MANAGEMENT SCIENCE
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0025-1909 | EISSN 1526-5501 | 00 | 0000 | 0001

INFORMS
DOI 10.1287/xxxx.0000.0000

Competition Avoidance vs Herding in Job Search: Evidence from Large-scale Field Experiments on an Online Job Board

Andrey Fradkin* Monica Bhole† John J. Horton‡

Information about the number of applicants to a job vacancy might simultaneously signal the degree of competition and vacancy quality. We study how this information affects job search. To do so, we conduct three experiments on a large online job platform in which the treatment varies what information is shown to job seekers. Information about the number of prior applicants to a vacancy increases the number of applications and redirects them to vacancies with few prior applications. Information about vacancy age increases application rates, especially to new vacancies. To further investigate the causal mechanisms, we conduct and analyze a survey choice experiment. We conclude that job seekers prefer to avoid competition rather than use the popularity of a vacancy as a signal of quality.

1. Introduction

People use information about the behavior of others to make decisions about what to consume, where to apply, and how to search. For example, signals of song popularity (downloads) caused participants to listen to more popular songs (Salganik et al. (2006)). Herding behavior influenced by prior popularity signals has also been observed in social media, microlending, crowdfunding, petitions, job acceptance, and online platform adoption.[1] People may follow signals of popularity for various reasons, including because these signals contain information about the quality of an option. However, in labor, housing, and other constrained markets, popularity is also correlated with the degree of competition. Applications to more competitive options are less likely to succeed because the number of slots is fixed and there is an advantage to applying earlier (Van Ours and Ridder (1992)). This creates a tradeoff for individuals—apply to a more popular option, which may be better, but have a lower chance of success.

* Boston University, fradkin@bu.edu

† Meta, mbhole@meta.com

‡ MIT Sloan, John.Joseph.Horton@gmail.com

[1] See Coffman et al. (2017), Gee (2019), Muchnik et al. (2013), Tucker and Zhang (2010), Van de Rijt et al. (2014), and Zhang and Liu (2012).

Platform designers must choose whether and how to display signals of popularity and competition information to their users, and this decision critically depends on how users perceive that information. The provision of information may improve market efficiency if it causes searchers to redirect applications to less competitive options, as predicted by some models of search (Wright et al. (2019)). It may also improve market efficiency if it directs users to higher quality options. On the other hand, information about popularity may be harmful to users if it induces herding and results in wasted search. Which of these effects dominates is unknown, and major job search platforms differ in their information designs.[2] A notable paper by Gee (2019) found that displaying the number of prior applicants did not statistically significantly redirect applications to less competitive vacancies, although it increased overall application rates. We build on this work by conducting three large-scale field experiments on a different job platform.

We investigate how job seekers use information that simultaneously conveys vacancy popularity and competition in the labor market. We find that, when given information on the number of prior applicants to a vacancy, job seekers redirect their search towards vacancies with few prior applications. In a related treatment, we find that job seekers prefer recently posted jobs to jobs that are on the platform for a longer duration, a preference that can be rationalized with competition avoidance. Our main results demonstrate that participants in the labor market prefer to avoid competition, even if it means applying to less popular jobs.

We motivate the paper with a simple horse-race between the competition avoidance and signaling roles of information about the prior number of applicants. If competition avoidance dominates, then searchers will prefer vacancies with few prior applications, holding all else equal. If herding dominates because applications are a signal of quality, then searchers will prefer vacancies with relatively more applications.

To conduct our study, we use three experiments encompassing millions of job seekers conducted on an online job board operated by Meta called 'Jobs on Facebook' (JOF).[3] JOF was global, mainly catering to full-time positions that did not require a college education. Users of the platform saw a list of vacancies and had access to a rich set of filters by which to refine their search. Searchers had the ability to apply to most vacancies using the platform. Consequently, the platform could also directly observe the number of applications sent to a particular vacancy in real-time and display information about the count of prior applications to searchers.

[2] As of 2022, Indeed and Google Jobs do not display the number of other applications on the search page, while LinkedIn does. AngelList Jobs does not display the number of prior applicants but does display that the employer is actively hiring. All platforms display vacancy age, but EconJobMarket also displays the deadline for applications.

[3] In December 2022 Meta made an announcement about phasing out JOF. See here: `https://www.facebook.com/business/help/982945655901961`.

Starting in March 2019 and continuing through August 2019, Facebook conducted three experiments related to our research questions. 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 increased application rates to vacancies with fewer than five prior applications by 3.8%, with a range of .9% to 6.4%, depending on the experiment. In contrast, application rates to vacancies with many prior applications fell. This gradient in treatment effects is consistent with competition avoidance substantially dominating herding as a motive for application decisions.

To better understand the mechanisms behind our findings, we conducted a choice survey in which we asked online panel respondents to make choices over vacancies that vary in their wages and number of prior applications. We find that many respondents preferred to apply to vacancies that had fewer applications, even when the vacancy had a low wage. In contrast, very few respondents preferred a job with a low wage when the number of prior applicants was high. The survey analysis corroborates our field experimental finding that job seekers' responses to competition information are driven by competition avoidance, not herding.

Next, we consider the role of vacancy age in directing job searchers towards vacancies with less competition. Since vacancy age is positively correlated with the number of applications, it is a proxy for competition. Furthermore, since vacancy age information is available to job-seekers in many online settings, it is potentially a driver of directed search. The default JOF interface displayed vacancy age, but we removed this information in one of our experiments. In contrast to our other treatments in which the platform added signals to help direct search, removing vacancy age removed a signal that otherwise could be used to direct search. The experimental removal of the vacancy age helps us understand what job-seekers were trying to accomplish, even if it likely made job-seekers worse off.

Based on the experiment, we can determine that workers used vacancy age to decide whether and where to apply. Job seekers who did not have information about vacancy age clicked on 3% fewer vacancies and sent 1.8% fewer applications. The removal of vacancy age also had distributional effects. Treated users were less likely to apply to new vacancies and were more likely to apply to old vacancies. We also find that removing the vacancy age decreased applications to vacancies with few prior applications. Consequently, it increased the concentration of applications to popular vacancies. Note that the reaction of seekers to vacancy age is consistent with both a competition avoidance effect and also with other mechanisms, such as the desire to avoid vacancies that have already been filled.

Next, we consider the efficiency implications of the prior application treatments. One concern is that that competition avoidance could lead to substantial crowding out of applications, so that the total number of applications falls. We find that the total number of applications *increased* due to the treatment. Therefore, the additional applications to low-competition vacancies did not fully crowd out existing applications. The effects we find are large in the context of A/B tests on digital platforms.[4]

We then consider whether the success rate of applications changed due to the treatment. Given that applications increased, a natural concern is that the per-application success rate could have decreased. For example, searchers may have chosen to send applications to worse matching vacancies, which would have a lower likelihood of hiring. Alternatively, applications in the treatment group should benefit from being earlier. As a result, the sign of the effect of the treatment on application outcomes is theoretically ambiguous. We find that applications in the treatment were not substantially more or less likely to be viewed, contacted, or interviewed. Since the total number of applications increased but the outcomes were not harmed, the treatment likely helped treated applicants and made the market more efficient.

An important implication of our results is that behavior in the job market does not exhibit the type of social contagion based on popularity signals found in other social 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 market design interventions and analyses 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 occurred 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 indirectly influence matching through the information it displays and emphasizes.

In this vein, we build on the paper by Gee (2019), who varied whether the number of 'people who clicked to apply' to a vacancy was shown on the detailed view page on LinkedIn. Our study differs from Gee (2019) in several ways, including the analysis we conduct, the experimental treatment,

[4] For example, Azevedo et al. (2020) conduct a meta-analysis of experiments on the Bing search engine. The mean effect sizes are smaller than .02% across a variety of metrics. The standard deviation of the effect across experiments is just .036% for their main metric. Our effect of 3.8% is well outside of this typical range.

setting and interface, and findings about competition avoidance. We discuss these differences in detail in Section 8. Nonetheless, both our paper and Gee (2019) demonstrate that job seekers pay attention to competition information.

Other papers in the literature have shown that search is directed with respect to vacancy characteristics such as compensation (Belot et al. (2018), Banfi and Villena-Roldan (2019), Flory et al. (2015), Samek (2019)), workplace happiness (Ward (2023)), and employer preferences (Kuhn et al. (2020), Leibbrandt and List (2018), Ibañez and Riener (2018)). A particular focus of this literature has been gender differences in preferences regarding the competitiveness of compensation schemes. We study a different aspect of competitiveness: the amount of competition to get hired. We find no differences between men and women regarding their preferences towards such vacancies, using both field experimental and choice survey evidence.

More generally, 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. For example, Hensvik et al. (2020) show that personalized algorithmic recommendations on a job board similar to ours increase employment.

Several other papers have used data on search in digital labor platforms. Faberman and Kudlyak (2019) and Marinescu (2017) study how search evolves over the course of an unemployment spell, Marinescu (2017), Marinescu and Skandalis (2021), and 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. (2021) 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.

Lastly, our results also contribute to a literature on directed search in labor markets. Competition avoidance due to the signals corresponds to a common feature of directed search models, that workers care about the likelihood that their application is successful (Wright et al. (2019)). 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 their search towards newer postings—a phenomenon that we confirm using experimental variation regarding information about vacancy age.

