

Why Search Advertising Works, When It Works: Evidence from a Field Experiment*

Apostolos Filippas John J. Horton Prasanna Parasurama Diego Urraca

January 25, 2025

Abstract

A common concern about sponsored advertising is that advertisers are adversely selected, worsening the buyer experience and requiring the platform to trade-off between transactional and advertising revenue. We study the introduction of advertising in a large online labor market where workers could bid to receive one of three top positions in the employer’s application list. The platform randomly varied whether employers were exposed to advertising, and within the exposure group, whether employers could see that the application was sponsored, whether advertising changed the applicant’s ranking, and whether employers saw algorithmic recommendation labels. 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 show that advertising can substitute for algorithmic recommendation labels, and that new and experienced workers benefit equally from advertising. We find no difference in post-hire outcomes between the treatment and control groups. 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

In many online platforms, sellers can pay to display their listings more prominently to buyers. This “sponsored advertising” was pioneered by internet search companies, and has been adopted by online platforms with very different scopes and designs. Some examples of sponsored advertising include sellers promoting their products on Amazon, real estate agents promoting their listings on Zillow, course creators promoting their courses on Udemy, and companies promoting their job openings on LinkedIn. Today, sponsored advertising constitutes a substantial revenue stream for many of these platforms.¹

A large body of research has documented ambiguous effects for sponsored advertising. A common concern is that advertising only has a business-stealing effect, altering the platform’s “organic” rankings without improving consumer welfare. If advertisers are adversely selected, advertising may even reduce consumer welfare (Abhishek et al., 2022; Moshary, 2024).² On the other hand, advertising can create a separating equilibrium where high-quality sellers are more likely to advertise, because these sellers benefit most from advertising. Advertising then signals high quality and can increase consumer surplus, as well as help new or underpromoted sellers to be noticed (Nelson, 1974; Athey and Ellison, 2011; Sahni and Nair, 2020a; Dai et al., 2023). These contrasting results suggest that the overall effect of sponsored advertising depends on its mechanism design, the underlying market structure, and the relative importance of its signaling and ranking effects.

This paper provides direct empirical evidence on the effects of sponsored advertising by reporting the results of a field experiment in 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 in an algorithmically determined order, and then shortlist, interview, and hire candidates. During the experiment, all workers became eligible to bid to “boost” (sponsor) their applications when applying for a job. The platform randomly varied whether and how employers were exposed to boosted applications. The design of the experiment is key to disentangling the distinct effects of sponsored advertising.

The experiment had four treatment groups. 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, but 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

¹For example, Amazon reported a \$47B revenue from advertising in 2023. For more details, see <https://www.statista.com/statistics/259814/amazons-worldwide-advertising-revenue-development/>.

²FTC alleged that sponsored advertising on Amazon degrades customer experience by deprioritizing relevant, organic search results in favor of paid advertisements.

additional algorithmically determined recommendation label—which employers in all other groups saw.³ 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 are a particularly good match with a job, and/or they put more effort into applications that they choose to boost.

Second, in line with positive self-selection into advertising, 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 signalling 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.

³Algorithmic recommendation labels are common many platforms. For example, Amazon displays “Overall Pick” and “Best Seller” labels next to some listings in addition to sponsored listings.

Fifth, we find that the experimental assignment made no had no effect on employers: the ability to see boosted applications did not affect the employers’ hiring process including the number of interviews extended, the time to hire, and the number of hires made, as well as post-hire outcomes such as the amount of money spent, or the ratings given to workers. This result suggests that there was no efficiency gain on the buyer side from the introduction of sponsored advertising. On the other hand, it suggests that the introduction of sponsored advertising did not lead to a degradation of the buyer experience—in contrast with the common concern that employers encountering advertising for the first time on the platform might have been eager to try hiring advertising workers, only to be later disappointed by their performance.

An important difference of our study compared to previous work is that our empirical setting is a labor market. Compared to e-commerce platforms, labor markets have several distinct characteristics that make predicting the effects of sponsored advertising difficult. On the supply side, workers have private information about their match with a job that neither employers nor the platform has. Furthermore, advertising is highly targeted—workers choose exactly which jobs they advertise for. But because workers have finite capacity, employers may expect high-quality workers to have enough work to not need to advertise. On the demand side, employers’ consideration sets are typically smaller than in e-commerce—a job receives about 20 applications in our setting. Moreover, hiring is a high-stakes decision than purchasing a product, and engaging with all applicants can be beneficial in wage negotiations. Whether advertising could meaningfully affect hiring decisions at all is unclear. Despite these contrasting factors, we see positive selection into advertising and a positive effect of advertising on hiring outcomes.

