Information About Vacancy Competition Redirects Job Search[*]

Monica Bhole
Facebook Research

Andrey Fradkin
Boston University

John Horton
MIT & NBER

July 16, 2021

Abstract

Job seekers typically do not know the degree of competition they face for a particular vacancy. As a result, they may unwittingly send applications to vacancies with a lot of competition and may overlook vacancies with little competition. We study how providing information about competition for a vacancy redirects applications. To do so, we conduct three field experiments on a large online job platform in which treated job searchers are shown information about the number of prior applicants to a vacancy. This information increases overall applications and redirects applications to vacancies with few prior applications. Applications are sent to vacancies that receive fewer cumulative applications but result in similar outcomes to control applications. We use a complementary treatment to show that job seekers also use the age of the vacancy to direct search towards newer vacancies with relatively little competition. Our results are consistent with a model in which searchers have imperfect information about competition for a vacancy and redirect their search towards less competitive vacancies when they receive an improved signal.

1. Introduction

Job seekers cannot coordinate their applications. As a result, some vacancies may receive too many applications, while others may receive too few relative to the efficient allocation. This feature of search is a long-theorized source of matching frictions with potentially large welfare consequences (Petrongolo and Pissarides (2001)). Theories of directed search posit that coordination frictions can be mitigated but not eliminated when observable characteristics of a job are correlated with the degree of competition for a vacancy (Wright et al., 2019). A key assumption in these models is that information about competition is not directly observed by applicants—but this no longer needs to be true in a digital world.

We study how job seekers use information about the degree of competition for a vacancy to direct their applications on a digital platform. The provision of this information may reduce search frictions if it causes searchers to redirect applications to less competitive vacancies, as predicted by some models of search. On the other hand, these signals may induce herding, as in Salganik et al. (2006) and Muchnik et al. (2013), and thereby exacerbate search frictions.

To conduct our study, we use three experiments encompassing millions of job seekers conducted on an online job board. We find that information about the vacancy level degree of competition affects job search, and in particular, directs applications towards vacancies with few applications in line with the spirit of directed search models. The applications induced by the information do not fully crowd out other applications and exhibit similar outcomes to those in the control group.

The setting of our study is a job board operated by Facebook called 'Jobs on Facebook' (JOF). JOF is global, and, as of 2019, mainly caters to full-time positions that do not require a college education. Users of the platform see a list of vacancies and have access to a rich set of filters by which to refine their search. Searchers have the ability to apply to most vacancies using the platform. Consequently, the platform is able to directly observe the number of applications sent to a particular vacancy in real-time and to display information about the count of prior applications to searchers.

We begin the paper by documenting several motivating facts. First, we show that being an earlier applicant increases the likelihood that an application is viewed, responded to by an employer, and results in an interview. Second, we show that, even within the same commuting zones, some vacancies receive many applications while others receive few. This imbalance in job applications may simply reflect that some vacancies were a better match for the job seekers on the platform, but it might also be a sign of an inefficient allocation

applications which could be addressed by our treatments. Finally, we show that the number of applications received by a vacancy is highly correlated with the frequency with which it is shown to job seekers. This could be induced by the preferences of job-seekers, but it also suggests a role by the platform in which jobs it tends to promote—a factor that job-seekers might find it difficult to condition upon when searching.

Based on the above facts, we propose a theoretical model of job search, adapted to our platform setting. In this model, vacancies differ in their degree of exposure to job seekers and the number of hiring slots they have per vacancy. Searchers are exposed to vacancies and must make a decision of whether to apply to a job based on the expected number of other applicants and the number of hiring slots. In the baseline, searchers do not know the exposure level of the vacancy they see or the total number of other job seekers on the platform. We then consider a full information counterfactual in which searchers know about the exposure of the job to other job seekers and the number of other job seekers on the platform. Relative to the no information scenario, applications in the full information scenario are redistributed from high exposure vacancies to low exposure vacancies.

We test this theoretical framework in a series of experiments in which the treatments approximate the model based counterfactual by providing additional information to job seekers. Starting in March 2019 and continuing through August of 2019, Facebook conducted three experiments. All three experiments contained treatment arms that displayed information about the number of prior applicants to a vacancy in the search interface. The different treatment arms and experiments varied the frequency (every vacancy, every three vacancies, or every 10 vacancies) and color (grey vs blue) of the information.

We find that these treatments increase applications to vacancies with fewer than 5 applications by 3.8%, with a range of .9% to 6.4% depending on the experiment. As predicted by our theory, applications to vacancies with many applications fall. We also find that the total number of application increases due to the treatment. Therefore, the additional applications to low competition vacancies do not fully crowd out existing applications. The frequency and the color of information did not have first order effects on the rate of applications to these vacancies.

Next, we study the causal effect of the treatment on the quality of applications submitted. A key limitation of our setting is that we do not observe offers and hires. Therefore, we cannot directly measure match quality and must rely on proxies of application quality, the match between searcher and vacancy, and the response behavior of the employer. Applications in the treatment were more likely to be viewed, but not more likely to be contacted or interviewed through the platform.

Finally, we consider the role of vacancy age in directing job searchers towards less congested vacancies. In our experiment, vacancy age was always visible on the platform in our control groups. Vacancy age is positively correlated with the number of applications. As such, it is potentially a proxy for congestion and is available to job-seekers in many online settings. It also is informative about the probability that the employer has already made a decision, and thus whether an application is worth sending. In contrast to our other treatments in which the platform added signals to help direct search, removing vacancy age potentially removed a signal that otherwise could be used to direct search. But removing vacancy age helps us better understand what job-seekers were trying to accomplish.

We first show that vacancy age has an effect on the likelihood that an employer views an application, even conditional on the number of prior applications. As a consequence, searchers should attend to vacancy age when choosing whether to apply. Second, we show that workers use vacancy age to decide where to apply. We find that job seekers who do not have information about vacancy age clicked on 3% fewer vacancies and sent 1.8% fewer applications. This treatment also has distributional effects. Treated users were less likely to apply to new vacancies and were more likely to apply to old vacancies. Since vacancy age is correlated with the amount of competition for a vacancy, we find that removing vacancy age increases the concentration of applications. Furthermore, information about vacancy competition increases applications precisely to older vacancies.

We interpret the above results with the help of our theoretical framework. The redistribution of applications towards vacancies with fewer prior applications is efficient in our model, since the average applicant faces less competition. However, the effects on searcher welfare are ambiguous. Because applications exert negative externalities on other searchers, the overall level of applications may or may not be efficient relative to a social optimum (Hosios (1990)).

We do not find that behavior in the job market exhibits the type of social contagion based on popularity signals found in other digital settings (Salganik et al. (2006), Muchnik et al. (2013)). In other words, there is no danger of indicating a job is "popular" causing it to receive even more applications, as in some kind of social learning or information cascade scenario. In short, job-seekers view the application process more as a congestion game and all else equal, would prefer facing fewer competitors.

Our results provide causal evidence about several mechanisms present in models of labor market search. At a high level, job seekers should care about the probability of getting an offer and must infer that probability from job characteristics. We provide empirical support for a general class of directed search models (Wright et al. (2019)) based on this assumption.

Cheron and Decreuse (2017) and Albrecht et al. (2017) focus specifically on the importance of 'phantom vacancies', which are vacancies that have already been filled. In their models, workers rationally direct search towards newer postings — a phenomenon which we document empirically. We also show how competition information can further direct job search towards less congested vacancies.

Our treatments are enabled by the fact that digital job boards have a bird's eye view of the market. While this may seemingly limit the applicability of this approach, given that the labor market as a whole is decentralized, an increasing amount of job search occurs on digital job boards (Kuhn and Mansour, 2014; Baker and Fradkin, 2017; Kroft and Pope, 2014; Marinescu, 2017). These job boards make decisions that could ameliorate—or worsen—congestion. Due to the heterogeneity in preferences for vacancies-across seekers and vacancies, centralized matching is infeasible. Instead, the platform can influence matching indirectly through the information it chooses to display and emphasize.

The paper most closely related to ours is by Gee (2019), who varied whether the number of applicants to a vacancy was shown on LinkedIn. In Gee (2019), popularity information could only be learned after clicking on a job, whereas in our setting popularity information is shown directly in the search interface. The timing of information acquisition has previously been shown to matter greatly for outcomes in search markets (Branco et al. (2012); Hodgson and Lewis (2020); Gardete and Hunter (2020); Abaluck and Compiani (2020)). This difference in timing may explain why Gee (2019) found no significant interaction between the number of applicants a job had received and the decision to apply, whereas we did. In short, congestion avoidance appears to occur higher up the funnel.

Several other papers have used data on search in digital labor platforms. Faberman and Kudlyak (2019); Marinescu (2017) study how search evolves over the course of an unemployment spell, Marinescu (2017); Marinescu and Skandalis (2019); Baker and Fradkin (2017) study how unemployment insurance affects online job search. Azar et al. (2019) use data from CareerBuilder to build a demand model of applications and use it to estimate firms' market power in the labor market. Le Barbanchon et al. (2020) use data on search criteria declared to a public employment agency, as well as applications on a digital platform, to show that women care more than men about commuting when it comes to applying for jobs. Skandalis (2018) shows how job search is affected by news about a company's hiring needs.