The rest of the paper is organized as follows. Section 2 describes the empirical context. Section 3 discusses the design of our various interventions. Section 4 reports the effects of the treatment on search behavior. Section 5 discusses our results with respect to vacancy age. Section 6 reports on the differences in application outcomes across treatments. Section 7 demonstrates that applying earlier is beneficial to job-seekers. Section 8 discusses the implications of our findings for market design and relation to related work. Section 9 concludes.

2. Empirical context

JOF was an online job board that operated between 2016 and 2022. The product was global in nature and mainly catered to positions that 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.[5] Employers posted vacancies and job-seekers browsed vacancies and sent applications through the platform. The service was free for both sides, but job-seekers needed a Facebook account. Even before the launch of JOF, there was substantial job-search behavior on Facebook (Gee et al. 2017).

Job-seekers were exposed to JOF via the "News Feed" and via notifications.[6] They could also navigate to JOF by clicking on the "Explore" tab and then clicking on a briefcase icon labeled "Jobs." The JOF interface was similar to other job boards, though most of the job search occurred on mobile devices. That most use occurred on mobile presents opportunities—if the user had enabled location-tracking, vacancies within a given radius could easily be shown—but also challenges, in that there was a constrained space in which to present information.

When looking for work, job-seekers could 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 *status quo* job search interface (the job board)—as we will discuss at length, this presentation was modified by various treatments. As we can see in the figure, for each vacancy, the job-seeker could see the title of the job, whether it was part-time or full-time, the name of the employer, the 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 had to click on the "tile" for that opening. Clicking exposed a "detailed view" of the job that included the full job description written by the employer. It also included an "apply" button that the job-seeker could use to submit an application.

Employers interacted with the platform by posting jobs and reviewing applications. A posted vacancy was automatically live for 30 days, but employers could renew it manually. This resulted in an average duration during our testing period of 42 days.[7] With rare exceptions, vacancies did not have application deadlines that were stated in job descriptions. Employers were able to view applications for a job both on desktop and on mobile. The employer could choose to be alerted via

[5] According to publicly available data, approximately 10% of Facebook's overall user base is US-based. JOF's relatively high penetration in the US is likely due to the fact that JOF first launched in the US prior to being expanded globally (`https://perma.cc/6CAM-EKLR`).

[6] Which users are exposed to the Jobs product in the News Feed is determined by an algorithmic system and is not dependent on the treatment assignment of the experiments discussed in this paper.

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

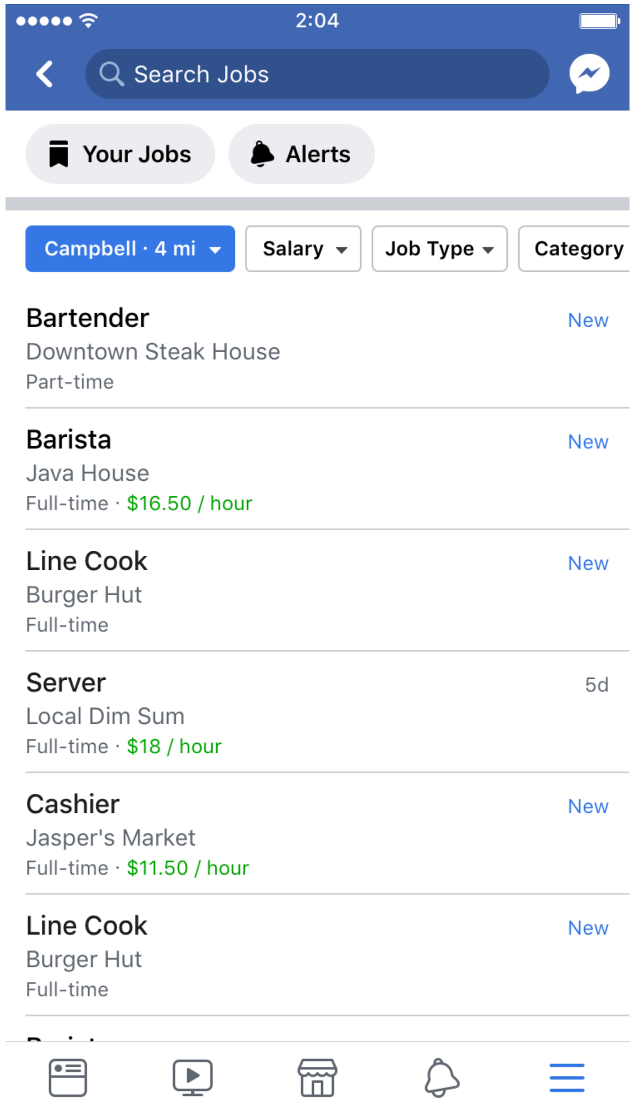

Notes: Interface shown to job-seekers on a mobile device.

messenger and/or via email whenever an application arrived. The employer also had access to an applicant tracking system (ATS), in which the applicants were listed in reverse chronological order by default and which had additional filters available. For each application, employers could send a custom message, send a summary response, or send an interview request.

[7] We calculate this duration for all vacancies created between March 2019 and August 2019. Note that Facebook reviewed job postings and could have removed 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.

2.1. Measurement of job-search behavior and vacancies

We now describe our measurement of vacancies and job-search behavior. We observe the job vacancies a user loaded onto the interface during their search, which is a function of how far they scrolled down the device, their location, and their search query parameters. For each job vacancy, we observe the date it was posted, the date it was closed, and various meta-data both posted by the employer and inferred by algorithms. For example, employers specify job location, job title in the native language, and the opening type (full-time, part-time, contract, internship).[8] We also observe searcher behavior. A "view" occurred when a vacancy tile appeared 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 and when the job-seeker applied to a particular opening.

When job-seekers started an application, their information was populated into an application, using data from their Facebook profile—educational history, past employment, contact information, and so on. Searchers could fill 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 'first party', 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 that 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 platform.

Table 1 displays the characteristics of a 10% sample of vacancies that were posted between March 3 and August 18 of 2019, the period during which our experiments ran. We break out summary statistics by the application flow ('native', which happens through the JOF applicant tracking system, and 'external', which happens fully outside of the platform). The share of vacancies that have a native application flow is 88%. Of these, about a third of vacancies are from the US, and they

[8] However, some important information is missing on JOF and other platforms. For example, we do not observe the number of slots available for a job.

Table 1: Vacancy characteristics

Application Flow		Mean	Median	P25	P75	P95	P99
Native	US Vacancy	0.35	0.00	0.00	1.00	1.00	1.00
	Num. Applications	15	1.00	0.00	8.00	80	196
	Third Party	0.29	0.00	0.00	1.00	1.00	1.00
	Page Age (Years)	7.49	6.25	4.23	11	13	15
	Has Wage	0.28	0.00	0.00	1.00	1.00	1.00
	Full Time	0.77	1.00	1.00	1.00	1.00	1.00
External	US Vacancy	0.87	1.00	1.00	1.00	1.00	1.00
	Page Age (Years)	4.91	4.15	3.75	4.19	13	13
	Has Wage	0.01	0.00	0.00	0.00	0.00	1.00
	Full Time	0.90	1.00	1.00	1.00	1.00	1.00

Notes: This table displays summary statistics for a 10% sample of vacancies between March 3 and August 18 of 2019. Native application flow vacancies are ones that allow seekers to apply through JOF. Third-party vacancies are ones that were syndicated off of JOF. 'Num. applications' is the number of applications that a vacancy has eventually received.

are linked to Facebook pages that have typically been around for about six years. Most vacancies do not have a posted wage and are for full-time positions, and the median vacancy receives one application. Facebook also used algorithms to infer the Standard Occupational Classification (SOC) codes from the job posting. The most common SOC codes were related to sales, driving, and fast food counter workers.[9]

3. Experimental provision of information

Our primary research interest is in testing the relative importance of competition avoidance and herding in job-search. We view this as a horse race between two explanatory factors. If competition avoidance dominates then applicants will prefer to apply to vacancies with few other applicants, all else being equal. If herding dominates, then applicants will prefer to apply to vacancies with more prior applicants all else equal. We formalize the intuition above in Section A.1.

We study this question by analyzing three experiments randomized at the job seeker level, which were conducted over a five-month span in 2019.[10] The authors of this paper provided input into the design of these experiments but these experiments were primarily conducted for the purposes of improving the JOF product. The final decisions regarding which treatment arms to run and when were determined by product managers and designers.

[9] Specifically, the top five 5-digit SOC codes are First-Line Supervisors of Sales Workers (41-101), Marketing and Sales Managers (11-202), Driver/Sales Workers and Truck Drivers (53-303), Sales Representatives, Wholesale and Manufacturing (41-401), Fast Food and Counter Workers (35-302). These comprise approximately 10% of all vacancies.

[10] 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)

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 in addition to a control group. While this is a lot of treatment arms relative to many academic experiments, it is typical at tech platforms to try many minor variations of a treatment to determine the best one (see Kohavi et al. (2020)). The 14 treatment arms relating to just the number of applicants are summarized in Figure A.4.

In experiment 1, there was an arm that showed prior applicant information on every tile. Vacancies that had more than 4 prior applications displayed one of the following labels: '5 - 20 applications', '21 - 50 applications', '50 - 200 applications', '201+ applications'. When the number of prior applicants was less than 5, the text 'Be one of the first' was shown. This labeling scheme was the same across all experiments. Another treatment arm in the first experiment showed just 'Be one of the first' when the number of applicants was fewer than five, but no information otherwise. Experiment 1 also had one treatment arm that just removed the vacancy age, and two treatment arms that removed vacancy age and added prior applicant information.

