as if they had not been boosted since boosting had no effect on the employer’s application list. We find boosted applications in the PLACEBO cell were 80.5% (6 pp) more likely to be interviewed and 101% (1.5pp) more likely to be hired compared to non-boosted applications in the PLACEBO cell. These effects are present even with worker fixed effects, indicating 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.

Second, we find that boosting an application increases the likelihood of a worker getting interviewed and hired. To estimate the causal effects 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 between cells differences-out the self-selection effect, and isolates the causal effects of boosting. We find that boosting an application increases the likelihood of a worker getting an interview by 28.7% (3.8 pp), and getting hired by 40.8% (1.2 pp).

Third, we decompose the total effects of boosting into the ranking and signaling effects.

⁴This type of non-ad-based recommendation label is also common in other platforms. For example, Amazon has recommendation labels such as “Overall Pick” and “Best Seller” next to some recommended listings in addition to sponsored listings.

To do so, we compare the difference in outcomes in the ADON cell (which captures both the ranking and signaling effects) cell to the ADNODISCLOSURE cell (which captures just the ranking effect). 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%. The ranking effect accounts for 79.8% of the total effects of boosting, and the signaling effect accounts for 20.2% of the total effects.

Fourth, we find that the effects of boosting are similar both in the presence and absence of the algorithmically determined recommendation label. We compare the outcomes in the ADON cell (where employers see an additional algorithmic recommendation label) to the ADNOREC cell (where employers do not see the additional recommendation 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%. Contrary to the concerns that sponsored advertising could compete with other algorithmic recommendation labels lowering the effectiveness of sponsored advertising, we find no evidence of this in our context.

Turning to the effects of boosted applications on employers, we consider a number outcomes: (1) number of jobs posted, (2) number of applications received, (3) number of invited applications, (4) average rating of the applicant pool, (5) number of workers hired, (6) average time to hire, (7) amount of money spent on the job, and (8) average feedback ratings. We do not find any statistically significant effects of boosted applications on any of the above employer outcomes.

We rationalize these results with an illustrative model of hiring and accompanying evidence that boosted applications increase the likelihood of a worker being interviewed, but not the likelihood of a worker being hired conditional on being interviewed. This implies 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 in the first place. This explains why we observe positive self-selection into advertising. Because there is positive selection into advertising, employers can use the information about whether a worker advertised to discern high-quality from low-quality workers. This explains why the disclosure that the worker advertised increases the likelihood of a worker getting an interview, and thus getting hired.

This paper makes a number of contributions to the sponsored advertising literature. First, we provide the first empirical evidence on the effects of sponsored advertising in a labor market setting. Second, we provide evidence of positive self-selection into advertising. Whereas previous studies have provided suggestive evidence of positive selection (e.g., [Moshary \(2021\)](#); [Sahni and Nair \(2020a\)](#)), our experimental design with the PLACEBO allows us to isolate and precisely estimate the self-selection effect. Third, we disentangle the ranking effect from the

signaling effect of sponsored advertising. In a typical sponsored ad setting, these two effects are usually confounded—advertisements are displayed more prominently and are labeled as such. [Sahni and Nair \(2020a\)](#) isolate just the signaling effect of advertising by randomly varying disclosure to the buyer, but their design does not allow them to measure the ranking effect of advertising. We measure both the ranking and signaling effects of advertising, which allows us to estimate the relative importance of these two effects. Lastly, our study provides evidence that sponsored advertising, at least in our context, does not compete with other algorithmically determined recommendation labels.

The rest of the paper is organized as follows. [Section 2](#) reviews the related literature. [Section 3](#) describes the empirical context for our study. [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 workers’ use of boosted applications. [Section 7](#) reports the effects of boosted applications on employers. [Section 8](#) provides an illustrative model to rationalize these results. Lastly, we discuss the implications of our results and conclude in [section 9](#).

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. [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, limiting the generalizability of the results.

Recent platform-based studies have overcome this issue by randomly varying (i) access to advertising for sellers or (ii) 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 the effects of 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. [Sahni and Nair \(2020a\)](#) study the effect of advertising in another 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, highlighting the signaling role of advertising (see also [Sahni \(2015\)](#); [Sahni and Nair \(2020b\)](#)). [Moshary \(2021\)](#) study the effects of sponsored advertising in an e-commerce platform and find that sponsored advertising increased the likelihood of a product being purchased. In contrast, [Joo et al. \(2024\)](#) find that ad disclosure in an e-commerce platform *decreased* click-through and conversion rates, whereas [Abhishek et al. \(2022\)](#) find that additional advertising increased purchases for some categories but not others.