The above literature has not focused on the market design possibilities available on job boards. Our paper shows that the information design of these job boards has large effects on behavior. This creates opportunities for additional information interventions, and the

study of their equilibrium effects using market level experiments. In this way, the market design innovations pioneered in other digital platforms, such as those for labor procurement (Horton (2017)), dating (Fong (2019)) and accommodations (Fradkin (2017)), can be used to improve outcomes on digital job boards.

The rest of the paper is organized as follows. Section 2 describes the empirical context. Section 3 presents evidence of competition and concentration in applications that motivates some of our treatments. Section 4 described our theoretical framework. Section 5 discusses the design of our various interventions. Section 6 reports the effects of the treatment on search behavior. Section 7 reports on the differences between applications across treatments. Section 8 discusses the our results with respect to vacancy age. Section 9 concludes.

2. Empirical context

JOF is an online job board that was launched in the US in 2016 and then worldwide in October 2018. The product is global in nature and mainly caters to positions which do not require a college education. The share of US users across our experiments ranged from 21% to 25%, and the median user in our experiments was between 31 and 33 years old across experiments. Employers can post vacancies and job-seekers can browse vacancies and send applications through the platform. The service is free for both sides, but job-seekers must have a Facebook account. A posted vacancy is automatically live for 30 days, but employers can renew it. This results in an average duration during our testing period of 42 days. [1] Even before the launch of JOF, there was substantial job-search behavior on Facebook (Gee et al., 2017).

Job-seekers are exposed to JOF via the "News Feed" and via notifications. [2] They can also navigate to JOF by clicking on the "Explore" tab and then clicking on a briefcase icon labeled "Jobs." The JOF interface is similar to other job boards, though most of job search occurs on mobile devices. That most use occurs on mobile presents opportunities—if the user has enabled location-tracking, vacancies within a given radius can easily be shown—but also challenges, in that there is a constrained space in which to present information.

When looking for work, job-seekers can enter a number of criteria to narrow their search, including their location and the type of position they are interested in. Figure 1 shows the

---

[1] We calculate this duration for all vacancies created between March 2019 and August 2019. Note that Facebook reviews job postings and may remove them from the platform prior to the scheduled expiration of the job. Therefore, this number is an overestimate of the length of time that a job may be visible on the platform.

[2] Which users are exposed to the Jobs product in the News Feed is determined by an algorithmic system.

status quo job search interface (the job board) —as we will discuss at length, this presentation was modied by various treatments. As we can see in the gure, for each vacancy, the job-seeker can see the title of the job, whether it is part-time or full-time, the name of the employer, number of days since it was posted, and if the employer has posted the wage, the hourly wage. To learn more about the vacancy, the job-seeker has to click on the "tile" for that opening. Clicking exposes a "detailed view" of the job that includes the full job description written by the employer. It also includes an "apply" button that the job-seeker can use to submit an application.

2.1 Measurement of job-search behavior

We observe the job vacancies a user loaded onto the job board during their search, which is a function of how far they scrolled down the device, their location, and their search query parameters. A "view" occurs when a vacancy appears on a screen. We also observe whether the user clicked on a vacancy to learn more, which we call a "detailed view." Finally, we observe whether the job-seeker applied to a particular opening and some information about that application.

When job-seekers start an application, information about them is populated into an application, using data from their Facebook prole — educational history, past employment, contact information and so on. Searchers can ll in additional information that is not already listed on Facebook. Our application measure is likely a lower bound on the number of job applications created, as in some cases, job-seekers would have enough information about the employer to apply directly. However, the convenience of simply submitting through the Facebook App makes this the most likely course of action.

After an application occurs, we have imperfect information about what happened. For vacancies created through the JOF platform, which we call `native', we can observe a variety of interactions including whether an employer viewed an application, whether the employer contacted an applicant through Facebook, and whether an employer told Facebook whether an interview was scheduled. Each step in this process is `leaky', so that we see a large share of applications are viewed, but a much smaller share have contacts and interviews. This partially occurs because at each step employers and applicants can choose to take the interaction off of the platform. There are also vacancies which are syndicated from other platforms, which we refer to as `third-party vacancies'. In the US, third party platforms can be applicant tracking systems. For some of these vacancies, we cannot measure interactions between applicant and employer because the interactions take place off of the

Figure 1: Status quo job search interface for Jobs on Facebook job-seekers

Notes: Interface shown to job-seekers on a mobile device.

platform.

3. Evidence on competition and concentration in applications

We now describe three empirical regularities which motivate our model and experimental analysis. First, earlier applicants have a higher likelihood to have their applications viewed and responded to. This justi es the desire by applicants to apply before others. Second, we show that some vacancies receive many applications whereas others do not and this occurs even conditional on location and time period. Of course, this inequality may occur simply because some vacancies are better than others. However, the presence of inequality in applications across vacancies suggests that there could be scope for information treatments to reduce congestion. Third, the imbalance in applications coincides with an imbalance in the likelihood that a vacancy is exposed to users in search. The advantages of applying earlier combined with imbalances in applications across vacancies point to a role of information interventions, which we discuss in the next section.

3.1 Effects of applying earlier

We now show that applying earlier (whether measured by the application order or time) results in a higher probability that an application receives a positive outcome. The simplest explanation for this is that an employer checks applications at some time that the applicants cannot condition on. As a result, when everything else is held constant, earlier applications are more likely to be considered. However, it is possible that better applicants apply more quickly, which could create the illusory appearance of application rank mattering, and so we have to be cautious in claiming a true causal effect.

   To measure the effect of application order, we can compare the outcomes of multiple applications to the same vacancy. In particular, we estimate regressions of the following form:

$$y_{ij} = \alpha_{order;ij} + \alpha_{age;ij} + \gamma_j + \epsilon_{ij} \tag{1}$$

where y_{ij} are application outcomes for vacancy j by application i, $\alpha_{order;ij}$ are xed effects for the order of the application (e.g. 1st application, 10th application), $\alpha_{age;j}$ are xed effects for the age of the vacancy at the time (e.g. 1 day, 3 days) at which the application was submitted, and γ_j are vacancy xed effects. We estimate the above equation on a 5% sample

Figure 2: Effects of application order on employer responses

Notes: Each point represents the estimated effect and each line present the 95% confidence interval for estimates of the effects of application order on whether an application is viewed, responded to, or results in an interview. Standard errors clustered at a vacancy level.

of vacancies posted between March 3 and August 18 of 2019, the period during which our experiments ran.

Figure 2 displays the estimated coefficients on application order from the above equation, where the first application is normalized to 0. Later applications have a lower probability of being viewed, contacted, and interviewed. To get a sense of the magnitudes involved, the baseline rate of interviews in this sample is 1.3%, which reflects the fact that many employers do not use the JOF interface to record interviews. We see that later applications experience an over 65% decline in interview rates for these employers. To summarize, application order matters for application outcomes, and this justifies our focus on providing application order information to job-seekers.

3.2   Imbalances across vacancies

To measure outcomes at a more aggregate level, we need an appropriate definition of a market, which for labor markets is typically a commuting zone (CZ). We use a global mapping of locations to CZs based on location data, which is constructed by Facebook (Facebook (2020)).

---

[2]Note that there exist some vacancies which do not allow the user to apply using the Facebook platform. These vacancies make up a small number of impressions and are typically syndicated from outside of the Facebook platform. We do not include these in the above calculations.

Figure 3: Prevalence of Under-subscribed vs Popular Vacancies by Commuting Zone - Week

Notes: Each point plots the share of vacancies with more than 100 applications against the share of vacancies with fewer than 5 applications within a month of posting by commuting zone and week. The sample includes weeks between `2019-02-26' and `2019-09-01'. Only commuting zones with at least 100 posted vacancies are included.

In Figure 3, we plot the share of all vacancies in a given CZ-week that received fewer than 5 applications in the rst month that it was open versus the share of all vacancies in a given CZ-week that received over 100 applications. The plot demonstrates that there is imbalance in applications across vacancies. Vacancies with many applications exist in the same commuting zones as vacancies with few applications. Of course, there are many factors that could explain this pattern that are not per se inef cient—job-seekers could be avoiding "bad" jobs and seeking out "good" jobs–but the general pattern is suggestive that there is an opportunity for a better distribution of applicants over vacancies.

The imbalance in applications coincides with the imbalance in the frequency with which vacancies are viewed by job seekers. We are able to calculate the number of views each vacancy received as well as the number of applications. There is a correlation of 0.47 between the number of views a job has and the number of applications it has in this sample. This may re ect both the fact that some vacancies get exposed more due to their location in a place with many searchers, due to the fact that vacancies vary in other characteristics related to search, or due to ranking algorithms. Furthermore, since such algorithms learn from prior users, there may be reverse causality between applications and exposures.

Figure 4: Imbalance in Views by Vacancy

Notes: Histogram of the share of normalized inter-quartile range of the number of views per job. The sample aggregates over all exposures between `2019-07-22' and `2019-08-18'. Only commuting zones with at least 100 posted vacancies are included.

Nonetheless, we find that the imbalance in applications across vacancies is also reflected in exposures. Figure 4 displays the inter-quartile range of exposures with a commuting zone normalized by mean exposures in that CZ. The difference in the inter-quartile range is typically greater than the mean number of views, and there is substantial variation across CZs. These facts regarding views and applications motivate our theoretical framework and experimental analysis.