The second experiment included eight treatment arms which varied the frequency of information shown (every three tiles or every ten tiles) and the color of information (grey or blue). Lastly, the third experiment had four treatment arms that varied information frequency and whether information was shown when the number of prior applications was greater than four. This third experiment was used to make a final decision on how competition information was shown on the platform.

Across all experiments, randomization was conducted at a job seeker level. There were no pre-conditions for eligibility in the experiment other than interacting with the jobs platform. That said, the experiments did not comprise the entire set of eligible users. For each experiment, a randomly chosen subset of $\leq 50\%$ of job seekers were eligible.[11] Note that other aspects of Facebook's systems, such as the ranking algorithm, did not use information about a user's treatment assignment. Any other experiments conducted by JOF were assigned in a matter that was orthogonal to our experiments. Due to this randomization, there were no systematic differences in the types of jobs seen across treatment arms or in the characteristics of users across treatment arms. Also note that since the experiments took place at different points in time, the same user could be in multiple experiments. 83% of users were only observed in one of the three experiments, while 15% were observed in two experiments, and 1.4% were observed in all three. Since randomization was conducted independently across the experiments, we treat each user by experiment observation as independent.

[11] Note that for reasons of confidentiality the company did not permit us to report the exact percentage. 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 2: Illustration of popularity information shown vacancy tiles

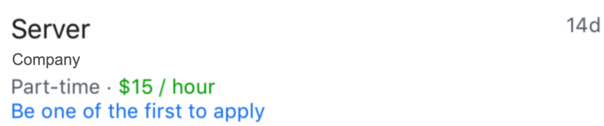

(a) Notice on a job tile that a job has few applicants (in blue)

(b) Control presentation of a job tile

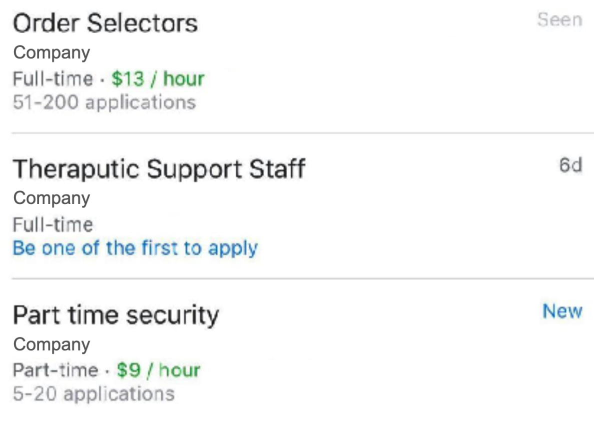

(c) Bin applicant count information

Notes: Job vacancy tile interfaces.

Figure 2 displays how the interface presented to job-seekers was altered by the treatments. Figure 2a displays one job tile when the number of applicants was less than 5. The color and the information varied depending on the treatment and the number of applications to the vacancy. Vacancies that had more than 4 prior applications displayed information in bins: '5 - 20 applications', '21 - 50 applications', '50 - 200 applications', '201+ applications'. Figure 2b displays the control tile. Note that the control tile occupied *less* vertical space than the treatment tile. This will be important for our subsequent results given the limited screen space available on mobile devices. Figure 2c displays how each tile was combined in the JOF product.

In Table A.1 we display the summary statistics of searchers in the control groups of each of the experiments. We see that across all experiments, seekers were mostly outside of the US and Android users. The typical seeker did not send an application or click on a vacancy, reflecting that many seekers were just curious about JOF but were not actively looking.[12] Table A.2 displays seeker characteristics conditional on sending at least one application. We see that those who send an application have broadly similar characteristics to the overall sample, although they tend to be younger and to have more friends. Those who send an application have a median of 6 to 7 detail vacancy views and between 100 and 200 exposures to vacancies in search. The mean number of applications sent for this group is over 10, reflecting a long tail of applications. Lastly, in Table A.3

we show how applications vary in frequency by their prior number of applications in each control group. We find that a similar number of applications are sent across prior application bins in experiment 1, but that in experiments 2 and 3, vacancies with more than 200 applications receive relatively fewer applications.

Given the number of treatments available, we primarily analyze the experiment by pooling similar treatment arms. This allows us to simplify the exposition and increase our statistical power. Appendix A.3 demonstrates that the specific manner in which a particular intervention was implemented within an experiment was not of first-order importance to the treatment effect on applications. Appendix A.3 also demonstrates that the treatment arms are balanced on pre-treatment covariates—indicative of successful randomization. Lastly, Appendix A.4 discusses differences in effect sizes across experiments and heterogeneous treatment effects.

One concern with our experimental design is that there may be violations of the Stable Unit Treatment Value Assumption (SUTVA). In particular, when a treated searcher applies to a vacancy due to the treatment, this may affect the competition faced by subsequent searchers and may affect the job posting behavior of employers who receive applications induced by the treatment. Our experiment is not designed to study these equilibrium spillovers. Instead, we focus on differences in individual job seekers' decisions about which vacancies to view and apply to. Note that at any given time period in our sample, the proportion of treated and control seekers who enter the platform is similar, so that any differences in their behavior are explained by the treatment. That said, our estimates represent seeker behavior under the market conditions observed during our experiments and may change under alternative market conditions.

4. Competition avoidance dominates herding: evidence from treatments about prior applications.

Directly displaying information about the number of prior applicants increases applications and redistributes them toward relatively under-subscribed vacancies. This is consistent with competition avoidance and not herding.

4.1. Overall job search 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 at a job seeker level, since the unit of randomization is also a job seeker. The estimating equation takes the following form:

$$Y_s = \gamma_{exp} + \beta_1 Treat_s + \epsilon_s \tag{1}$$

[12] In terms of representativeness, we can focus on the US subsample. US users tended to be younger (median age 35 to 39 across experiments) and more male (64% to 67% across experiments) than the overall US labor force. They had a similar usage of Apple devices (close to 50%).

where Y_s refers to outcomes for searcher, s, observed in the experiment and $Treat_s$ 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 fixed effects, γ_{exp}, for the experiment number (1 - 3). To get an effect in terms of percentages, we take a ratio of β_1 and the mean of Y in the control group. Figure 3 plots the treatment effects and standard errors[13] for the main variables of interest calculated across from the pooled sample consisting of 29,375,231 observations and Table A.4 displays the regression results in table form.

Total applications increase by 0.59% and the share of searchers with at least one application 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 find 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.[14] In Appendix A.3, we demonstrate this fact by showing that the treatment effect on views 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.

4.2. Evidence on competition avoidance

We now test whether job seekers respond to information by avoiding competition or herding. To do so, we consider the effect of the treatment on applications to jobs with differing amounts of prior applications. We find that the largest increases in applications occur for vacancies with few prior applications and that there are negative treatment effects for vacancies with many prior applications, although these are not precisely estimated. Our results favor the competition avoidance mechanism and not the herding mechanism.

Our estimation strategy follows Equation 1 with the outcome, $Y_{s,b}$, equal to the number of applications sent by seeker s to jobs with a number of applications at the time of exposure in a particular range. For example, if a job seeker sends two applications to vacancies that had fewer than five applications at the time, $Y_{s,<5} = 2$. We bin outcomes in a way that parallels the information treatment. Figure 4 displays the results.

[13] Standard errors for this object are calculated via the delta method. We considered using randomization inference but the computational costs were high with our large sample size and the bias of the asymptotic standard errors is likely to be low with a large sample.

[14] To see that the negative effect on views is plausible, suppose that a mobile phone screen can fit four vacancies on average and that each vacancy takes up three lines without the treatment. Adding one extra line takes up an extra $1 - (4*3)/(4*3+1) = 7.7\%$ space. Given the diversity of mobile phones and differences in search activity, it is plausible that this extra line can affect whether a vacancy is viewed.

Figure 3: Effects of revealing competition information on job search behavior and outcomes

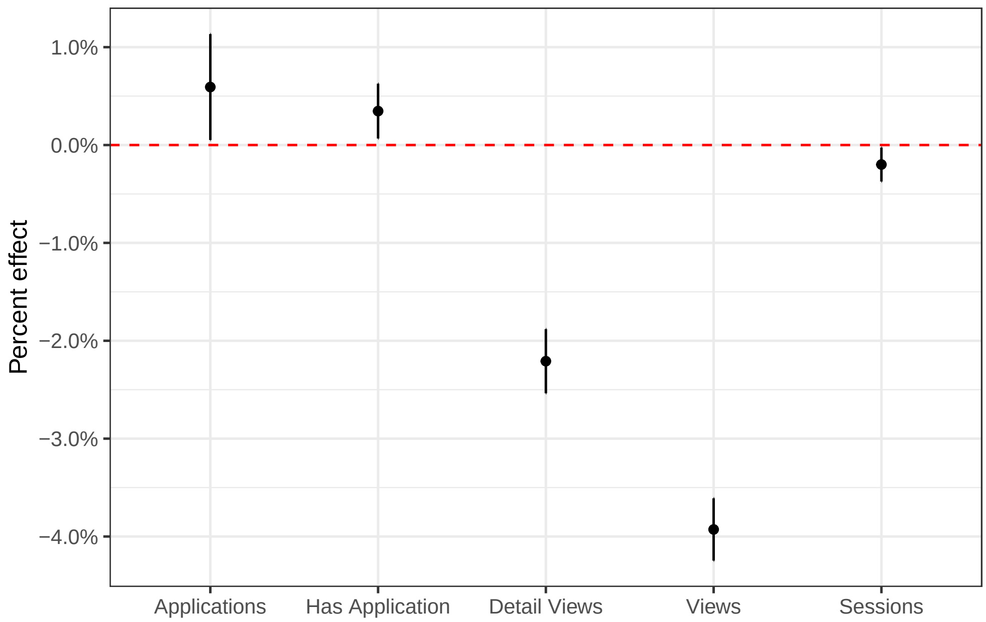

Notes: This figure plots the average effects in percent terms and their 95% confidence intervals. Each observation is a searcher in an experiment and all treatments that included information about prior applications were pooled. Standard errors are computed via the delta method. Table A.4 displays the regression estimates used to generate this figure and Figure A.2 displays the effects in levels.