One reason for these divergent results could be that whether advertising has a positive or negative effect on sellers is dependent on whether advertising sellers are positively or adversely selected ([Nelson, 1974](#); [Kihlstrom and Riordan, 1984](#); [Milgrom and Roberts, 1986](#)). If there is a separating equilibrium where high-quality sellers advertise at a higher level than low-quality sellers, then advertising will have a positive effect on sellers; and vice versa if low-quality sellers advertise at a higher level.

Experimental evidence on the effects of sponsored advertising on buyers remains limited. [Sahni and Nair \(2020a\)](#) find some suggestive evidence that ad disclosure improved buyer outcomes in that they were more likely to call higher-rated restaurants. [Sahni and Zhang \(2023\)](#) study the effects of search advertising on consumer preferences and find that users exposed to lower levels of advertising in a search engine reduced their search engine usage.

Our study is related to these streams of literature in that we study the effects of sponsored advertising, both on the sellers (workers) and buyers (employers). But unlike the above studies, we study it in a labor market setting, where there are substantive differences in the role that advertising could play.

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, better matches, and more diverse hires. 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 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. In this paper, we study how incorporating this private information (through advertising) 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 may invite desirable workers to apply for their jobs, and 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 (Horton, 2017, 2019), barriers to entry (Pallais, 2013; Stanton and Thomas, 2016), and information asymmetries (Benson et al., 2019; Filippas et al., 2021).

3.1 The status-quo job application process

The platform secures revenues by taxing the wage bill from formed matches. However, as in all labor markets, workers likely do not fully internalize the benefits of their job search efforts (Diamond, 1982; Hosios, 1990; Mortensen and Pissarides, 1994). Consequently, it is in the platform’s interest to incentivize job search—to the extent that this leads to more high-quality matches. However, encouraging job search alone does not necessarily lead to better matches, for reasons including over-application and crowd-out (Crépon et al., 2013; Horton and Vasserman, 2020).

Over-application is considered enough of a problem that the platform sets a price for applying to jobs. To encourage greater job search, the platform could decrease the price of applications, but at the same time, it would discourage job-seekers from being selective in their applications. The platform has to balance these two countervailing effects when setting the application price—a difficult task to optimize. This motivated the introduction of a new signaling mechanism—boosted applications—which encourages greater job search by reducing the cost of application but still allowing job-seekers to send a costly signal of their interest in a job.

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—cost of application—is determined by the platform using a proprietary formula that takes into account only job-

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.

specific attributes, such as the anticipated job duration and earnings. Before the experiment was conducted, the cost for applying to a job ranged from 2 to 6 coins. During the experiment, the cost for applying was lowered to 1 coin. Employers may also invite workers to apply to jobs. A job application following an employer invite uses up no coins (Filippas et al., 2022).

Table 1: Pre-experimental summary statistics

	Mean	Median	SD	Min	Max
<i>Employer/Job statistics</i>					
number of posts per employer	4.49	2	10.11	1	1,261
number of apps per job	20.10	13	28.65	0	2,243
number of invited apps per job	5.19	1	58.76	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	256.91	0	1362.00	0	251,551
<i>Worker statistics</i>					
number of applications	5.08	2	17.62	1	3,399
number of invites received	0.47	0	2.28	0	183
number of contracts formed	0.17	0	0.75	0	81
number of contracts per application	0.03	0	0.13	0	1
average hourly asking wage	22.16	15	34.75	0	999
average fixed asking wage	790.95	75	12036.29	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.

We provide some summary statistics on this status quo job search and matching in Table 1. Our sample includes jobs posted in the pre-experimental period. A job received 20.1 applications on average, of which 5.19 were invited applications. The average job fill rate was 0.53, meaning that about half of the jobs were filled. On the worker side, workers submitted 5.08 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

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.

Boosted applications introduced experimental variation in the employers’ ATS. Employers were allocated randomly to one of four treatment groups upon posting a job during the experimental period. The ATS version each employer saw depended fully on their treatment assignment. These changes are summarized in Table 2, and described below.

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	✓	✓		✓

- ADON: Boosted applications were displayed at the top of the employer’s ATS, and included a “Highly Interested” label next to the name. Hovering over the label revealed that the worker paid more 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 did not include the algorithmically determined “Best Match” label.
- 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.

Workers did not know which treatment group each job belonged to. This experimental design 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. Second, we can estimate the causal effect of the boosted applications on employer and worker outcomes, by comparing the ADON and PLACEBO cells. Third, we can separate the ranking effect from the signaling effect (“Highly Interested” label) of a boosted application, by comparing the ADON and 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 ADNOREC cells.

4.2 The boosted application auction format

There were three boosted application slots. To compete for these slots, workers could bid using coins, 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. 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.