4.   Theoretical Framework

We now describe a simple theory of search and competition for vacancies on the platform. This theory allows us to derive comparative statics for interventions that affect the information known by searchers and guides our empirical analysis. The intuition behind the model is that, in the absence of information about the congestion of a vacancy, searchers must choose whether to apply based on an expectation about congestion in the market. Once searchers learn the level of congestion in the market, they increase their application rates to less competitive vacancies and decrease their application rates to more competitive vacancies.

Consider a setting with a mass S of searchers and a mass 1 of vacancies. Each searcher, i, arrives at the platform and is exposed to one vacancy, j. Which vacancy is viewed depends on an exposure parameter, $e_j \sim U[0; 1]$. Searchers view a job of type j with an exogenous

density.

$$R\frac{e_j}{\int_0^1 e_t \, dt} = 2\,e_j \tag{2}$$

Each searcher receives a utility from applying to a job of $u_{ji} = v - k_j a_j - c_{ij}$. In this equation, v represents the expected utility of applying to the job when there are no other applicants. a_j is the number of applications to the vacancy and k_j is a parameter that governs how sensitive the utility of applying to a job is to the number of other applicants. k_j is known to searchers and is meant to capture the fact that some vacancies, such as for Amazon delivery drivers, have many slots while others, such as for baristas at the local coffee shop, may only have one slot. c_{ij} represents an idiosyncratic application cost net of idiosyncratic match quality. We assume that $c_{ij} \sim F$. Lastly, assume that workers have a belief about the mass of searchers in the market, $\hat{S} = S$.

Proposition 1: When the worker has no information (NI) about vacancy level exposure, application rates are decreasing in k, decreasing in S, and increasing in v. All vacancies with the same k have the same application rate.

Proof: See Appendix A

Proposition 2: Consider the equilibrium when workers are given information about exposure, e_j. Call this equilibrium EI (exposure information). For vacancies with $e_j < p$, applications increase relative to equilibrium NI, and for vacancies with $e_j > p$ they decrease, where $p = \frac{2}{3}$. The change in applications is moderated by k, vacancies with low k are less affected by the policy than vacancies with high k.

Proof: To see this, consider the fact that under NI, the decision is made over the expected number of other applications to an exposed job. This coincides with the number of applications submitted in scenario EI when $e_j = \frac{2}{3}$. Furthermore, the difference in applications rates, $r^I(k_i; e_i) - r^{NI}(k_j; e_j)$ is decreasing. Details in Appendix A.

An important feature of JOF is that it is a relatively new and fast growing platform. As a result, searchers may not know how many other searchers are using the platform. In that case, disclosing information about the number of prior applicants may also reveal information about overall usage on the platform. We conceptualize this as searchers learning the true number of searchers, S, in the market.

Proposition 3: When searchers overestimate the number of other searchers in the market, revealing information about the true number of searchers in addition to exposure information increases application rates.

Proof: To see this, note that application rates increase when searchers believe that fewer other people are applying for the job. Details in Appendix A.

Figure 5: Example equilibrium of congestion model

Notes: This gure plots three scenarios of the congestion model. We assume that the true mass of searchers is 10, the biased belief about the number of searchers is 11, that k is .1, and that v equals 2. Lastly, we assume that costs are normally distributed.

To summarize, Figure 5 plots the equilibrium application rates in the scenarios: no information (NI), exposure information (EI), and market size information along with exposure information (MSI). The gure is consistent with our key results. Giving searchers information about exposure increases application rates for low exposure vacancies and decreases them for high exposure vacancies. Giving searchers information about the number of other searchers in the case when searchers overestimate the number of other searchers increases application rates relative to the EI scenario.

Our theory does not yield unambiguous predictions regarding the effects of information on the total number of applications or welfare. The reason for this is that information redistributes applications across vacancies. Searchers who would have applied to high exposure vacancies regardless of information are better off since they face less competition. On the other hand, searchers who would have applied to low exposure vacancies are worse off since they now face more competition. In our model, these comparative statics depend on the distribution of application costs. In a more general model, they would also depend on the distribution of exposure probabilities across vacancies. Lastly, in our experiment, we compare treated searchers to control searchers, rather than comparing two distinct equilibria, as in the above model. This requires a nuanced interpretation of our experimental results.

5.   Experimental provision of information

Our primary research interest is in how information about the competition for a particu-
lar vacancy affects job-seeker decision-making. We study this question by analyzing three
experiments conducted over a ve month span in 2019. [3]

The experiments all varied the information job-seekers had about a vacancy when they
viewed it. In total, there were 17 treatment arms across the experiments. On platforms with
vast numbers of users, it is typical to try many variations of a treatment to determine the
best one, given the amount of statistical power available. A random chosen subset of   50%
of job seekers were eligible for each of the experiments. [4] Across the experiments, treatment
cells differed in (a) what information was displayed, (b) how spaced out the information
was (with information only presented for a subset of tiles), (c) the color (grey vs blue) of the
information (to make information more or less salient).

Figure 6 displays how the interface presented to job-seekers is altered by the treatments.
Figure 6a displays one job tile when the number of applicants is less than 5. The color and
the information varies depending on the treatment and the number of applications to the
vacancy. Vacancies that have more than 4 prior applications have messages that display
one of the following texts: `5 - 20 applications', `21 - 50 applications', `50 - 200 applications',
`201+ applications'. Figure 6b displays the control tile. Note that the control tile occupies
lessvertical space than the treatment tile. This will be important for our subsequent results
given the limited screen space available on mobile devices. Figure 6c displays how each tile
is combined in the JOF product.

Given the number of treatments available, we primarily analyze the experiment by pool-
ing similar treatment arms. This allows us to simplify the exposition and increase our sta-
tistical power. Appendix B demonstrates that the speci c manner in which a particular
intervention was implemented within an experiment was not of rst order importance to
the treatment effect on applications. This appendix also demonstrates that the experiments
are balanced on pre-treatment covariates—indicative of a successful randomization.

---

[3]The authors of this paper provided substantial input into the design of these experiments. Experiment I was
conducted from 2019-03-26 to 2019-05-09 (44 days). Experiment II was conducted from 2019-05-31 to 2019-06-28
(28 days). Experiment III was conducted from 2019-07-22 to 2019-08-18 (27 days)

[4]Tech companies often allocate only part of the universe of users for an experiment in order to isolate the effects
of potentially interacting concurrent experiments and in order to mitigate risk (Bakshy et al. (2014)).

Figure 6: Illustration of popularity information shown vacancy tiles

(a) Notice on a job tile that a job has few applicants (in blue)

(b) Control presentation of a job tile

(c) Bin applicant count information

Notes: Job vacancy tile interfaces.

6. Effects of information about the number of prior applicants

Directly displaying information about the number of prior applicant increases applications and redistributes them towards relatively under-subscribed vacancies. The effect of the treatment on applications comes from information displayed about a particular vacancy. Spillovers onto vacancies for which information is not shown are minimal.

6.1 Overall job application intensity

We begin by analyzing the aggregate job search effects of the pooled treatment before discussing its heterogeneous effects across vacancies. We estimate these effects by running regressions of the following form:

$$Y_i = \beta_{exp} + \beta_1 Treat_i + \epsilon_i \tag{3}$$

where Y_i refers to outcomes for searcher i observed in the experiment and $Treat_i$ refers to an indicator variable for whether the searcher was in the treatment group that provided information on the number of prior applicants. We also include xed effects, β_{exp} for the experiment number (1 - 3). To get an effect in terms of percents, we take a ratio of β_1 and the mean of Y. Figure 7 plots the treatment effects for the main variables of interest calculated across from the pooled sample consisting of 29,375,533 observations.

Total applications increase by 0.59% and the share of searchers with at least one applications increases by 0.35%. This demonstrates that the treatment effects are coming from both the extensive and the intensive margin. In contrast to the effect on total applications, we nd relatively large decreases in the number of views and detail views. These decreases in views are mostly a mechanical consequence of the fact that the information provided by the treatment takes up more space in the interface. In Appendix B, we demonstrate this fact by showing that the treatment effect size is correlated with how frequently information is shown in a given treatment. The effect on the number of search sessions is less pronounced at -0.2%, which may explain why we nonetheless see increases in overall applications.

6.2 Redirection of applications to vacancies with fewer prior applicants

The largest increases in applications occur for vacancies with few prior applications and there are negative treatment effects for vacancies with many prior applications. To show

Figure 7: Effects of revealing congestion information on job search behavior and outcomes

Notes: This gure plots the average effects in percent terms and their 95% con dence intervals. Each observation is a searcher in an experiment and all treatments that included information about prior applications were pooled. Standard errors are computed via the delta method.

this, we bin outcomes by prior applications at the time of the application in a way that parallels the information treatment. Figure 8 displays the results. We nd positive treatment effects for applications to vacancies with fewer than 5 prior applications. We nd negative treatment effects for vacancies with more than 20 applications and in the case of the 51 - 200 prior applications bin, this effect is statistically signi cant. This is consistent with out theoretical framework, which predicts a redistribution of applications from vacancies that get many applications to those who get few applications.