A striking pattern in Figure 4 is that the effect on applications is largest for vacancies with fewer than five applications. This is consistent with competition avoidance being focused on being one of the first applicants. We provide evidence for why this may be rational in Section 7.

We do not see any difference in application rates for vacancies with 5-20, 21-50, or 201+ prior applications. Suggesting either that seekers on average do not update based on these application counts, or that some applicants update positively while some update negatively. We also find a small negative effect on application to vacancies that have 51-200 prior applications. Given the number of hypothesis tests we conduct, we do not place that much weight on this evidence.

One concern with the above results is that they may be driven by a particular treatment arm. However, this is not the case. The effect on competition information is not driven by the frequency with which information is shown or by whether information about other application bins is shown. Figure A.7 shows that the effect on the < 5 category is similarly sized for seven different treatment arms.

We also consider whether competition avoidance is driven by particular types of jobs or seekers. In Figure A.11, we investigate heterogeneity by a number of factors including gender, age, and device, and fail to detect statistically significant differences. We also fail to find heterogeneity with regards to the skill requirements of the vacancy (Table A.6). To summarize, we fail to detect meaningful heterogeneity in effect sizes.

Figure 4: Effects of competition information on applications to different status vacancies

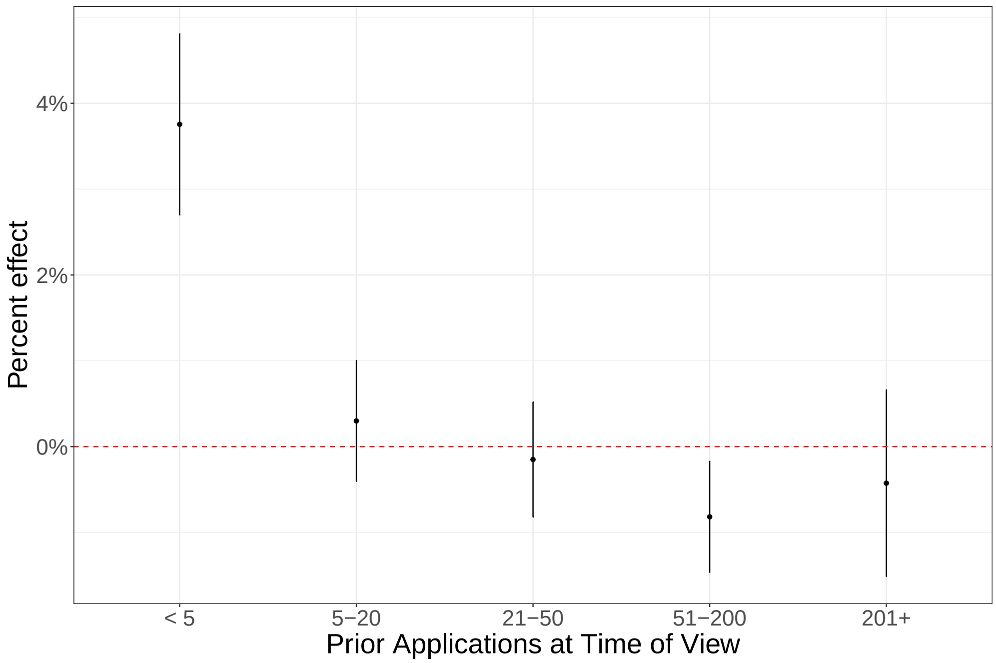

Notes: This figure plots the average effects in percent terms and their 95% confidence intervals. Each observation is a searcher in an experiment and all treatments that included information about prior applications were pooled. For each searcher, we observe the number of applications in the experiment sent to jobs with a number of prior applications in a given bin. This calculation is done at the time the seeker is first exposed to a particular vacancy. Standard errors are computed via the delta method. Table A.5 displays the regression estimates used to generate this figure.

The fact that we don't detect gender heterogeneity contrasts other papers that have found gender differences in job search (Cortés et al. (2021)) and selection into competitive environments (Niederle and Vesterlund (2007), Van Veldhuizen (2022)). This difference may be explained by the fact that we study a different aspect of job search, choosing where to apply as a function of the number of prior applicants, rather than choosing what offers to accept or whether to apply to jobs where compensation is structured as a tournament. In particular, both men and women may find it equally beneficial to apply before others to a given vacancy.

4.3. Survey evidence on competition avoidance

Our field experimental results show that, on average, job seekers prefer to apply to jobs with few other applicants. However, these results tell us little about heterogeneity in these preferences or causal mechanisms. To further investigate heterogeneity and causal mechanisms, we designed a survey choice experiment. This experiment corroborates our prior results that the vast majority of job seekers avoid competition due to information about the number of applicants. Herding is not quantitatively important.

The pre-registered survey experiment[15] is designed as follows. Participants recruited from the online platform Prolific were faced with three choice scenarios (see Figure A.12). In each scenario, they chose which of two jobs they value applying to the most were they to be unemployed. The jobs in each choice differed in their names (Blank Co or Brown Co), wages (with the higher wage always assigned to Brown Co), and current applications. For each choice and participant, we randomized with equal probability which of the two jobs was associated with the higher number of prior applicants. For example, choice 1 was between a job at Blank Co at an $18 wage or a job at Brown Co at a $20 dollar wage. Either 'Currently has: 200+ Applications' was displayed for Blank Co and 'Currently has: 5 - 20 Applications' was displayed for Brown Co, or vice versa. Other choices had differing levels of current applications (0 - 5 vs 5 - 20 and 0 - 5 vs 200+). There were a total of 592 participants in the sample comprising the main analysis. These participants had a median age of 35 and were 52% male. We also included a variation in which we displayed the "AI probability you get an offer". We discuss this variation in Section A.5

Figure 5 shows that participants were much more likely to choose the lower wage option when it has fewer current applicants. Almost no one choose the lower wage option when it had more current applicants. Section A.5 contains additional analysis of the survey, including analysis of textual responses, regression analysis, and heterogeneity analysis by gender and recent job search experience.[16] Consistent with our field experiment, we find no heterogeneity of the effects by gender.

5. Evidence for competition avoidance from the effects of displaying vacancy age

We have already 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 that is correlated with vacancy competition. Job-seekers often know the age of a vacancy because it is provided directly by the platform and JOF provided this information by default. To understand how this information is used, we study how job-seekers without access to this vacancy age information responded. We find that job searchers use the vacancy age to direct applications to new vacancies, and that as a byproduct this results in applications to vacancies with few prior applicants. When vacancy age information exists, competition information helps to attract applicants to older vacancies that have relatively little competition.

Our empirical strategy is to study the effects of a treatment arm in which vacancy age is not displayed but everything else is held constant. This was the case in one treatment arm of Experiment

[15] MIT's Committee on the Use of Humans as Experimental Subjects determined the experiment to be exempt from the IRB. The pre-registration for the experiment is available here: `https://www.socialscienceregistry.org/trials/9344`.

[16] We find no substantial treatment effect heterogeneity by gender either in our field or survey experiments. It seems that in the context of information about prior applicants, both men and women seek to avoid competition at equal rates. Note that since Gee (2019) does not find evidence for competition avoidance, her results on gender differences do not correspond to our setting.

Figure 5: Distribution of survey responses

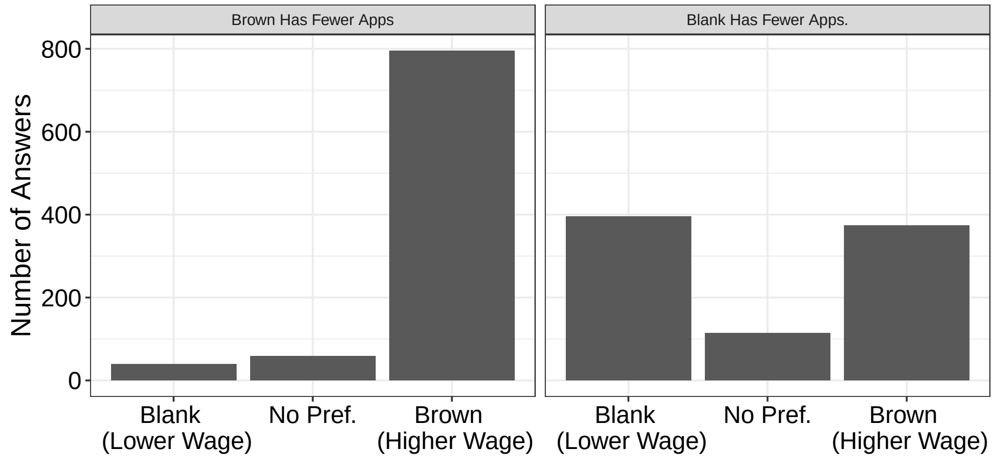

Notes: This figure plots the distribution of responses in the choice survey. The left side contains responses for which Brown Co is displayed as having fewer applications and having a higher wage. The right side contains responses for which Blank Co has fewer responses and a lower wage.