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, 37,616 (35.22%) were allocated to ADON, 37,417 (35.04%) were allocated to PLACEBO, 15,838 (14.83%) were allocated to ADNOREC, 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 with the same experimental group throughout the experiment. The platform did not inform the employers that they received different treatments. The experimental groups are seemingly well-balanced.

In Appendix [A.1](#), we report two-sided t-tests for various employer-level attributes, and plot allocations over time.

5 Effects of boosted applications on workers

5.1 Estimation strategy

We estimate the treatment effect of boosted applications 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 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 differences in 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 outcome between 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 effect into 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 capture the causal effect of boosted applications on workers' outcomes in each of the three active treatment groups.

5.2 Treatment effects of boosted applications

Our outcomes of interest are: (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 each outcome in Table 3 and Table 4, respectively. In each table, Column (1) reports regression estimates for the specification of Equation 1. Column (2) adds worker fixed effects, Column (3) adds job posting fixed effects, and Column (4) adds both worker and job posting fixed effects.

A few things are worth noting from the results. First, In Column (1), the intercept (PLACEBO) is the baseline probability of a non-boosted application being interviewed (Table 3) and hired (Table 4) in the PLACEBO cell. Second, BOOST is positive for both outcomes, meaning that boosted applications are more likely to be interviewed and hired compared to non-boosted applications in the PLACEBO cell, which in turn means that boosted applications are positively selected. Third, ADON, ADNODISCLOSURE, and ADNOREC are all negative, meaning that non-boosted applications are less likely to be interviewed and hired in the three active treatment cells compared to the PLACEBO cell. This indicates a crowd-out effect of boosted applications on non-boosted applications. Lastly, the interaction terms are all positive, indicating that boosting an application causally increases the likelihood of a worker being interviewed and hired.

Column (2) adds worker fixed effects, which exploits within worker variation in boosted applications. That BOOST is positive with worker fixed effects indicates 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, boosting is specific and targeted. Column (3) adds job posting fixed effects, which accounts for heterogeneity in boosted applications and outcomes across jobs. For example, it accounts for such scenarios as when workers boost their applications at a higher rate for jobs that are more congested and thus have a lower baseline probability of being hired, which could bias the treatment effect estimates. Column (4) adds both worker and job posting fixed effects. This is our preferred specification, as it most closely answers the question: for a given job, what is the effect of a worker boosting their application? Based on this specification, we find that boosting an application increases the likelihood of a worker getting an interview by 28.7%⁶ (3.8 pp), and getting hired by

⁶This is relative to the average probability of a boosted application being interviewed in the PLACEBO.

Table 3: Treatment effect estimates of boosted application on the likelihood of being interviewed

	(1)	(2)	(3)	(4)
PLACEBO	0.0741*** (0.0022)			
ADON	-0.0047‡ (0.0024)	-0.0044* (0.0017)		
ADNoREC	-0.0052‡ (0.0028)	-0.0045* (0.0021)		
ADNoDISCLOSURE	-0.0054‡ (0.0028)	-0.0050* (0.0021)		
BOOST	0.0597*** (0.0017)	0.0507*** (0.0015)	0.0247*** (0.0009)	0.0108*** (0.0010)
ADON × BOOST	0.0352*** (0.0021)	0.0350*** (0.0019)	0.0376*** (0.0014)	0.0384*** (0.0015)
ADNoREC × BOOST	0.0334*** (0.0026)	0.0330*** (0.0025)	0.0365*** (0.0019)	0.0372*** (0.0020)
ADNoDISCLOSURE × BOOST	0.0250*** (0.0027)	0.0246*** (0.0025)	0.0259*** (0.0019)	0.0272*** (0.0020)
<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 interviewed. 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.

40.8%⁷ (1.2 pp).

To better visualize these results, we plot the estimated effects of boosting as a percent increase from the baseline hire rate in Figure 1.

Appendix B reports additional robustness checks, where we use a more restrictive outcome variable (“hired and earned > 0 from the job”), and where we include invited workers in the

⁷This is relative to the the average probability of a boosted application being hired in the PLACEBO.