Conditional on viewing a vacancy, treatment users are more likely to click on the vacancy and apply to it. Figure 9 shows that this effect is especially large for vacancies with fewer than 5 prior applications. This supports our theory that searchers prefer vacancies that have few applications. [5]

6.3 Spillovers

Finally, we consider whether the effect of the treatment comes solely from the information about a particular vacancy or whether there are spillovers onto vacancies for which no information is provided. A positive signal about a particular vacancy may draw away applications from other vacancies or it may induce additional applications to other positions due to learning. We can study these mechanisms by considering treatments which display

[5]The overall increase in the treatment ratios can be explained by the fact that the treatment reduced views and detail views.

Figure 8: Effects of congestion information on applications, by prior applications to a vacancy

Notes: This gure plots the average effects in percent terms and their 95% con dence intervals. Each observation is a searcher in an experiment and all treatments that included information about prior applications were pooled. Standard errors are computed via the delta method.

Figure 9: Effects on detail views / views by prior applications to a vacancy

Notes: This gure plots the average effects in percent terms and their 95% con dence intervals. The outcome is the ratio of detail views to views. Each observation is a searcher in an experiment with at least one view or detail view for a job in the bin and all treatments that included information about prior applications were pooled. Standard errors are computed via the delta method.

information every three tiles rather than every tile. If information causes substitution, then we should see that the treatment has negative effects on vacancies that appear on tiles that are not multiples of three. On the other hand, if there is learning, we should see positive effects for these tiles.

Figure 10 plots the treatment effects on applications based on the position in which they were shown. The estimates are pooled across two treatment arms for which information is displayed every third tile and only when the vacancy has fewer than ve applications. We see a positive and statistically signi cant effect for vacancies in a position divisible by 3. In contrast, for vacancies in other positions, we see a negative but small in magnitude and insigni cant effect. As a result, we conclude that the applications induced by the treatment come from the information learned about particular vacancies and that spillovers onto other vacancies are minimal.

Figure 10: Effects of congestion information on applications, by position of vacancy

Notes: This gure plots the average effects in percent terms and their 95% con dence intervals. The outcome is the number of applications. Each observation is a searcher in a treatment arm where information is displayed every 3 tiles or in the control. Standard errors are computed via the delta method.

7.   Effects on application outcomes

Our theory predicts that searchers in the treatment will, on average, apply to vacancies with less competition. However, we do not have unambiguous predictions about the quality of applications in the treatment vs the control. We study these questions empirically by comparing the applications submitted by treated and control users. We focus on the treatments in experiment 3, since they had the largest effects on application rates and give us the high-

Table 1: Application Order and Characteristics

	Log(Order)	Log(Order)	Log(Eventual)	Third Party	Viewed
Treatment	-0.023***	-0.005***	-0.020***	0.003***	
	(0.002)	(0.001)	(0.001)	(0.000)	
Log(Order)					-0.066***
					(0.000)
N	6070247	6070247	6070247	6070247	5472586
Vacancy FE	No	Yes	No	No	No

Notes: This table contains results for a linear regression of applications outcomes on the treatment in experiment 3. `Order' refers to the order in which the application arrived and `Eventual' is the cumulative applications ever received by a vacancy. `Third party' refers to a vacancy syndicated from a third party platform.

est possibility of detecting effects on application outcomes.

Treated applicants experience a lower degree of competition. Column 1 of Table 1 reports the results of a regression of the log of the application order on the treatment. We see that treated applications have an order of submission that is 2% lower. This may be due to changes in the types of vacancies applied to or changes in the speed with which applications are sent. Column 2 informs this by reporting results from the same regression but with vacancy xed effects. Conditional on a vacancy, treated applications are submitted .4% faster. This demonstrates that most of the effect of the treatment on speed is due to applications to different vacancies. We investigate this further in columns 3 and 4, which show the estimates from regressions of the eventual applications received by vacancies and whether vacancies were syndicated from a third party. We see that treated applications are sent to vacancies that nish with fewer applications, suggesting that they are less popular. We also see that treated applications are more likely to be sent to third party applications, suggesting that information about competition is especially important for these vacancies.

Given the fact that treated applications arrive earlier to vacancies with less competition, we expect them to have better outcomes. We can conduct a back of the envelope calculation. Table 1 shows that treated applications arrive 2.3% earlier, and that a 1% increase in application order decreases the likelihood of a view by 6.6% (Column 5). Multiplying these two together will get us the expected increase in view probabilities in the treatment of 1.5%. However, this back of the envelope calculation may not be valid if the treatment also changes the types of applications submitted, which is what we investigate next.

Table 2: Differences in Application Outcomes

	View	Contact	Interview
Treatment	0.001**	0.000	0.000
	(0.001)	(0.001)	(0.000)
Constant	0.576***	0.381***	0.021***
	(0.000)	(0.000)	(0.000)
N	5472586	5472586	5472586

Notes: This table contains results for a linear regression of applications outcomes on the treatment in experiment 3. `Viewed' is an indicator whether the employer viewed the application, `Contact' is an indicator for whether an employer sent an applicant a message, and `Interview' is an indicator for whether an employer marked that an interview was conducted.

Table 2 displays outcomes across treated and control applications for applications sent to native vacancies, which are the ones for which we can measure outcomes. Column 1 demonstrates that treated applications to native vacancies are more likely to be viewed by the employers. However, columns 2 and 3 show that we do not detect any effects on the rates at which employers contact or interview applications. This may be due simply to a lack of statistical power or due to the fact that marginal applicants from the treatment are adversely selected. We nd some evidence of selection, although this effect is non-monotonic in application order (Table C.3).

To summarize, we nd that treated applications face less competition, and that this is a function of both applying to different vacancies and to applying to applying earlier. We nd that treated applications are more likely to be viewed by native employers, but nd no substantial differences between treated and control users in terms of other outcomes.

8.   Vacancy age effects

In conventional job markets, job-seekers can potentially know the age of a vacancy because it is provided directly by the platform. To understand how this information is used, we show how job-seekers without access to this vacancy age information responded.

Before we present these results, it is useful to step back and consider what we know from the other cells. To summarize, we have shown that job seekers value information about the level of competition when applying for vacancies and respond by applying to vacancies with fewer prior applications. We now investigate whether searchers use other information as a

Figure 11: Effects of vacancy age at the time of application on employer responses

Notes: Each point represents the estimated effect and each line present the 95% con dence interval for estimates of the effects of vacancy age at the time of application on whether an application is viewed, responded to, or results in an interview. Standard errors clustered at a vacancy level.

proxy for vacancy competition.

We nd that job searchers use the vacancy age as a proxy for competition when other information is not available and that, when vacancy age information exists, competition information helps to attract applicants to older vacancies.

We begin by showing that vacancy age affects the probability of responses from employers. To measure this, we plot the estimated coef cients for vacancy age from Equation 1 in Figure 11. We nd that, even conditional on the order of the application, vacancy age is negatively correlated with employer responses, whether these are measured by views, contacts, or interviews. Although the coef cients are not as large in magnitude as those we found for application order, they are nonetheless statistically signi cant.

Given that searchers should care about vacancy age, we now study what happens when vacancy age is not displayed but everything else is held constant. This was the case in one treatment arm of Experiment I. In the control group, the tile displayed `New' in blue if the vacancy was 5 or fewer days old, and would display `xd' in grey otherwise, where `x' is the number of days the vacancy has been posted (See Figure 1).

Figure 12 displays the overall effects of removing vacancy age. We see that treated users submitted fewer applications and clicked on fewer vacancies. These effects are of comparable magnitude to the effects of including information on prior applicants. Consequently, searchers are using vacancy age to choose to use the platform and to direct their search.

Figure 12: Treatment effects on job search outcome
Removing vacancy age

Notes: This figure plots the average effects in percent terms and their 95% confidence intervals. Each observation is a searcher in experiment I who was either in the control group or in the treatment group for which vacancy age was removed.

We investigate these effects further by calculating the treatment effects based on the age of the vacancy at the time of the view. The left panel of Figure 13 displays the treatment effects split by whether the vacancy had the text `New' or not. We see that removing vacancy age decreased applications and detail views to new vacancies and increased them for vacancies older than five days. This heterogeneous effect suggests that users prefer applying to new vacancies when both new and old vacancies are identifiable, but otherwise cannot perfectly predict vacancy age based on observed information.

Since vacancy age is correlated with the number of prior applications, we consider whether vacancy age information helps direct searchers to vacancies with less competition. The right panel of Figure 13 displays how removing vacancy age affects applications to vacancies with different levels of competition. The treatment causes the largest falls in applications for vacancies with fewer than 5 prior applicants and it has no effects on vacancies that receive over 200 applications.

If vacancy age already allows searchers to find low competition vacancies, then why do the information treatments have an effect? One reason may be that this information allows seekers to identify older vacancies with little competition. Figure 14 shows the treatment effects of competition information for applications to vacancies with different ages. We find that the treatment increases applications to older vacancies, confirming this conjecture.

Figure 13: Treatment effects by job type
Removing vacancy age

Notes: This figure plots the average effects in percent terms and their 95% confidence intervals. Each observation is a searcher in experiment I who was either in the control group or in the treatment group for which vacancy age was removed.