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 6 displays the overall effects of removing vacancy age. We see that treated users submitted fewer applications, were less likely to apply to any job, and clicked on fewer vacancies. These effects are of comparable magnitude to the effects of including information on prior applicants. These effects could be caused by differences in information or could be caused by a change in the 'look and feel' of the platform that may cause seekers to leave. Next, we measure heterogenous treatment effects to show that the results are inconsistent with a simple story in which a change in the platform simply causes users to leave.

We calculate the treatment effects of removing the vacancy age split by the age of the vacancy at the time of the view. The top panel of Figure 7 displays the treatment effects on actions relating to either vacancies that were less than 5 days old or to vacancies that were more than 5 days old. Removing vacancy age decreased applications and detail views to newer (≤ 5 days old) 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 direct search towards newer vacancies based on observed information.[17] In particular, in the absence of vacancy age information, some older vacancies look particularly attractive. The increase in applications to older vacancies is inconsistent with a simple story in which a change in the look of the platform causes users to uniformly submit fewer applications.

[17] Seekers may also try to infer vacancy age from other job characteristics or the ranking of the result. In this sense, our estimates represent a lower bound on the effects of vacancy age on search.

Figure 6: Treatment effects on job search outcome
Removing vacancy age

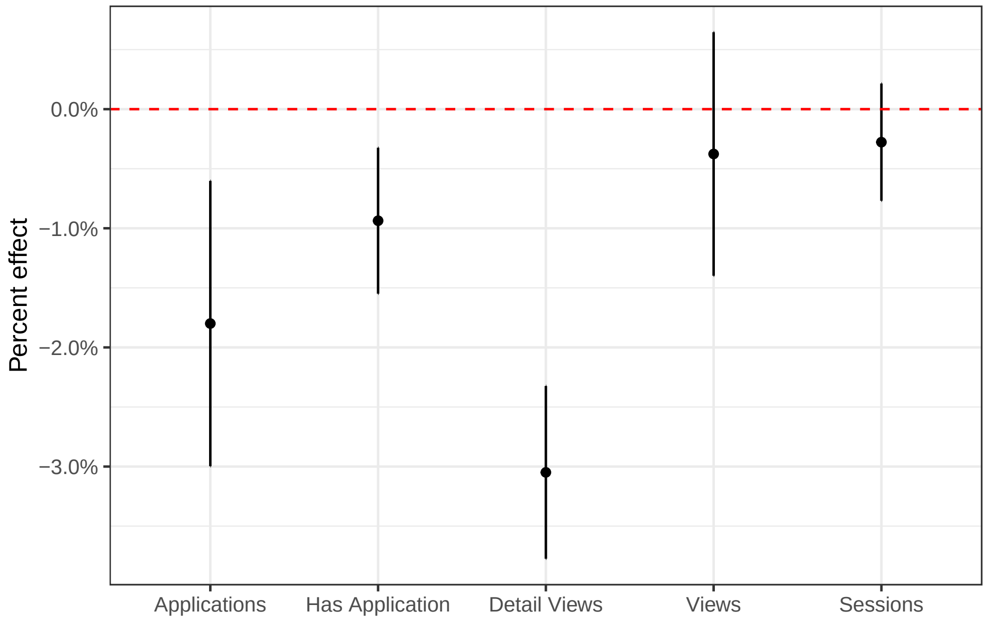

Notes: This figure plots the average effects in percent terms and their 95% confidence intervals. Each observation is a searcher in experiment I, who was either in the control group or in the treatment group for which vacancy age was removed.

Since vacancy age is correlated with the number of prior applications, it is natural to suppose that vacancy age information helps direct searchers to vacancies with less competition. We test this be calculating the heterogeneous treatment effects of removing vacancy age by different levels of prior applications. The results of this exercise are shown in the lower panel of Figure 7. The treatment causes the largest drops in applications for vacancies with fewer than 5 prior applicants and it has no effects on vacancies that receive over 200 applications. This confirms that vacancy age information allows job seekers to reduce their exposure to competition. At the same time, vacancy age can have other roles, such as signaling the selectivity of the vacancy or the likelihood it has already been filled. Our estimated treatment effects likely represent a combination of these factors.

If vacancy age already allowed searchers to find low-competition vacancies, then why did the prior applicant information treatments have an effect? One reason may be that this information allows seekers to identify older vacancies with little competition. Figure A.3 shows the treatment effects of competition information for applications to vacancies with different ages. We find that the treatment increased applications to older vacancies, confirming this conjecture.

Note that our experiment 1 also included two treatment arms that removed vacancy age and added prior applicant information (see Figure A.6). These arms undid some of the negative effect of removing vacancy age on applications to jobs with less than 5 prior applications. While the

Figure 7: Treatment effects by job type
Removing vacancy age

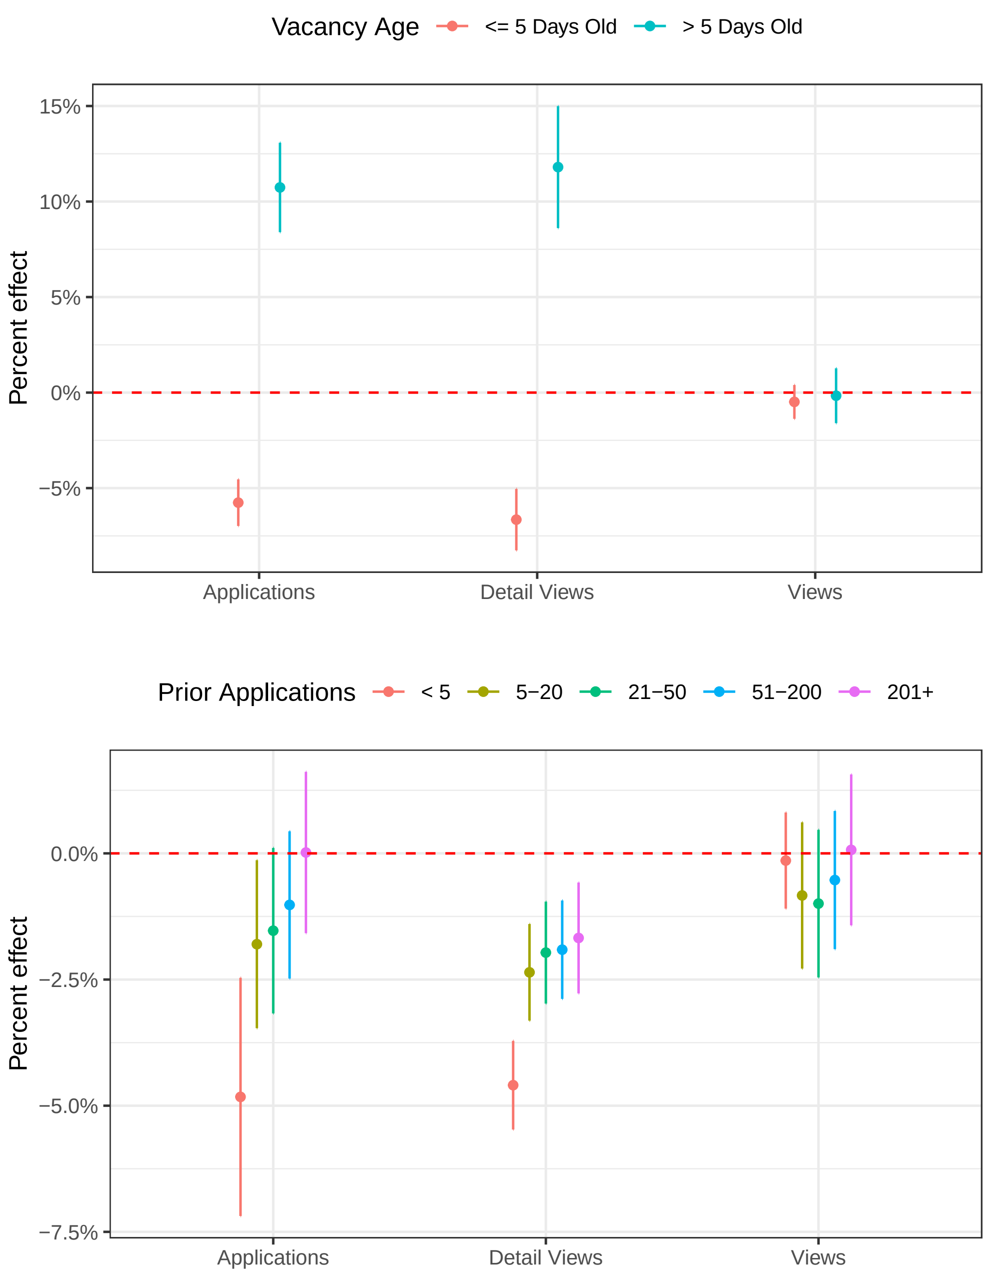

Notes: This figure plots the average effects in percent terms and their 95% confidence intervals. Each observation is a searcher in experiment I, who was either in the control group or in the treatment group for which vacancy age was removed.

treatment arm that removed vacancy age reduced applications to jobs with less than 5 prior applications by approximately 5%, the arms that also added information about prior application had reductions on the order of 2.5%. Meaning that information about vacancy age and competition are substitutable but not fully so.[18] The sign of this interaction effect supports the competition avoidance mechanism.

6. Effects on application outcomes

Signals of vacancy competition help searchers direct their applications to vacancies with less competition. This is primarily a behavioral finding but does not tell us whether displaying competition information is good for the platform or the searchers. Below, we study this question by comparing application outcomes submitted across treated and control users. Treated users did not experience worse application outcomes, and applied more often and to less competitive vacancies relative to control users. The combination of these results suggests that the treatment was successful in accomplishing the platform's goals.