Table 4: Treatment effect estimates of boosted application on the likelihood of being hired

	(1)	(2)	(3)	(4)
PLACEBO	0.0150*** (0.0008)			
ADON	-0.0021* (0.0009)	-0.0020** (0.0006)		
ADNoREC	-0.0025** (0.0009)	-0.0022*** (0.0007)		
ADNoDISCLOSURE	-0.0019* (0.0009)	-0.0020** (0.0007)		
BOOST	0.0152*** (0.0007)	0.0103*** (0.0007)	0.0089*** (0.0005)	0.0035*** (0.0006)
ADON \times BOOST	0.0112*** (0.0009)	0.0114*** (0.0009)	0.0119*** (0.0008)	0.0123*** (0.0008)
ADNoREC \times BOOST	0.0115*** (0.0011)	0.0112*** (0.0011)	0.0124*** (0.0010)	0.0125*** (0.0011)
ADNoDISCLOSURE \times BOOST	0.0089*** (0.0011)	0.0088*** (0.0011)	0.0094*** (0.0010)	0.0098*** (0.0011)
<i>Fixed-effects</i>				
Worker		✓		✓
Job posting			✓	✓
<i>Fit statistics</i>				
Observations	3,403,094	3,403,094	3,403,094	3,403,094
<i>Clustered (employer) standard-errors in parentheses</i>				
<i>Signif. Codes: ***: 0.001, **: 0.01, *: 0.05, †: 0.1</i>				

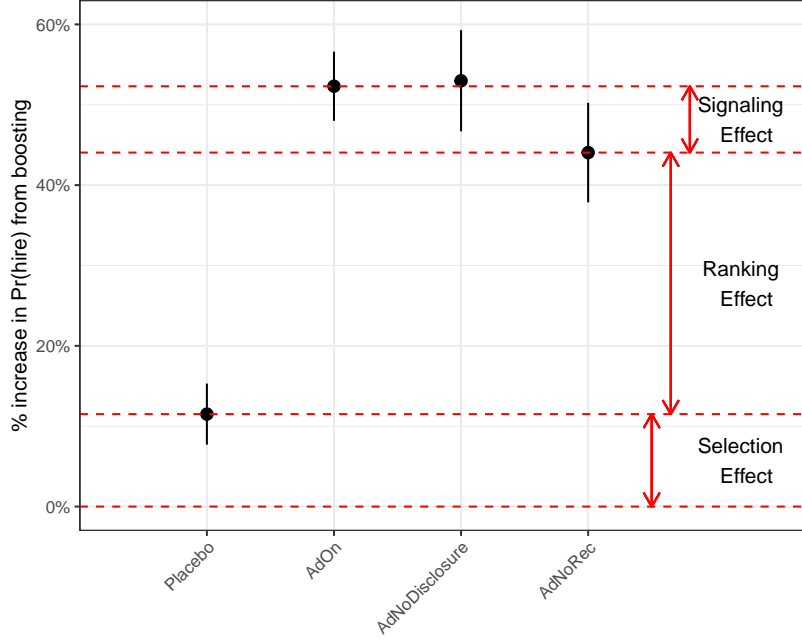
Notes: This table reports the OLS estimates of the effect of boosted application on the likelihood of being hired. The independent variables are: (i) a binary indicator for whether an application is a boosted application, (ii) treatment indicators for the experimental assignment of the employer, and (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.

analysis. The results remain highly similar.

5.3 Decomposing the ranking and signaling effects of boosted applications

We decompose the total effects of boosted applications into two components: the signaling effect stemming from the “highly interested” label of a boosted application, and the ranking effect stemming from the change in the position of a boosted application in the employer’s application list. We do this by comparing the outcomes of applications in the ADNoDISCLO-

Figure 1: Effects of boosted application on the likelihood of being hired



Notes: This figure plots the estimated effects of boosting as a % increase from the baseline probability of being hired. The estimates are based on the Worker + Job FE model.

SURE cell to the outcomes of applications in the ADON cell. Recall that in the ADON cell, employers could see both the ranking change and the “highly interested” label for boosted applications. But in the ADNODISCLOSURE cell, employers could see only see the ranking change. The difference in treatment effects between these cells thus provides the causal effect of the “highly interested” label on the likelihood of a worker being hired—i.e., it isolates the signaling effect. Formally, in model specification 1, β_5 captures the total effect of boosted application on the likelihood of a worker being hired, β_6 captures the ranking change effect of boosted application, and $\beta_5 - \beta_6$ captures the signaling effect.

As reported in Table 4, the treatment effect of boosted application on the likelihood of a worker being hired is 1.2 percentage points – attributable to both the ranking and signaling effects. In the ADNODISCLOSURE cell, the treatment effect is 1 percentage points – attributable only to the ranking effect. The difference between these two estimates (0.2 percentage points) is the signaling effect of boosted application. A linear hypothesis test ($\beta_5 - \beta_6 = 0$) confirms that the signaling effect is statistically significant (p -value = 0.027). In relative terms, the ranking effect accounts for 79.8% of the total effects, and the signaling effect accounts for 20.2% of the total effects.