We explore why positive selection into advertising and advertising equilibrium exist in our setting through a simple model of hiring and advertising. In our model, the multi-stage nature of the hiring process provides the “repeat interaction”—similar to repeat purchases in signaling models of advertising—that discourages low-quality workers from advertising at the same level as high-quality workers. In other words, advertising is a signal of high-quality because it only increases the likelihood of a worker getting an interview, but not the likelihood of getting hired conditional on being interviewed. We provide empirical evidence that this is the case 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 sponsored advertising literature. First, we document strong evidence of positive selection into advertising. Second, we disentangle the distinct but often concomitant effects of sponsored advertising into: (1) the self-selection effect, (2) the signaling effect, (3) the ranking effect, and (4) the interactive effects with other algorithmic recommendation labels. Third, we provide the first empirical evidence on the effects of sponsored advertising in a labor market setting, and we show how repeat

interactions in the hiring process makes it possible for an advertising equilibrium to exist.

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 sketches a parsimonious model 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. [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 \(2024\)](#) 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.

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, 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.”⁴ 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

⁴Coins 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.

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., 2022).

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, of which 5.19 were invited applications. The average job fill rate was 0.53, meaning that employers made a hire in about half of the jobs. 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 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” 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.

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.

- 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

There were three boosted application slots per job. 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, 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 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 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 outcomes in Table 3 and for hire outcomes in Table 4. In each table, Column (1) reports regression estimates for the specification of Equation 1. Column (2) adds worker fixed effects to control for within worker variation in boosted applications. Column (3) adds 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. Column (4) adds both worker and job posting fixed effects. This 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. We discuss the results next.

5.2 Boosting applicants were positively selected

In Table 3 and Table 4, the estimate for the variable BOOST captures the difference in outcomes between boosted and non-boosted applications within the PLACEBO group. Looking at Column (1), 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 (2) and (4), 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.

5.3 The aggregate effect of boosting

The interaction terms in Table 3 and Table 4 capture the causal effect of boosting on the likelihood of being interviewed and hired. Looking at Column (4), 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

Table 3: Treatment effect estimates of boosting 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.

PLACEBO cell. The estimates for the that 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

Table 4: Treatment effect estimates of boosting 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.


the treatment. The results remain highly similar to those of Table 4.

5.4 The ranking effect of boosting

In Tables 3 and 4, the estimates for the interaction term $ADNoDISCLOSURE_j \times BOOST_{i,j}$ 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 employer being interviewed by 20.3% (2.7

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 4 of Table 4).

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 both 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 -value = 7.4e-08), and being hired by 8.3% (0.2 pp;

p -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.⁵ 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 5.

In Column (1) of Table 5, 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.

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.

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

⁵Note that we have to conduct this analysis at the job-category level, rather than at the job level, because there is no variation in treatment within a job. There are 176 such job categories, such as Article & Blog Writing, 3D Animation, Back-End Development, and so on.

Table 5: 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.

in the two estimates is not statistically significant. This suggests that the 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 into one of 12 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.

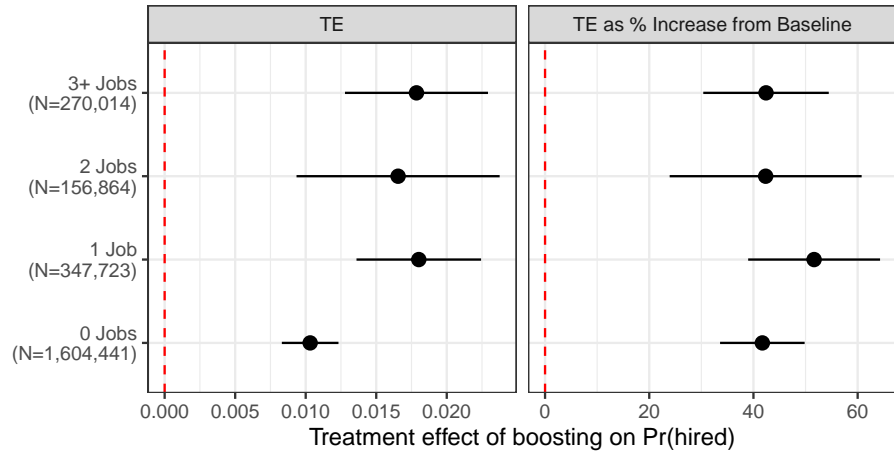
The effects of boosted applications could also vary by worker experience on the platform. We replicate the analysis of Figure 2a, but now splitting our data by worker experience on the platform. We measure worker experience by counting the total number of jobs held by the worker in the pre-experiment 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

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.

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 and non-boosting applicants. We estimate 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 Table 6. 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.