Our evidence in this section comes from a series of regressions. Column 1 of Table 2 reports the results of a regression of the log of the application order on the pooled treatment indicator. We see that the order of treated applications was 1.5% lower than the order of control applications. This difference in application order could be due to selection — namely changes in the types of vacancies applied to – or changes in the speed with which applications are sent. In column 2, we report results from the same regression but with vacancy-fixed effects. Conditional on a vacancy, treated applications were submitted .2% faster. The conditional effect is much smaller than the unconditional effect — demonstrating that most of the reduction in competition is due to redirecting applications to vacancies with fewer other competitors.

We investigate this selection effect further in columns 3 and 4. Column 3 shows that treated applications are sent to vacancies that receive fewer eventual applications. Column 4 shows that these applications are also more likely to be sent to third-party applications, showing that information about competition is especially important for these vacancies.

We now investigate whether the treatment induces better outcomes. There are several ways in which the treatment could have affected outcomes. First, since treated applications arrived earlier, they should have had higher success rates, all else equal. Second, since treated applications went to different vacancies on average, these vacancies could have had different proclivities to hire or not. Lastly, since application rates changed, treated applicants may have sent their applications to better or worse matching vacancies. The sign of these effects is theoretically ambiguous.

Table 2 shows that treated applications arrive 1.5% earlier, and that a 1% increase in application order decreases the likelihood of a view by 7.8% (Column 5). Multiplying these two together will get us the expected increase in view probabilities in the treatment of 1.2%.

[18] Our theoretical model in Section A.1 describes why signals about competition and age not redundant with each other from the perspective of a job seeker.

Table 2: Application order and characteristics

	Log App Order (1)	(2)	Log Eventual Apps (3)	Third Party Vac. (4)	Viewed (5)
Treatment	-0.0147***	-0.0021**	-0.0127***	0.0017***	
	(0.0027)	(0.0009)	(0.0028)	(0.0005)	
Log App Order					-0.0777***
					(8.96×10^{-5})
R^2	0.040	0.834	0.047	0.006	0.125
Observations	13,846,246	13,846,246	13,846,246	13,846,246	12,765,262
Experiment fixed effects	✓	✓	✓	✓	✓
Vacancy fixed effects		✓			

Notes: This table contains results for a linear regression of applications outcomes on treatment (competition information), where all applications sent in the experimental sample are observations. '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. 'Viewed' refers to whether the application was viewed by the employer. Standard errors are clustered at the applicant level. ** $p < 0.05$; *** $p < 0.01$.

Table 3: Differences in application outcomes

	Viewed (1)	Contact (2)	Interview (3)
Treatment	0.0008	-0.0003	-2.25×10^{-5}
	(0.0006)	(0.0004)	(0.0001)
Mean of Y:	0.455	0.273	0.017
R^2	0.052	0.046	0.001
Observations	12,765,262	12,765,262	12,765,262
Experiment fixed effects	✓	✓	✓

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. Standard errors are clustered at the applicant level. ** $p < 0.05$; *** $p < 0.01$.

The treatment had negligible effects on the success rate of applications. Table 3 displays outcomes across treated and control applications for applications sent to vacancies for which we can measure outcomes. Column 1 demonstrates that treated applications to vacancies were no more likely to be viewed by employers. The magnitude of this effect is .18%, which is very small and smaller than our expected 1.2% effect based on the application order effect. This suggests the presence of selection to vacancies with lower view rates and differences in application quality for marginal applicants. We also consider the effects of the treatment on interviews and hires—outcomes that are more directly related to what the applicant cares about. Columns 2 and 3 of Table 3 show tiny and not statistically significant effects on these outcomes. These effects are precisely close to zero — the 95% confidence interval excludes effects on the order of more than .25% ($\frac{-.0003-1.96*.0003}{.273}$) in magnitude for contacts by employers.

To summarize, we find that treated applications face less competition. This is a function of both applying to different vacancies and applying earlier. We find precisely small effects on application outcomes. These two results combined suggest that the treatment was successful in accomplishing the platform's goals: Treated users did not experience worse outcomes, and at the same time, applied more often and to less competitive vacancies.

7. Employer behavior and the effects of applying earlier

In the previous section, we showed that job seekers who see competition information apply more, but that they have similar per application outcomes. This null effect is due to a combination of factors, including that applying earlier is helpful, that jobs with fewer prior applications tend to have differing response rates, and that marginal applications may be worse. In this section, we isolate the first of these factors, that applying earlier is helpful to applications.

Applying earlier may be helpful for one of several reasons. If employers evaluate candidates in batches and tend to check the first batch early, then it helps to be in the first batch. If employers continuously evaluate candidates as applications flow in, then earlier applications are more likely to be considered. On the other hand, if vacancies have few qualified applications or if employers check applications long after the vacancy opens, then applying earlier may not help much.

In this section, we show that employers tend to notice applications quickly and job seekers benefit from applying earlier. These results are enabled by a unique aspect of our dataset — that we can measure when employers view, message, and set up an interview for an applicant. Our measures of views and messages are more reliable since these activities very frequently happen on the platform, while our measures of interviews are sparse and noisy, since most employers continue conversations with applicants off the platform.

Table 4 displays summary statistics about employer actions and first applications for a 10% sample of vacancies between March 3 and August 18 of 2019, the period during which our experiments ran. We consider the time of the first application and the time of the first action of an employer towards any application. We find that the median vacancy receives the first application one day after being posted. An application is viewed and contacted almost immediately, typically within one day of being sent. In the 1.3% of cases where we observe an interview, we find that it typically occurs six days after the initial application was received. This data shows that typical employers are responding quickly to applications.

Next, we measure the causal effect of sending an earlier application in a two-way fixed effects regression. We take advantage of the fact that we observe the outcomes of multiple applications to the same vacancy, and the outcomes of multiple applications by the same seeker across vacancies.

Table 4: Employer responses to applications

	Mean	Median	P25	P75	Max	Missing (%)
Days Job Open to First App.	5.76	1	0	3	145	0
Days First App. to First View	14.36	0	0	2	1316	45
Days First App. to First Contact	18.64	1	0	7	1244	59
Days First App. to First Interview	20.27	6	2	17	1071	99

Notes: This table displays summary statistics for vacancies that receive at least one application. The first row measures the days between the date that a vacancy was posted and when the first application arrived. The later three rows describe the number of days between the arrival of the first application and three outcomes: an employer's first view of any application to the vacancy, an employer's first contact of any application to the vacancy, and the first observed interview time for the vacancy.

In particular, we estimate regressions of the following form:

$$y_{i(s)j} = \beta_{\text{order},i(s)j} + \delta_{\text{age},i(s)j} + \kappa_j + \mu_s + \epsilon_{i(s)j} \tag{2}$$

where $y_{i(s)j}$ are application outcomes for vacancy j by application i from searcher s, $\beta_{\text{order},i(s)j}$ are fixed effects for the order of the application (e.g., 1st application, 10th application), $\delta_{\text{age},i(s)j}$ are fixed effects for the age of the vacancy at the time (e.g., 1 day, 3 days) at which the application was submitted, κ_j are vacancy fixed effects, and μ_s are seeker fixed effects. To ensure that the sample is balanced, we consider only the first 100 applications for vacancies that received at least 100 applications.[19]

The key identification assumption in the above specification is that the applications seekers send when they happen to arrive later (e.g. the fifth application) are similar to applications they send when they arrive earlier (e.g. the first application), conditional on the age of the vacancy. This assumption is likely to hold since the seekers did not observe the prior number of applicants when they apply under the status quo.

Figure 8 displays the estimated coefficients on application order from the above equation, where the first application is normalized to 0. Later applications are less likely to be viewed, contacted, and interviewed, with effects for views and contacts materializing even as soon as the first five applications. To get a sense of the magnitudes involved, we can compare effect sizes to baseline numbers. The sample's baseline rate of interviews is approximately 1%, which reflects the fact that many employers do not use the JOF interface to record interviews. We see that later applications experience a reduction in interview rates of almost .05 percentage point. To summarize, application order matters for application outcomes, and this justifies our focus on providing application order information to job-seekers.

[19] This keeps 26% of observations.

The fact that views and contacts start dropping immediately, suggests that job seekers should care about being one of the first applicants. A theoretical justification for this can be seen as follows. Suppose that the vacancy has already received n applications and the job seeker is considering sending application n + 1. Furthermore, suppose that the vacancy is equally likely to hire each one and only hires one applicant. Then the hire rate is highest when there are few other applications and the derivative of the hire rate is highest at the second application and diminishes with the number of other applicants.