5.4 Treatment effects in the absence of algorithmically determined “Best Match” label


We compare the relative effects of the boosted applications in the presence and absence of the algorithmically determined “Best Match” label by comparing the ADON and ADNOREC cells. Specifically, we test 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., $\beta_5 - \beta_7 = 0$, and find that they are not statistically different (p -value = 0.856). This suggests that the effects of boosted applications on the likelihood of being hired dominate any effects of the algorithmically determined “Best Match” label.

5.5 Heterogeneity analysis

We perform heterogeneity analysis to understand whether the effect of boosted application on the probability of being hired varies by job category. Each job is classified into one of 12 job category types. We sub-sample the data by job category type, estimate the above Worker + Job posting fixed effects specification for each job type separately, and report the ADON vs. PLACEBO treatment effect estimates in Figure 2. The left panel shows the nominal treatment effect estimates. Since the baseline hire rate could vary across job category types, we report the relative treatment effect estimates as a percentage increase from the baseline hire rate in the right panel. While there is some heterogeneity across job types, the mean estimates are positive across all job types.


Lastly, we replicate the analysis sub-sampling by worker experience on the platform and report the results in Figure 3. We measure experience by counting the total number of jobs held by the worker in the pre-experiment period and bucket 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. 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. Note that 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. When reporting treatment effects as a percent increase from the baseline hire rate, the relative treatment effect is slightly higher for less experienced workers. However, this difference is not statistically significant.

Figure 2: Treatment effect estimates of boosted application on the probability of being hired by job type



Notes: This plot shows the treatment effect estimates of boosted application on the probability of being hired, sub-sampled by job type. We estimate the worker + job posting FE model for each job type separately and report the ADON vs PLACEBO treatment effects. The left panel shows the nominal treatment effect estimates. The right panel shows the relative treatment effect estimates as a percentage increase from the baseline hire rate of the respective job type.

Figure 3: Treatment effect estimates of boosted application on the probability of being hired by worker experience on the platform



Notes: This plot shows the treatment effect estimates of boosted application on the probability of being hired, sub-sampled by worker experience. We estimate the worker + job posting FE model for each experience group separately and report the ADON vs PLACEBO treatment effects. The left panel shows the nominal treatment effect estimates. The right panel shows the relative treatment effect estimates as a percentage increase from the baseline hire rate of the respective experience group.

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

We estimate the effects of boosted applications on post-hire outcomes for workers conditional on being hired using the following regression specification:

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

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 [5](#). We do not find statistically significant effects of boosted applications on post-hire outcomes conditional on being hired. This suggests that workers who were hired by boosting their applications were not worse off after they were hired.

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. Invited workers are excluded from the analysis, since invitees were not eligible to submit a boosted application.

6 Workers use of boosted applications

6.1 Intensity of using boosted applications

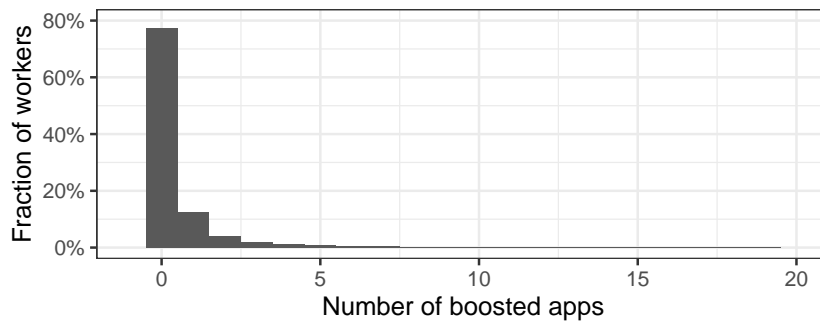
In this section, we examine how often workers use boosted applications. Figure 4a plots the distribution of boosted applications for workers who submitted at least one application during the experimental period. We can see that the majority of workers (77.1%) did not submit a boosted application during the experiment period. Figure 4b restricts the sample to the 22.9% workers who boosted at least one of their applications and plots the distribution of their likelihoods to boost an application. The median worker who submitted at least one boosted application, boosted 33.3% of their applications. However, this figure is skewed by the fact that a large number of workers submitted only one or two applications, and those applications were all boosted. To account for this, Figure 4c further restricts the sample to those workers who submitted at least 5 applications during the experiment period. Among those workers, the median worker boosted 17.7% of their applications.

6.2 Differences between boosters and non-boosters


We showed in Section 5 that boosted applications are positively selected. This raises a question: do workers boost high-quality applications/put more effort into applications, or are workers who boost their applications inherently higher quality? To understand whether there are quality differences at the worker-level, we compare the attributes of workers who had at least one boosted application (boosters) to workers who had no boosted applications (non-boosters), and report the results in Table 6. boosters, on average, had more invitations from employers, applied to more jobs, asked for higher wages, and were hired more often than non-boosters – both in the experimental and pre-experimental periods. These results suggest that there are inherent quality differences (that are not just job-specific) between boosters and non-boosters.