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 3a 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. We then restrict our sample to the 22.9% workers who boosted at least one of their applications, and we plot the distribution of their likelihoods to boost an application in Figure 3b. The median worker who submitted at least one boosted application, boosted 33.3% of their applications. However, this result can be skewed by the fact that workers who submitted only one or two applications during the experimental period may have boosted all their applications. To account for this, Figure

Table 6: 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.

3c 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. The distribution closely resembles the lognormal distribution, which is consistent with what we would expect.


6.2 Differences between boosters and non-boosters

We showed that boosted applications are positively selected in Section 5. 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 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 7 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 inherently different than non-boosters, and the positive

Figure 3: 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.

selection that we see cannot be explained by job-specific reasons.

7 The effects of boosted applications on employers

We not 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

Table 7: 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.

of the effects of boosting on employers. We estimate the following specification:

$$y_k = \beta_0 + \beta_1 \text{ADON}_k + \beta_2 \text{ADNO DISCLOSURE}_k + \beta_3 \text{ADNO REC}_k + \epsilon_k, \quad (3)$$

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 8a, and the estimates for the first job posting an employer makes once they are allocated to a treatment cell in Table 8b. 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.

Table 8: 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.

For all outcomes, we do not find any statistically significant effects of boosted applications on employers in neither sample. This suggests that boosted applications do not affect employers outcomes in our setting—even though they affect worker outcomes. Furthermore, 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.

8 Why does an advertising equilibrium exist in our setting?

We showed empirically that boosting—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. We next sketch out a parsimonious model of hiring to illustrate why this equilibrium exists in our context.

Consider a labor market comprising employers, and workers of high and low quality. Workers are privately informed about their quality, and can choose to self-advertise by boosting their applications. The employer hiring process has two stages: a costless screening stage, and a costly interviewing stage. During screening, the only signal available to the employer is the worker’s decision to advertise; during interviewing however, high-quality workers are more likely to emit a positive signal than low-quality workers.

The key aspect of this model 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, which results in a separating equilibrium. It is worth noting that classical signaling models of advertising, such as the one developed by [Nelson \(1974\)](#), rely on “repeat purchases” as the mechanism that discourages low-quality sellers from advertising. In a labor market, the multi-stage nature of the hiring process plays the role of the repeat interaction. For more details on our model, see [Appendix C](#).

Our data allows us to test directly the main assumption of our model—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 9](#), 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 markets rely heavily on employers and platform matching technologies to assess job-worker fit. Yet, workers have information about their fit with jobs that neither employers or the platform can observe, but which can be useful in the matching process. We reported

Table 9: Treatment effect estimates of boosting on being interviewed and 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
<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. Invited workers are excluded from the analysis, since invitees were not eligible to submit a boosted application.

on the results of a large-scale field experiment in an online labor market that can help us understand whether advertising via boosted applications 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. Through a simple model of hiring, we illustrated why this positive signal exists in equilibrium. 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. 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: we 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, advertising offers an alternative mechanism to mitigate the “cold start” problem in labor markets.

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 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 (Long et al., 2022).

We did not find any significant differences in employer outcomes due to boosting. 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 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.
- 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.
- Athey, S. and Ellison, G. (2011). Position auctions with consumer search. *The Quarterly Journal of Economics*, 126(3):1213–1270.
- 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.
- 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., 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.
- 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 Ipeirotis, 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.
- Moshary, S. (2024). Does sponsored search advertising augment organic search? evidence from an e-commerce platform. *Management Science*.
- 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.
- 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 10, 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 4 plots the allocation of employers to the treatment and control cells over time.

Table 10: 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	0.18
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.

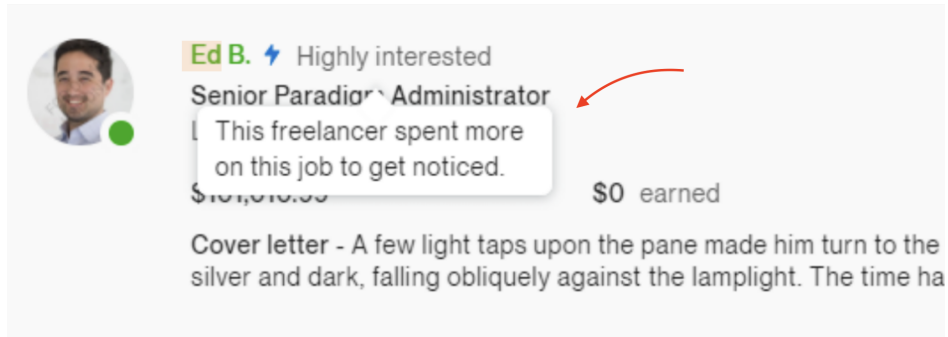
Figure 4: Employers allocated to the experimental groups over time.



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 5: 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 6: 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 11: 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 12: 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 quality 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.