Figure 8: Relationship between application order on employer responses

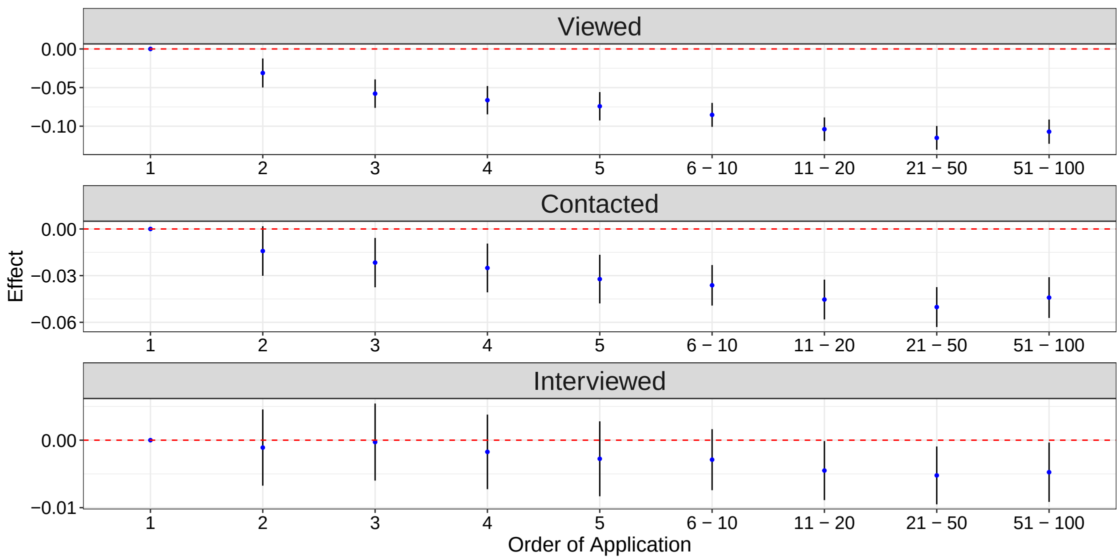

Notes: Each point represents the estimated effect and each line presents the 95% confidence interval for estimates of the effects of application order on whether an application is viewed, responded to, or results in an interview. The coefficient on the first application is normalized to 0. Standard errors are clustered at a vacancy level.

We also consider the effect of vacancy age on the likelihood of application success. We plot the estimated coefficients for vacancy age from Equation 2 in Figure 9. We find that even conditional on the order of the application, vacancy age negatively affects employer views, contacts, and interviews.

The above results show that there is at least some sequential search by employers, contrary to some prior results in labor economics (Van Ours and Ridder (1992)). This may be due to a few reasons. First, modern applicant tracking systems alert employers immediately when applications arrive. Second, many employers on JOF are small, and many have recruitment strategies that differ from large firms.[20]

A related question is whether and how job-seekers know about the negative relationship between application order, vacancy age, and interview rates. We believe that this relationship is likely to

[20] Van Ours and Ridder (1992) considered recruiting in the 1980s and for firms with at least 10 employees.

Figure 9: Relationship between vacancy age at the time of application and employer responses

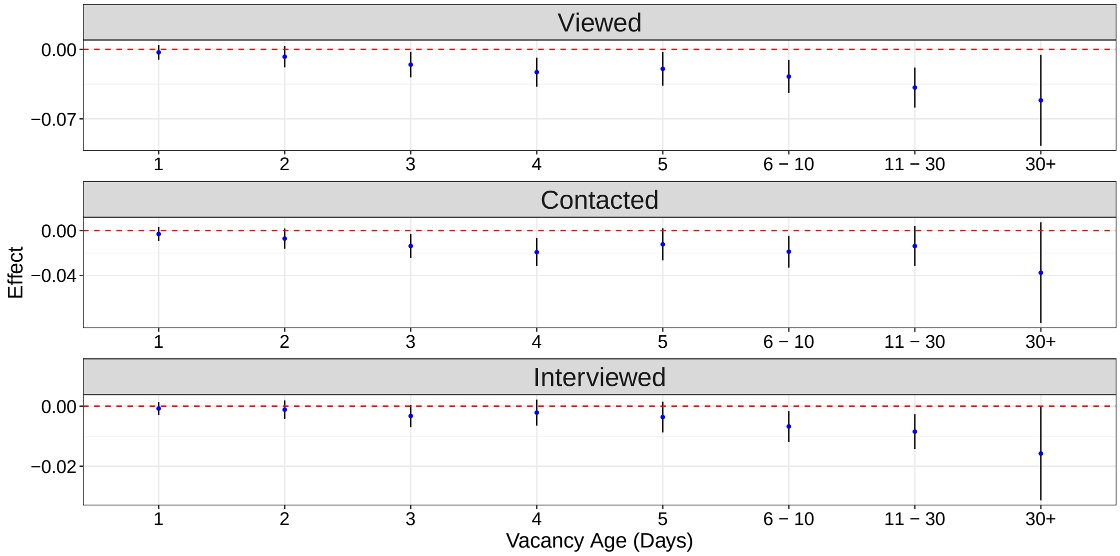

Notes: Each point represents the estimated effect and each line represents the 95% confidence interval for estimates of the effects of vacancy age at the time of application on whether an application is viewed, responded to, or results in an interview. The coefficient on the first application is normalized to 0. Standard errors clustered at a vacancy level.

hold for most vacancies, regardless of platform, although the magnitudes may vary. As a result, anyone searching for a job before may have had a chance to learn about it. Furthermore, many people have been on the hiring side of the market and could have observed that earlier applications get more attention.

8. Implications for Market Design and Discussion of Related Evidence

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. This evidence suggests that job platforms should explore interventions that surface this information. To the extent this information is simply a proxy for job offer probabilities, platforms can also give signals of predicted job offer probabilities. Below, we discuss some other aspects of job platform design and how they relate to our findings.

One aspect of the design that we have not explored is putting information about the number of applicants further in the funnel. The placement of information and the timing of its acquisition has been shown to matter for outcomes in search markets (Branco et al. (2012), Hodgson and Lewis (2020), Gardete and Hunter (2020), Abaluck and Compiani (2020)). Most relevant for our work is Gee (2019), which considers an intervention on LinkedIn where information about prior applicants is displayed after job seekers click on a job listing. Gee (2019) finds that competition information increases overall applications, but does not find evidence of competition avoidance. The lack of evidence for competition avoidance may be caused by the fact that its hard to direct search when the competition information is hidden on the search page. Another explanation is that the estimates of competition avoidance in Gee (2019) have wide confidence intervals.

Our study differs from Gee (2019) in a variety of ways in addition to just the placement of information, and this makes it hard to attribute any differences in our findings to one factor. For example, the platform interfaces and the relative share of search on mobile devices differ across the two papers. Our sample tends to be less educated and more international (42% US based in Gee (2019) vs 20% to 25% in our experiments). The types of jobs on JOF also tend to have lower skill requirements than those on LinkedIn circa 2013, in which high-tech and finance were over-represented. That said, we find competition avoidance effects even for vacancies that require a high skill level (Table A.6), meaning that differences are not entirely explained by skill requirements. We interpret the combination of evidence from our papers as showing that job seekers pay attention to competition information, but that the exact effects of this information depend on context.

Another market design decision is how frequently to display information. Displaying information for one item but not another may reduce cognitive load for searchers and may take up less screen space. At the same time, searchers may infer something about vacancies without information from vacancies with information as in Pennycook et al. (2020), which studies misinformation labels. Searchers may also infer something about the entire platform from information about prior applicants, as in Fong (2019) and Tucker and Zhang (2010).

We can study whether information about an item affects choices about another item by considering the treatment arms in our experiment where information is not always shown. We focus on the treatments that display 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 learning is important, 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 five applications. We see a positive and statistically significant effect for vacancies in a position divisible by 3. In contrast, for vacancies in other positions, we see a negative but small magnitude and insignificant effect. This coefficient is consistent with some level of negative spillovers between ads with and without information. If we take the point estimates at face value, then the negative spillovers on the two positions without information (2*.5%) are smaller than the positive effect on the vacancy with the information (2.5%), although this difference is not statistically significant.

In the same figure, we also display results just on applications to vacancies with few prior applications, which drive our main effects. When 'be one of the first to apply' is shown due to the

Figure 10: Effects of competition information on applications, by the position of vacancy and prior applications

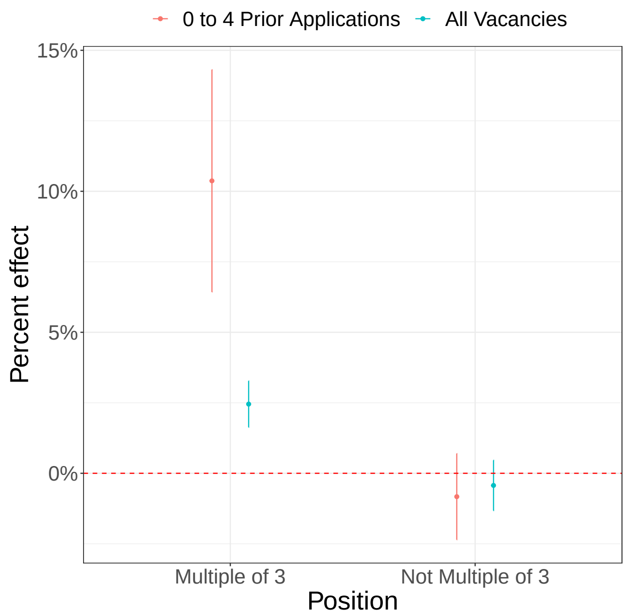

Notes: This figure plots the average effects in percent terms and their 95% confidence intervals. The outcome is the number of applications. Each observation is a searcher either in a treatment arm where information is displayed every 3 tiles or in the control. Red lines denote effects for vacancies which had 0 to 4 prior applications at the time of view and blue lines denote effects for all vacancies. In the first column, the difference considered is for a third tile with information minus a third tile without information. In the second column, the difference considered is for a non-third tile with information minus a non-third tile without information. Standard errors are computed via the delta method.

vacancy being on a 3rd tile, the increase in application rates is approximately 10%. This increase is not offset by declines in applications to similar vacancies for which information is not shown. This provides some evidence against job seekers making an inference about the competition for vacancies without competition information.