Figure 14: Treatment effects on vacancy age of applications
Competition information

Notes: This figure plots the average effects in percent terms and their 95% confidence intervals. Each observation is a searcher in the experiments who was either in the control group or in the treatment group for which competition information was shown.

9.   Conclusion

Job seekers often lack information about the competition they face for a particular vacancy and their likelihood of getting hired. As a result, they may direct their search in ways that would have been, from their perspective, sub-optimal had they known more information. If signals of job desirability are correlated across searchers, there may be an inef cient imbalance in applications across vacancies. Digital platforms are in the position to correct this imbalance with the right intervention.

We use data from JOF to study the ow of applications across vacancies and how it is affected by information about the degree of competition for a vacancy. We rst show that the degree of competition for a position, as measured by the number of prior applications, affects the likelihood that an employer views, responds to, and interviews a candidate. This suggests that job seekers should value being one of the rst to apply to a job. We also show that there is an imbalance in both applications and exposures across vacancies in the same commuting zone. This means that there is a possibility that some job-seekers can direct their job-search to less competitive vacancies and improve market ef ciency by doing so. We propose a simple model that rationalizes the above facts and has predictions for the role of information in this type of market.

We test our model by conducting three experiments over the course of ve months in 2019. Treated users were shown a simple measure of competition, the number of prior applicants, and this information caused them to redirect their job search. The treatments increased applications to vacancies with little competition and reduced applications to vacancies with a lot of competition. Although our treatments varied the frequency and color with which information was shown, these variations did not have rst order implications for our aggregate effects.

While re-balancing applications towards less competitive vacancies is likely ef cient, applications induced by the information may be worse. We show that the likelihood that an application was responded to or interviewed didn't change between the treatment and control group. This suggests that marginal applicants were not substantially worse than applicants in the control group.

Even in the absence of competition information, applicants were able to direct their search towards less competitive vacancies. We use an experiment to show that this is due to the fact that vacancy age is displayed on the platform and is highly correlated with the number of prior applications. When vacancy age was removed, applications to less competitive vacancies dropped. As a corollary, when we added information about prior applicants,

applications increased precisely to older vacancies with few prior applications.

We studied just a few of the many information design decisions by the platform. For example, the platform could add or remove information or change where and how information is displayed. The platform could also create better signals of competition and match quality and display these signals to job seekers. Information design decisions may also interact in important ways with other platform design decisions such as ranking and user acquisition strategies.

The welfare implications of information design decisions are difﬁcult to study with user level experiments. A key reason for this difﬁculty is that actions by seekers and employers exert externalities on each other, meaning that treatment effects from a user level experiment may not be indicative of market outcomes were everyone to be treated. Another difﬁculty is that the platform has very noisy information about hiring and the quality of those hires. This means there is no direct measure of welfare. Future work may be able to address these limitations with commuting zone level experiments, better measurement, and structural models of job search in digital platforms.

Finally, no single platform has a bird's eye view of the entire labor market. Both searchers and employers multi-home across a variety of platforms. As a result, measures of competition on one platform may not fully reﬂect the true level of competition and optimizations made on one platform might not improve outcomes in the entire labor market. The implications of this fragmentation in labor market platforms are important to understand for market designers and policymakers.

References

Abaluck, Jason and Giovanni Compiani , "A Method to Estimate Discrete Choice Models that is Robust to Consumer Search," Technical Report, National Bureau of Economic Research 2020.

Albrecht, James, Bruno Decreuse, and Susan Vroman , "Directed search with phantom vacancies," 2017.

Azar, José, Steven Berry, and Ioana Elena Marinescu , "Estimating labor market power," Available at SSRN 3456277, 2019.

Baker, Scott R and Andrey Fradkin , "The impact of unemployment insurance on job search: Evidence from Google search data," Review of Economics and Statistics, 2017, 99

(5), 756–768.

Bakshy, Eytan, Dean Eckles, and Michael S Bernstein , "Designing and deploying online eld experiments," in "Proceedings of the 23rd international conference on World wide web" 2014, pp. 283–292.

Barbanchon, Thomas Le, Roland Rathelot, and Alexandra Roulet , "Gender differences in job search: Trading off commute against wage," 2020.

Branco, Fernando, Monic Sun, and J Miguel Villas-Boas , "Optimal search for product information," Management Science, 2012, 58 (11), 2037–2056.

Cheron, Arnaud and Bruno Decreuse , "Matching with phantoms," The Review of Economic Studies, 2017, 84 (3), 1041–1070.

Faberman, R Jason and Marianna Kudlyak , "The intensity of job search and search duration," American Economic Journal: Macroeconomics , 2019, 11 (3), 327–57.

Facebook , Methodology: Commuting Zones 2020.

Fong, Jessica, "Search, Selectivity, and Market Thickness in Two-Sided Markets," Selectivity, and Market Thickness in Two-Sided Markets (September 23, 2019), 2019.

Fradkin, Andrey , "Search, matching, and the role of digital marketplace design in enabling trade: Evidence from airbnb," Matching, and the Role of Digital Marketplace Design in Enabling Trade: Evidence from Airbnb (March 21, 2017) , 2017.

Gardete, Pedro and Megan Hunter , "Guiding Consumers through Lemons and Peaches: An Analysis of the Effects of Search Design Activities," 2020.

Gee, Laura K. , "The more you know: Information effects in job application rates by gender in a large eld experiment," Management Science, 2019, 65.

_ , Jason Jones, and Moira Burke , "Social Networks and Labor Markets: How Strong Ties Relate to Job Finding on Facebook's Social Network," Journal of Labor Economics , 2017, 35 (2), 485–518.

Hodgson, Charles and Gregory Lewis , "You can lead a horse to water: Spatial learning and path dependence in consumer search," 2020.

Horton, John J , "The effects of algorithmic labor market recommendations: Evidence from a eld experiment," Journal of Labor Economics , 2017, 35 (2), 345–385.

Hosios, Arthur J , "On the ef ciency of matching and related models of search and unem-
   ployment," The Review of Economic Studies, 1990, 57 (2), 279–298.

Kroft, Kory and Devin G Pope , "Does online search crowd out traditional search and im-
   prove matching ef ciency? Evidence from Craigslist," Journal of Labor Economics, 2014,
   32 (2), 259–303.

Kuhn, Peter and Hani Mansour , "Is Internet job search still ineffective?," The Economic
   Journal, 2014, 124 (581), 1213–1233.

Marinescu, Ioana , "The general equilibrium impacts of unemployment insurance: Evi-
   dence from a large online job board," Journal of Public Economics, 2017, 150, 14–29.

Marinescu, Ioana Elena and Daphn é Skandalis, "Unemployment insurance and job search
   behavior," Available at SSRN 3303367, 2019.

Muchnik, Lev, Sinan Aral, and Sean J Taylor , "Social in uence bias: A randomized experi-
   ment," Science, 2013, 341 (6146), 647–651.

Petrongolo, Barbara and Christopher A Pissarides , "Looking into the black box: A survey
   of the matching function," Journal of Economic literature, 2001, 39 (2), 390–431.

Salganik, Matthew J, Peter Sheridan Dodds, and Duncan J Watts , "Experimental study of
   inequality and unpredictability in an arti cial cultural market," science, 2006, 311 (5762),
   854–856.

Skandalis, Daphn é, "Breaking News: Information About Firms' Hiring Needs Affects the
   Direction of Job Search," Technical Report, mimeo 2018.

Wright, Randall, Philipp Kircher, Benoit Julien, and Veronica Guerrieri , "Directed Search
   and Competitive Search: A Guided Tour," Journal of Economic Literature, 2019.

A.   Theoretical Appendix

In this section, we derive the proofs for the propositions in section 4.

Proposition 1:   When the worker has no information (NI) about vacancy level exposure, application rates are decreasing in k, decreasing in S, and increasing in v. All vacancies with the same k have the same application rate.

Proof: The searcher must make a decision to apply but does not know e_j or a_j. As a result, they apply if $v - kE[a_j | k_j] > c_{ij}$. We now characterize the expected number of applications conditional on k.

$$\hat{a}(k_j) = E[a(k_j)] = SE[2e_j | Exposure]F(v - k_j E[a(k_j)])$$
$$= \frac{4}{3}SF(v - k_j \hat{a}(k_j)])$$

We can then derive the realized number of applications for vacancy, i, when searchers have no information (NI) about potential congestion as: $a(k_j ; e_j)^{NI} = 2e_j SF(v - k_j \hat{a}(k_j))$

The above equation has a unique fixed point since the left hand side is increasing in a and the right hand side is decreasing in E[a] and we can remove the expectations since there is no randomness at the aggregate level.