Figure 4: Distribution of boosted applications across workers


(a) Among workers with at least 1 application



(b) Among workers with at least 1 boosted application



(c) Among workers with at least 5 applications and 1 boosted application



Notes: This figure shows the distribution of boosted applications across workers. Subfigure (a) plots the distribution of the number of boosted applications submitted by workers. Subfigure (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. Subfigure (c) further restricts the sample to workers that submitted at least five applications and one boosted application.

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

	Non-Boosters	Boosters	Diff. in Means	p -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 Effects of boosted applications on employers

We estimate the treatment effects of boosted application on a wide range of employer outcomes using the following regression specification:

$$y_k = \beta_0 + \beta_1 \text{TRTADON}_k + \beta_2 \text{TRTADNODISCLOSURE}_k + \beta_3 \text{TRTADNOREC}_k + \epsilon_k, \quad (3)$$

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

We consider a range of 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 ensures that the estimates don't conflate the long-term indirect effects (e.g., employer changing their hiring behavior after being exposed to the treatment) of boosted application on employer behavior and outcomes, whereas the former is net of indirect effects.

For all the above outcomes, whether we only consider the first post or all posts, we do not find any statistically significant effects of boosted applications on employers.

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 How can advertising be a positive signal in equilibrium?

Our empirical results show that boosting an application (i.e., advertising) sends a positive signal to the employer. 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. Using a simple 2-stage model of hiring, we outline the conditions under which a separating equilibrium exists. A key aspect of the model (and an assumption that we test) 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 in the first place. This creates the separating equilibrium. Whereas the classical signaling models of advertising use “repeat purchases” as the linkage that discourages low-quality sellers from advertising, in the hiring setting, the multi-stage nature of the hiring process naturally provides the repeat interaction.

To test whether advertising affects the chances of getting an interview but not the chances of getting hired conditional on an interview, we estimate the effect of boosted applications on these two outcomes. Columns (1) and (2) of Table 8 report the treatment effect estimates of boosted application on the likelihood of getting an interview, and columns (3) and (4) on the likelihood of being hired conditional on getting an interview, respectively. 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 workers who advertise are positively selected, and why advertising can be a positive signal in equilibrium within the hiring setting.

These findings also rationalize why we don’t see increases in the final match success measures between employers and workers. Because hiring is a multi-stage process and the advertising happens in the first stage, all the information there is to be extracted from advertising is extracted in the first stage. This means that advertising is primarily increasing the matching efficiency in the first stage, either by increasing the quality of the shortlist or by reducing the amount of time employers spend reviewing and screening workers. This can explain why we don’t see significant increases in the final match success measures (such as feedback scores) between employers and workers.

Table 8: Treatment effect estimates of boosted application on the likelihood of getting an interview and being hired

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 × BOOST	0.0376*** (0.0014)	0.0384*** (0.0015)	-0.0028 (0.0048)	-0.0043 (0.0076)
ADNoREC × BOOST	0.0365*** (0.0019)	0.0372*** (0.0020)	-0.0002 (0.0060)	-0.0156 (0.0096)
ADNoDISCLOSURE × 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

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: (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. Invited workers are excluded from the analysis, since invitees were not eligible to submit a boosted application.

9 Discussion and Conclusion

The traditional design of online labor markets and labor market intermediaries relies heavily on employers and the platform’s matching technologies to assess job-worker fit. Yet, there is private information that workers have about their fit to jobs that is not observable to employers or the platform, which can be useful in the matching process. We conducted a large-scale field experiment to understand whether advertising via boosted applications can help workers get hired and whether the platform can use this information to improve the matching process for employers. We show that boosted applications are positively selected, and boosted applications increase the likelihood of a worker being hired. This effect is driven by both the ranking and signaling effects of boosted applications. We use a simple model of hiring to illustrate why this positive signal can exist in equilibrium. Lastly, we find no significant effects of boosted applications on a wide range of employer outcomes. These results have several implications for the design of online labor markets and labor market intermediaries.

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 especially helpful to new workers, who struggle to find jobs without having built a reputation on the platform. Indeed our results show that boosted applications are just as effective for less experienced workers as they are for experienced workers. While there is extensive technical literature on ways to mitigate the “cold start” problem in recommendation systems, our work provides an alternative mechanism design approach to mitigate the “cold start” problem in labor markets.