Define the application rate as $r_i = \frac{a_i}{2Se_j}$, which represents the share of exposures for job i that result in applications.

$$r_i^{NI}(k_j) = F(v - k_j \hat{a}(k_j)) = F(v - 2Se_j k_j \hat{r}(k_j)) \tag{4}$$

$$\frac{@r^{NI}(k_j)}{@v} = f(\cdot)(1 - 2Se_j k_j \frac{@r^{NI}(k_j)}{@v}) \tag{5}$$

$$\frac{@r^{NI}(k_j)}{@v} = \frac{f(\cdot)}{1 + 2Se_j k_j f(\cdot)} > 0 \tag{6}$$

$$\frac{@r^{NI}(k_j)}{@k} = -2Se_j f(\cdot)(r^{NI}(k_j) + k_j \frac{@r^{NI}(k_j)}{@v}) \tag{7}$$

$$\frac{@r^{NI}(k_j)}{@k} = -\frac{2Se_j f(\cdot)r^{NI}(k_j)}{1 + 2Se_j k_j f(\cdot)} < 0 \tag{8}$$

$$\frac{\partial \tilde{f}^{NI}(k_j)}{\partial S} = 2k_j e_j f(:)(r^{NI}(k_j) + S\frac{\partial \tilde{f}^{NI}(k_j)}{\partial S}) \tag{9}$$

$$\frac{\partial \tilde{f}^{NI}(k_j)}{\partial S} = \frac{2ke_j f(:)r^{NI}(k_j)}{1+2Se_j)k_j f(:)} < 0 \tag{10}$$

Proposition 2: Consider the equilibrium when workers are given information about exposure, e_j. Call this equilibrium I. For vacancies with $e_j < \frac{2}{3}$, applications increase relative to equilibrium NI, and for vacancies with $e_j > \frac{2}{3}$ they decrease. The change in applications is moderated by k, vacancies with low k are less affected by the policy than vacancies with high k. Lastly, the total number of applications decreases and the total welfare in the market increases.

Proof: To see that the number of applications in the two scenarios is equal at $\frac{2}{3}$, notice that the solution to the equation for the expected number of application corresponds to the solution to the number of applications when $e_j = \frac{2}{3}$.

$$\hat{a}(k_j) = E[a(k_j)] = SE[e_j jExposure]F(v \quad k_j E[a(k_j)])$$

$$= \frac{4}{3}SF(v \quad k_j \hat{a}(k_j)])$$

$$a(k_j; e_j)^{NI} = 2e_j SF(v \quad k_j \hat{a}(k_j))$$

$$a(k_j; e_j)^{I} = 2e_j SF(v \quad k_j a(k_j; e_j)^{I})$$

Now consider the difference in application rates and take the derivative. Notice that the application rate under the NI scenario is not a function of e_j since searchers do not know e_j.

$$(k_j; e_j) = r^{I}(k_j; e_j) \quad r^{NI}(k_j)$$

$$= F(v \quad k_j a(k_j; e_j)^{I}) \quad F(v \quad k_j \hat{a}(k_j)))$$

$$\frac{\partial (k_j; e_j)}{\partial \rho} = S(f(v \quad k_j a(k_j; e_j)^{I}))\frac{\partial (k_j; e_j)^{I}}{\partial \rho} < 0 \tag{11}$$

Proposition 3: When searchers overestimate the number of other searchers in the market, revealing information about the true number of searchers in addition to exposure information increases application rates.
Proof:
Denote $\hat{S}$ as the belief about the number of searchers and $S < \hat{S}$ as the true number of

searchers. The number of applications conditional on $(\hat{S}; S; e_j)$ is then determined by the equations below. The derivative of applications with respect to believed searchers is negative, which implies that if the belief is adjusted downward, the number of applications increases.

$$\hat{a}(k_j; e_j)^I = 2e_j \hat{S}F(v - k_j \hat{a}(k_j; e_j; \hat{S})^I)$$

$$\frac{@\hat{a}^I}{@\hat{S}} = \frac{2e_j F(:)}{1 + 2e_j k_j f(:)} > 0 \tag{12}$$

$$a(k_j; e_j)^I = 2e_j SF(v - k_j \hat{a}(k_j; e_j; \hat{S})^I)$$

$$\frac{@a}{@\hat{S}} = -2e_j Sk_j f(:)\frac{@\hat{a}^I}{@\hat{S}} < 0 \tag{13}$$

$$\tag{14}$$

B.   Description and analysis by each treatment arm.

In this section, we describe and analyze the treatment arms of each of the three experiments. We begin with by describing the set of treatments tried across all three experiments. Figure B.1 shows the set of treatment parameters and the experiments which contained an arm with those parameters.

The rst dimension along which treatments differed was in whether every vacancy tile was eligible to show competition information. Column 1 contains the set of arms where information could be shown on every tile, while columns 2 and 3 contain arms where information could be shown either every 3 tiles or every 10 tiles (beginning with the rst tile on the screen). Next, row 1 displays the set of treatments where information about competition was shown only for vacancies that had fewer than 5 prior applications. For these vacancies, the text `Be one of the  rst to apply' was displayed.  [6]  Row 2 displays the set of treatment arms for which competition information could also be shown for vacancies with more than 4 applications. For these vacancies, the following text could be shown, where appropriate: `Be one of the rst to apply', `5 - 20 applications', `21 - 50 applications', `50 - 200 applications', `201+ applications'. Finally, treatment arms differed by whether they displayed this information in blue (vs grey) always (`All'), just for vacancies with $< 5$ applications (`First'), or never ('None'). Finally, the grid excludes one treatment arm from Experiment 3, in which some signals were eligible to be shown every tile, while those relating to vacancies with < 5 vacancies could only be shown on every third tile.

---

[6]The text was also translated into the appropriate language for each locale.

Figure B.1: Treatment Arms for Experiments

Notes: This figure displays the experiments during which each combination of treatments appeared. Information was presented either every tile, every 3 tiles (starting with tile 1), or every 10 tiles (starting with tile 10). Treatment arms varied by whether only under-subscribed vacancies (< 5 prior applications) were marked with competition information, or whether all eligible vacancies were marked with congestion information. Lastly, in certain cases competition information was given a blue color. Values of `First' in the `Blue' column denote that only signals for under-subscribed vacancies were given a blue color.

Experiment 1 also contained three treatment arms that varied whether the vacancy age was displayed. One of these arms was discussed in section 8. Two other arms removed vacancy age, but added competition signals (either just `Be the first to apply' or all congestion signals).

To check that the randomization was properly conducted, we preformed a set of balance tests. Figure B.2 displays these tests, where the p-value for the difference in means on pre-treatment covariates is displayed for every treatment arm in each experiment. Across four covariates (Age, Android User, Gender, and US user), we find differences that are not statistically significant at a 5% p-value. This evidence suggests a proper randomization of the treatment arms by Facebook.

B.1   Effects by Treatment Arm

Next, we discuss the by-arm treatment effects for each treatment and experiment. We begin with Experiment (Figure B.3). Columns 1, 4, and 5 of of the figure plot the treatment effects where the vacancy age is hidden. Columns 2 - 5 plot treatments where competition infor-

Figure B.2: Covariate balance test p-values across experiments

mation is added. Broadly, the treatments where vacancy age is hidden experience drops in views, detail views and applications. Columns 2 and 3, where competition information is added but vacancy age remains. The two treatments have similar effects on our outcomes.

Figure B.4 displays the effects of the separate treatment arms of experiment 2. Broadly, the effects are of similar magnitude across arms. The clearest difference is that there is a bigger drop in views when competition information is displayed every 3 tiles rather than every 10 tiles. This is expected since the competition information takes up an additional line of text and therefore fewer vacancies can be shown in the 'every 3' treatments.

Finally, Figure B.5 displays the effects of the separate treatment arms of experiment 3. As in the other experiments, the effects on applications and sessions are similar across treatment arms. As before, the more frequently competition information is shown, the fewer vacancies are seen by the searchers.

Figure B.4: Treatment Effects for Experiment 2

Figure B.5: Treatment Effects for Experiment 3
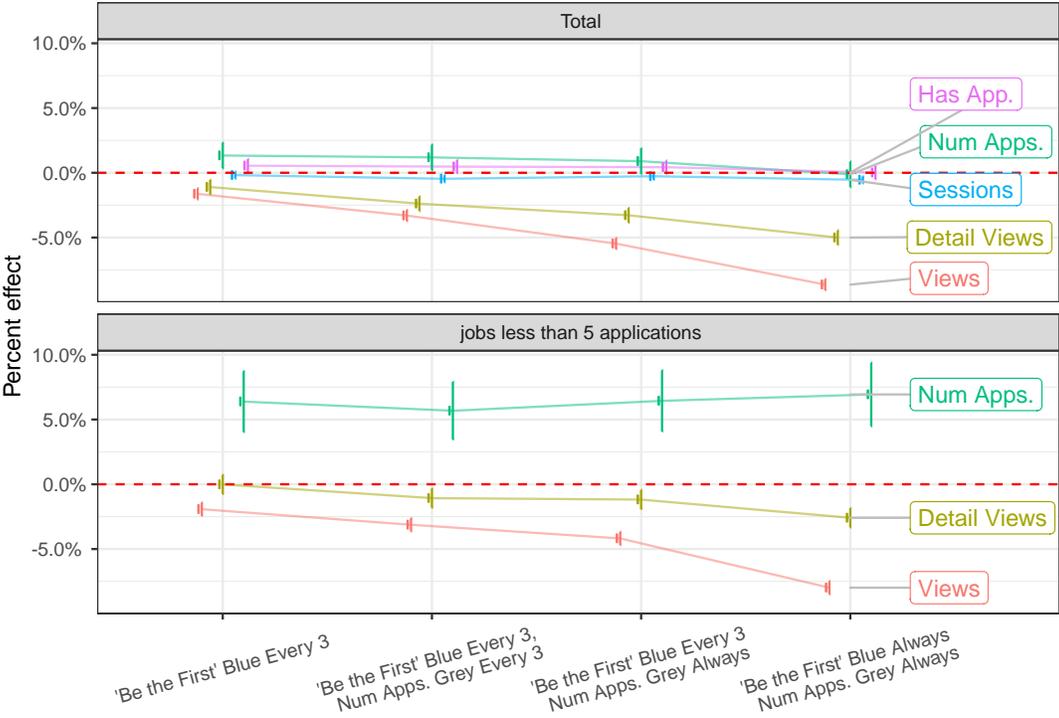
C.   Why does the effect size vary across experiments?

We now investigate why the effects of competition information on applications vary so greatly across the three experiments. We show that the details of the treatment implementation, changes in the demographics of users, and changes in market tightness do not explain the differences in treatment effects.

C.1   Differences in treatment

As explained in Appendix B, each of our three experiments had several treatment variations. One concern is that our main results are driven by differences in the exact implementation of the treatment across experiments. In this section, we compare two *identical* treatment arms across experiments 2 and 3 and show that the differences in experimental treatment effects persist even for identical treatments.

The first repeated treatment is one in which the 'Be one of the first to apply' signal is eligible to be shown in blue every three tiles. The estimates and 95% confidence intervals for this treatment are shown in Comparison A of Figure C.1. The effect of the treatment on applications to under-subscribed vacancies is more than twice as large in experiment 3 than it is in experiment 2 (pval: 0.016). There are also substantial differences in other outcomes, such as the number of detail views and views for all vacancies.

Similarly, there are differences in the effects of the other repeated treatment between experiment 2 and 3. This treatment displayed competition information every 3rd tile for all types of information. Furthermore, 'Be one of the first to apply' is shown in blue. The effect of the treatment on applications to under-subscribed vacancies is more than twice as large in experiment 3 than it is in experiment 2 (pval: 0.07). There are also substantial differences in other outcomes, such as the number of detail views and views for all vacancies. As a result, we conclude that the differences in experiments are not driven by the specific implementation of the competition signal.

C.2   Differences in observable user characteristics and market conditions.

Another reason for the differences in treatment effects across experiments may be that the user composition or market conditions are changing. JOF is a fast growing and global platform, so it is conceivable that these factors could change over a period as short as a month.

Table C.1 reports summary statistics for user characteristics for the three experiments.

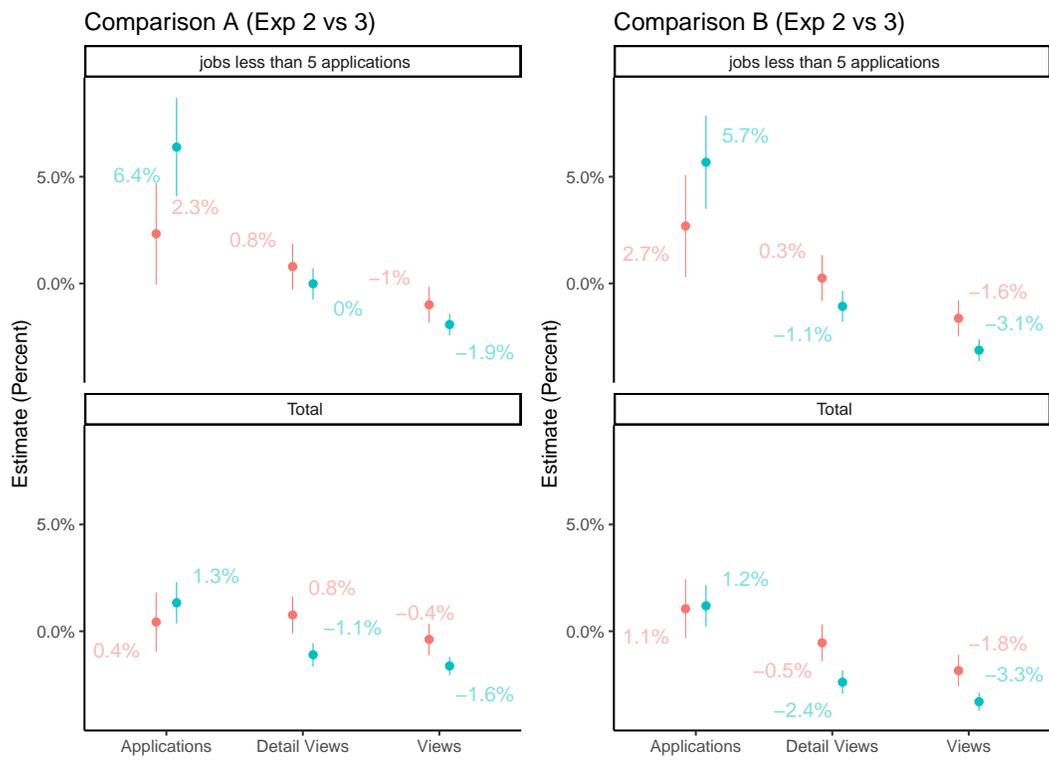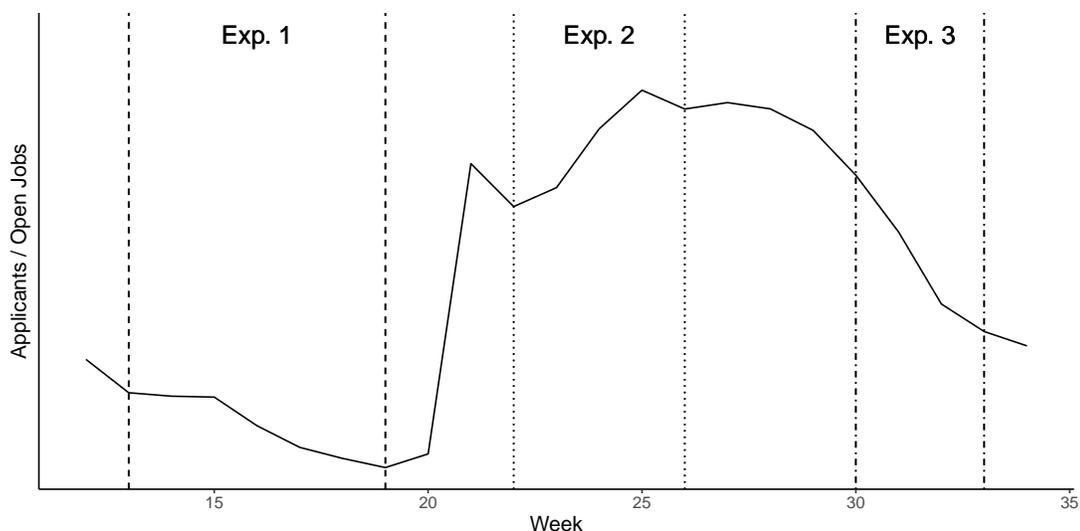Figure C.1: Effects of the same treatment across experiments

Figure C.2: Evolution of market tightness over time

There are some compositional differences across the experiments—for example, by Experiment III, the fraction of users who are from the US has declined, as has the fraction that are female. Furthermore, Experiment III has a lower share of users who had used the Jobs product in the two weeks prior to the experiment than Experiment II. We can also measure the market tightness of each commuting zone in our sample - defined by the prior week's number of applications divided by number of vacancies. Figure C.2 plots the evolution of this quantity over time and by region. We see that tightness increases after Experiment I and falls after Experiment II.

Next, we test for heterogeneous effects based on these factors, and find that they are not large enough to explain the differences between experiments. We estimate separate regressions interacting a dichotomized version of each variable with the treatment, where the outcome variable is applications to under-subscribed jobs. The results of these regressions are reported in Figure C.3. We see that there is some heterogeneity in treatment effects for those who've used the product before and for US users. However, this heterogeneity is not precisely estimated.

Table C.1: Control group demographics over the three experiments

	25th	Median	75th	Mean	StDEv	exp
Exp I (n = 1,763,735)						
Age	26	33	44	35.95	13.63	1
US User	NA	NA	NA	0.25	NA	1
Friends	219	444	873	746.74	896.28	1
Used Jobs Pre Exp.	NA	NA	NA	0.27	NA	1
iOS User	NA	NA	NA	0.34	NA	1
Male	NA	NA	NA	0.59	NA	1
Exp II (n = 863,215)						
Age	25	32	42	34.45	13.18	2
US User	NA	NA	NA	0.21	NA	2
Friends	203	437	922	785.59	969.06	2
Used Jobs Pre Exp.	NA	NA	NA	0.38	NA	2
iOS User	NA	NA	NA	0.27	NA	2
Male	NA	NA	NA	0.54	NA	2
Exp III (n = 3,265,172)						
Age	24	31	42	34.23	13.74	3
US User	NA	NA	NA	0.20	NA	3
Friends	173	399	873	747.61	962.37	3
Used Jobs Pre Exp.	NA	NA	NA	0.32	NA	3
iOS User	NA	NA	NA	0.31	NA	3
Male	NA	NA	NA	0.51	NA	3

Notes: User characteristics by experiment. 'Used Jobs Pre Exp.' is an indicator for whether the user used the job board in the two weeks prior to the experiment.