Second, our results show that the ranking of workers in the ATS still plays a crucial role in the matching process. Whereas this result is well established in e-commerce settings, it is less understood in labor markets. In contrast to an e-commerce market, engaging with all the 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 ranking of workers. This suggests that even absent advertising, labor platforms could use the propensity to advertise (or past advertising behavior) to improve the ranking algorithm as proposed by [Long et al. \(2022\)](#).

Lastly, we do not find any significant differences in employer outcomes due to boosted applications. By providing employers and platform with more information about the fit of workers to jobs, boosted applications should, in theory, increase matching efficiency. One reason why we don’t 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 the advertising happens in the first stage, all the information there is to be extracted from advertising is likely being extracted in the first stage. This means that advertising is likely increasing the matching efficiency in the first stage, either by increasing the quality of the shortlist or by reducing the amount of time employers spend reviewing and screening workers, neither of which are directly observable in our data.

In sum, our study highlights the potential of sponsored advertising in online labor markets. By giving workers the ability to send a costly signal of interest and fit for specific jobs, boosted applications provide a new information channel that can be used to improve the matching process.

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.
- Animesh, A., Viswanathan, S., and Agarwal, R. (2011). Competing “Creatively” in Sponsored Search Markets: The Effect of Rank, Differentiation Strategy, and Competition on Performance. *Information Systems Research*, 22(1):153–169.
- 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.
- Diamond, P. A. (1982). Aggregate demand management in search equilibrium. *Journal of political Economy*, 90(5):881–894.
- Filippas, A., Horton, J., Noor, S., and Sorokin, D. (2022). The “coy seller” problem: A market design to reveal willingness to trade. *Working paper*, 0(0):000–00.
- 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]*.

- 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.
- Horton, J. J. and Vasserman, S. (2020). Job-seekers send too many applications. *Working paper*.
- Hosios, A. J. (1990). On the efficiency of matching and related models of search and unemployment. *The Review of Economic Studies*, 57(2):279–298.
- Jeziorski, P. and Moorthy, S. (2018). Advertiser Prominence Effects in Search Advertising. *Management Science*, 64(3):1365–1383.
- 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 Ipeiritis, P. G. (2023). The Good, the Bad, and the Unhirable: Recommending Job Applicants in Online Labor Markets. *Management Science*, 69(11):6969–6987.
- 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.
- Mortensen, D. T. and Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The review of economic studies*, 61(3):397–415.
- Moshary, S. (2021). Sponsored Search in Equilibrium: Evidence from Two Experiments.
- Narayanan, S. and Kalyanam, K. (2015). Position Effects in Search Advertising and their Moderators: A Regression Discontinuity Approach. *Marketing Science*, 34(3):388–407.

- 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.
- Rutz, O. J., Bucklin, R. E., and Sonnier, G. P. (2012). A Latent Instrumental Variables Approach to Modeling Keyword Conversion in Paid Search Advertising. *Journal of Marketing Research*, 49(3):306–319.
- 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.
- Sahni, N. S. and Zhang, C. (2023). Are consumers averse to sponsored messages? The role of search advertising in information discovery. *Quantitative Marketing and Economics*.
- 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.

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 \bar{X}_{ADON}	PLACEBO mean \bar{X}_{PLACEBO}	p-value
<i>Post Characteristics</i>			
number of posts	1.79	1.86	0.176
amount spent	69.55	72.1	0.362
invites sent	9.19	9.19	0.59
fill probability	1.4	1.4	0.986
<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.

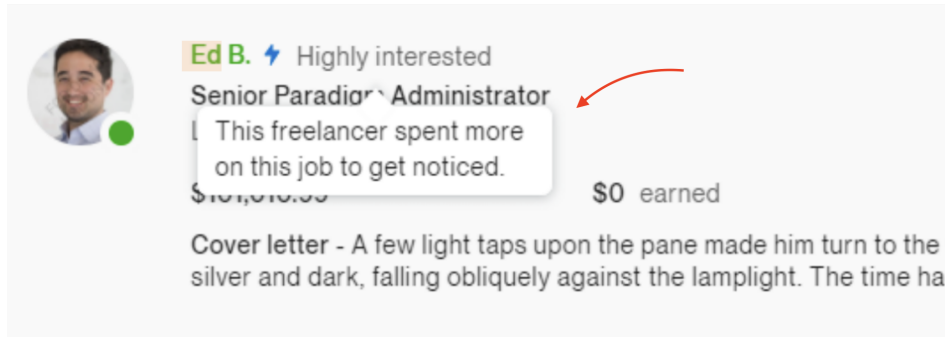
Figure 5: Employers allocated to the experimental groups over time.