Figure C.3: Heterogeneous Treatment Effects - Applications to Under-subscribed Vacancies
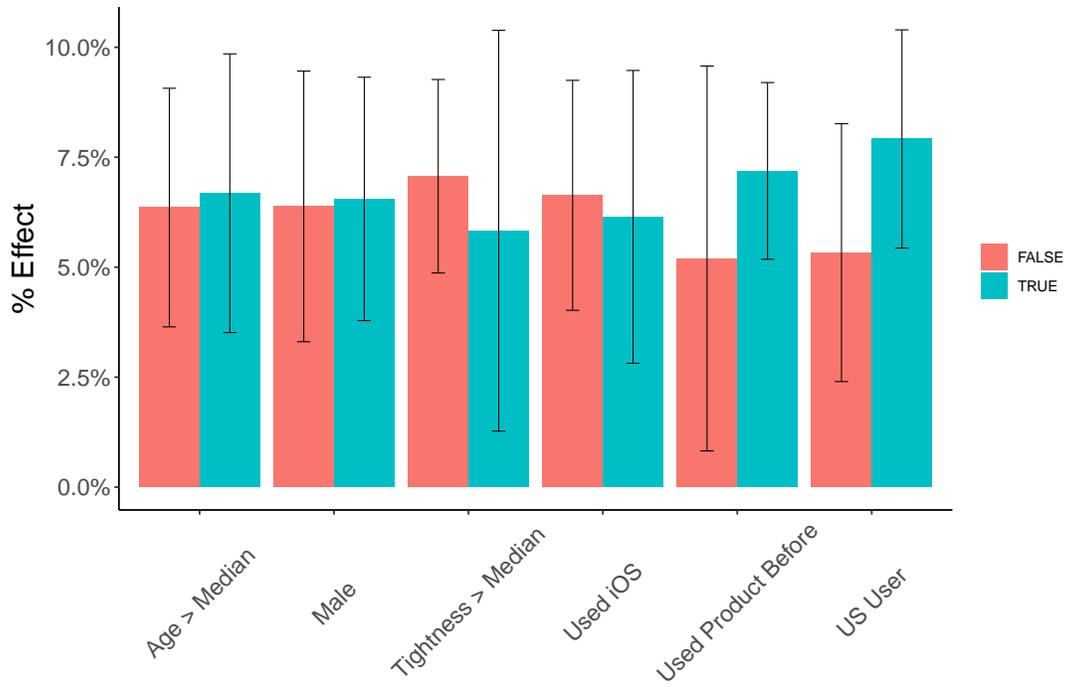

Table C.3: Differences Conditional on Application Order

	Viewed	Contact	Interview
Treatment	0.001	0.002*	0.000
	(0.001)	(0.001)	(0.000)
Log(Order)	-0.018***	-0.004***	-0.003***
	(0.000)	(0.000)	(0.000)
Treat * Log(Order)	0.000	-0.001**	0.000
	(0.000)	(0.000)	(0.000)
N	5472586	5472586	5472586
Vacancy FE	X	X	X

Notes: This table contains results for a linear regression of applications outcomes on the treatment in experiment 3. 'Viewed' is an indicator whether the employer viewed the application, 'Contact' is an indicator for whether an employer sent an applicant a message, and 'Interview' is an indicator for whether an employer marked that an interview was conducted.

C.3 Differences in viewed vacancies across experiments

We now consider whether differences in the vacancies shown to job seekers can explain the differential effects of the experiments. As in the proceeding section, we first measure whether there is any shift in vacancy characteristics over time. We then test whether there are heterogeneous treatment effects as a result of these characteristics.

Figure C.4 displays the average values of our measured variables across the three experiments (colors) and display position in the interface (x-axis). Our first measure of job characteristics is whether the candidate can apply to the job through Facebook. We see that for all experiments, the share of exposures to these jobs is above 85%, although there are some differences. Our second measure is whether the job was created through Facebook's portal (native) or whether it was syndicated from a third-party. We see that most vacancies are native and that there are fewer third-party vacancies in experiment 1. Our next two-measures incorporate information that occurs on the platform. Namely, we can observe the prior number of applications and vacancy age at the time of viewing. We see some differences, but these are not monotonic in experiment launch date. In particular, Experiment II is less likely to have under-subscribed vacancies be viewed, while Experiment I has a lower vacancy age exposed.

Next, we consider whether these characteristics can explain treatment effect heterogeneity in Experiment III. Table C.4 displays regressions where the outcome is whether the user applied to the first vacancy viewed in the experiment. We exclude vacancies shown in positions below the first because the depth of search is an endogenous outcome and we exclude vacancies that could not be applied to through Facebook.[7]

The explanatory variables are vacancy and vacancy by user characteristics. Column (1) displays the results while including the variables considered in Figure C.4. We see that these variables do not explain treatment effect heterogeneity. Column (2) further adds the characteristics of the vacancy poster - in particular whether the poster had a 'local' page on Facebook and whether they had a 'business' page on Facebook. It also adds an indicator for whether the vacancy and the searcher are located in the same city. The interactions are not statistically significant for any of these variables.

Finally, in column (3), we add the algorithmic score which predicts whether a given searcher will apply to a particular job. This score is generated using a proprietary machine learning algorithm and does substantially predict treatment effect heterogeneity. Furthermore, this effect is an order of magnitude larger than the average treatment effect. This het-

---

[7]This designation sometimes changed over time and a negligible fraction of these jobs did have an application.

Figure C.4: Characteristics of Vacancies Viewed Across Experiments
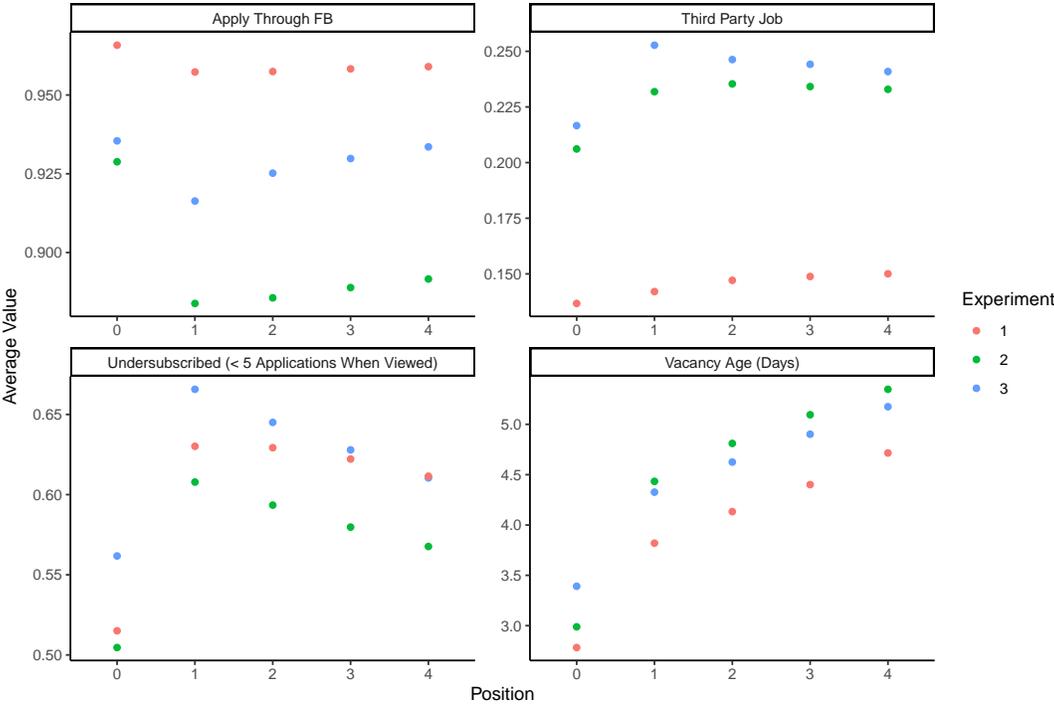
erogeneity demonstrates that there is scope for the match between searcher and vacancy to change over time and to cause treatment effect heterogeneity.

Table C.4: Characteristics of Vacancies Viewed Across Experiments

	Has Application		
	(1)	(2)	(3)
Treatment	0.0004***	-0.0002	-0.0006
	(1e-04)	(4e-04)	(0.0004)
Treatment * Third Party	0.0000	0.0001	0.0001
	(1e-04)	(1e-04)	(0.0001)
Treatment * Age > 5 Days	0.0000	0.0000	0.0002
	(1e-04)	(1e-04)	(0.0001)
Treatment * Same City		0.0001	0.0001
		(1e-04)	(0.0001)
Treatment * FB Local		0.0000	0.0001
		(1e-04)	(0.0001)
Treatment * FB Business		0.0005	0.0006*
		(4e-04)	(0.0004)
Treatment * Algorithmic Match Score			0.0053***
			(0.0012)
Num.Obs.	9020389	9020389	9020389
R2	0.000	0.000	0.004
R2 Adj.	0.000	0.000	0.004

Notes: Regression of whether the user applied to the first vacancy shown as a function of the treatment and interactions. Non-interacted covariates are not shown. Observations for which the first job shown could not be applied to through Facebook were excluded.