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

C An illustrative model of hiring

Consider a labor market in which workers could advertise themselves for a job by paying an advertising cost c . Workers are characterized by their type $t \in \{H, L\}$. H means high-type and L means low-type. The employer prefers to hire H workers but does not directly observe worker type. The employer tries to find out the worker type during the hiring process, which is composed of 2 stages: (1) screening and (2) interview. Interviewing is costly, so the employer shortlists a subset of candidates in the screening stage to interview. In the first stage, the employer observes a noisy signal of worker type, s_1 and whether the worker advertised or not. In the second stage, the employer interviews the shortlisted candidates, and observes a signal $s_2 \in \{\hat{H}, \hat{L}\}$. H workers emit a \hat{H} signal with probability π , and \hat{L} signal with probability $1 - \pi$. Similarly, L workers emit a \hat{L} signal with probability π , and \hat{H} signal with probability $1 - \pi$. A key assumption here is that advertising affects the probability of getting an interview, but not the probability of getting hired conditional on an interview.

Worker's decision to advertise

A worker advertises if the utility from advertising exceeds the utility from not advertising—i.e., $U_{ad} > U_{NoAd}$.

$$U_{Ad} = P(I|s_1, Ad) \cdot P(H|s_2, I) \cdot w - c \quad (\text{A1})$$

$$U_{NoAd} = P(I|s_1, NoAd) \cdot P(H|s_2, I) \cdot w \quad (\text{A2})$$

In the utility function, the first term is the probability of getting an interview. The second term $P(H|s, I)$, is the probability of getting hired (or being a H type) conditional on an interview and signal s_2 . The product of the two is the overall probability of getting hired. w is the wage.

Conditions for a Separating Equilibrium

For a separating equilibrium to exist, H workers advertise and L workers do not. This requires:

1. For H workers, the benefits from advertising must outweigh the costs.

$$U_{Ad}^H > U_{NoAd}^H$$

$$U_{Ad}^H = P(I|s_1, Ad) \cdot [\pi P(H|\hat{H}, I) \cdot (1 - \pi)P(H|\hat{L}, I)] \cdot w - c \quad (\text{A3})$$

$$U_{NoAd}^H = P(I|s_1, NoAd) \cdot [\pi P(H|\hat{H}, I) \cdot (1 - \pi)P(H|\hat{L}, I)] \cdot w \quad (\text{A4})$$

This implies:

$$[P(I|s_1, Ad) - P(I|s_1, NoAd)] \cdot [\pi \cdot P(H|\hat{H}, I) + (1 - \pi) \cdot P(H|\hat{L}, I)] \cdot w > c \quad (\text{A5})$$

2. For L workers, the costs from advertising must outweigh the benefits.

$$U_{Ad}^L < U_{NoAd}^L$$

$$U_{Ad}^L = P(I|s_1, Ad) \cdot [\pi P(H|\hat{L}, I) \cdot (1 - \pi)P(H|\hat{H}, I)] \cdot w - c \quad (\text{A6})$$

$$U_{NoAd}^L = P(I|s_1, NoAd) \cdot [\pi P(H|\hat{L}, I) \cdot (1 - \pi)P(H|\hat{H}, I)] \cdot w \quad (\text{A7})$$

This implies:

$$[P(I|s_1, Ad) - P(I|s_1, NoAd)] \cdot [(1 - \pi) \cdot P(H|\hat{H}, I) + \pi \cdot P(H|\hat{L}, I)] \cdot w < c \quad (\text{A8})$$

Let Δ_{Ad} be the change in the probability of getting an interview due to the interview.

$$\Delta_{Ad} = P(I|s_1, Ad) - P(I|s_1, NoAd)$$

Combining the two conditions, we get:

$$\Delta_{Ad} \cdot [(1 - \pi) \cdot P(H|\hat{H}, I) + \pi \cdot P(H|\hat{L}, I)] \cdot w < c < \Delta_{Ad} [\pi \cdot P(H|\hat{H}, I) + (1 - \pi) \cdot P(H|\hat{L}, I)] \cdot w \quad (\text{A9})$$

which dictates that the cost of advertising must be low enough for H workers to find it worthwhile to advertise, and high enough for L workers to discourage from advertising. The above inequality further dictates that:

1. $\pi > \frac{1}{2}$. H worker is more likely to emit a \hat{H} signal, and L worker is more likely to emit a \hat{L} signal.
2. $P(H|\hat{H}, I) > P(H|\hat{L}, I)$. The probability of hiring a worker that emits a high signal is greater than the probability of hiring a worker that emits a low signal, independent of whether they advertised or not.

3. $\Delta_{Ad} > 0$. Advertising increases the chances of getting an interview. This will be true in equilibrium since H workers advertise and L workers do not.