Information design decisions may also interact in important ways with other platform design decisions, such as ranking and user acquisition strategies. For example, ranking can be used to boost jobs with fewer prior applicants in a way that may be complementary to information designs. During the time of our study, JoF was experimenting with a variety of ranking algorithms. So our results should be interpreted as being conditional on the mix of algorithms in use at the time.

A critical market design issue is the presence of spillovers across seekers, vacancies, and employers. For example, displaying information about competition may result in fewer low competition vacancies in equilibrium, changing the benefits to job seekers of frequently monitoring the platform for new vacancies to apply to. Information policies may also change the benefits to posting vacancies, for example by allowing vacancies to accumulate the first few applications more quickly. The quality of these applications is also an equilibrium object.

When considering application quality, the platform has very noisy information about whether a hire is made and whether it was a good hire. Natural proxies for welfare such as duration on the job or wages are unavailable in most cases. As a result, the welfare effects of these information policies cannot easily be studied. That said, we believe there is much more to be done by using long-running market level experiments and structural models.[21]

9. Conclusion

Social signals are used across platforms to influence search behavior and other activity. We have investigated the role of such signals in the context of job search, where the sign of the effect of social influence is ex-ante uncertain. On the one hand, telling job seekers that few people have applied could be a signal that the job is low quality. On the other hand, this information can also convey that the job seeker has a higher chance of getting the job. We find that the latter effect dominates. Job seekers prefer to apply to jobs with very few prior applicants.

Even in the absence of competition information, applicants were able to direct their search toward less competitive vacancies. In a complementary treatment, we find that the job seekers strongly preferred new vacancies when information on vacancy age was available. Information about vacancy age greatly increased applications on the platform and redirected those applications away from popular but old vacancies. Both vacancy age and competition information increase usage of the platform and did not harm application outcomes. These results point to the positive effects of displaying this information. Contrary to work on other types of labor market interventions, we fail to find gender differences in how job seekers respond to this competition information.

Our work cannot measure how the design of one job platform affects the overall labor market. Both searchers and employers multi-home across a variety of platforms and no single platform has a bird's eye view of the entire labor market. 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. Policymakers and platform designers should consider the implications of this fragmentation when making decisions.

References

J. Abaluck and G. Compiani. A method to estimate discrete choice models that is robust to consumer search. Technical report, National Bureau of Economic Research, 2020.

[21] Karrer et al. (2021) describe an infrastructure for running clustered experiments at Meta which was broadly used, including for ranking experiments in the JoF product. Besbes et al. (2023) builds a structural model of competition signals in a services marketplace.

J. Albrecht, B. Decreuse, and S. Vroman. Directed search with phantom vacancies. 2017.

J. Azar, S. Berry, and I. E. Marinescu. Estimating labor market power. *Available at SSRN 3456277*, 2019.

E. M. Azevedo, A. Deng, J. L. Montiel Olea, J. Rao, and E. G. Weyl. A/b testing with fat tails. *Journal of Political Economy*, 128(12):4614–000, 2020.

S. R. Baker and A. Fradkin. The impact of unemployment insurance on job search: Evidence from google search data. *Review of Economics and Statistics*, 99(5):756–768, 2017.

E. Bakshy, D. Eckles, and M. S. Bernstein. Designing and deploying online field experiments. In *Proceedings of the 23rd international conference on World wide web*, pages 283–292, 2014.

S. Banfi and B. Villena-Roldan. Do high-wage jobs attract more applicants? directed search evidence from the online labor market. *Journal of Labor Economics*, 37(3):715–746, 2019.

M. Belot, P. Kircher, and P. Muller. How wage announcements affect job search-a field experiment. 2018.

O. Besbes, Y. Fonseca, I. Lobel, and F. Zheng. Signaling competition in two-sided markets. *Available at SSRN*, 2023.

F. Branco, M. Sun, and J. M. Villas-Boas. Optimal search for product information. *Management Science*, 58(11):2037–2056, 2012.

A. Cheron and B. Decreuse. Matching with phantoms. *The Review of Economic Studies*, 84(3):1041–1070, 2017.

L. C. Coffman, C. R. Featherstone, and J. B. Kessler. Can social information affect what job you choose and keep? *American Economic Journal: Applied Economics*, 9(1):96–117, 2017.

P. Cortés, J. Pan, L. Pilossoph, and B. Zafar. Gender differences in job search and the earnings gap: Evidence from business majors. Technical report, National Bureau of Economic Research, 2021.

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

J. A. Flory, A. Leibbrandt, and J. A. List. Do competitive workplaces deter female workers? a large-scale natural field experiment on job entry decisions. *The Review of Economic Studies*, 82(1):122–155, 2015.

J. Fong. Search, selectivity, and market thickness in two-sided markets. *Selectivity, and Market Thickness in Two-Sided Markets (September 23, 2019)*, 2019.

A. Fradkin. 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.

P. Gardete and M. Hunter. Guiding consumers through lemons and peaches: An analysis of the effects of search design activities. 2020.

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

L. K. Gee, J. Jones, and M. Burke. Social networks and labor markets: How strong ties relate to job finding on facebook's social network. *Journal of Labor Economics*, 35(2):485–518, 2017. doi: 10.1086/686225. URL `https://doi.org/10.1086/686225`.

L. Hensvik, T. Le Barbanchon, and R. Rathelot. Do algorithmic job recommendations improve search and matching? evidence from a large-scale randomised field experiment in sweden. 2020.

C. Hodgson and G. Lewis. You can lead a horse to water: Spatial learning and path dependence in consumer search. 2020.

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

M. Ibañez and G. Riener. Sorting through affirmative action: three field experiments in colombia. *Journal of Labor Economics*, 36(2):437–478, 2018.

B. Karrer, L. Shi, M. Bhole, M. Goldman, T. Palmer, C. Gelman, M. Konutgan, and F. Sun. Network experimentation at scale. In *Proceedings of the 27th acm sigkdd conference on knowledge discovery & data mining*, pages 3106–3116, 2021.

R. Kohavi, D. Tang, and Y. Xu. *Trustworthy online controlled experiments: A practical guide to a/b testing*. Cambridge University Press, 2020.

K. Kroft and D. G. Pope. Does online search crowd out traditional search and improve matching efficiency? evidence from craigslist. *Journal of Labor Economics*, 32(2):259–303, 2014.

P. Kuhn and H. Mansour. Is internet job search still ineffective? *The Economic Journal*, 124(581):1213–1233, 2014.

P. Kuhn, K. Shen, and S. Zhang. Gender-targeted job ads in the recruitment process: Facts from a chinese job board. *Journal of Development Economics*, 147:102531, 2020.

T. Le Barbanchon, R. Rathelot, and A. Roulet. Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics*, 136(1):381–426, 2021.

A. Leibbrandt and J. A. List. Do equal employment opportunity statements backfire? evidence from a natural field experiment on job-entry decisions. Technical report, National Bureau of Economic Research, 2018.

I. Marinescu. The general equilibrium impacts of unemployment insurance: Evidence from a large online job board. *Journal of Public Economics*, 150:14–29, 2017.

I. Marinescu and D. Skandalis. Unemployment insurance and job search behavior. *The Quarterly Journal of Economics*, 136(2):887–931, 2021.

L. Muchnik, S. Aral, and S. J. Taylor. Social influence bias: A randomized experiment. *Science*, 341(6146): 647–651, 2013.

M. Niederle and L. Vesterlund. Do women shy away from competition? do men compete too much? *The quarterly journal of economics*, 122(3):1067–1101, 2007.

G. Pennycook, A. Bear, E. T. Collins, and D. G. Rand. The implied truth effect: Attaching warnings to a subset of fake news headlines increases perceived accuracy of headlines without warnings. *Management science*, 66(11):4944–4957, 2020.

M. J. Salganik, P. S. Dodds, and D. J. Watts. Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762):854–856, 2006.

A. Samek. Gender differences in job entry decisions: A university-wide field experiment. *Management Science*, 65(7):3272–3281, 2019.

D. Skandalis. Breaking news: Information about firms' hiring needs affects the direction of job search. Technical report, mimeo, 2018.

C. Tucker and J. Zhang. Growing two-sided networks by advertising the user base: A field experiment. *Marketing Science*, 29(5):805–814, 2010.

A. Van de Rijt, S. M. Kang, M. Restivo, and A. Patil. Field experiments of success-breeds-success dynamics. *Proceedings of the National Academy of Sciences*, 111(19):6934–6939, 2014.

J. Van Ours and G. Ridder. Vacancies and the recruitment of new employees. *Journal of Labor Economics*, 10(2):138–155, 1992.

R. Van Veldhuizen. Gender differences in tournament choices: Risk preferences, overconfidence, or competitiveness? *Journal of the European Economic Association*, 20(4):1595–1618, 2022.

G. Ward. Workplace happiness and employee recruitment. In *Academy of Management Proceedings*, volume 2023, page 11933. Academy of Management Briarcliff Manor, NY 10510, 2023.

R. Wright, P. Kircher, B. Julien, and V. Guerrieri. Directed search and competitive search: A guided tour. *Journal of Economic Literature*, 2019.

J. Zhang and P. Liu. Rational herding in microloan markets. *Management science*, 58(5):892–912, 2012.

Appendix A: Appendix

A.1. Theoretical model

This section presents a stylized model of job search and hiring that incorporates competition aversion, herding, and vacancy age. Job seekers would like to obtain the job that gives them the highest utility. However, some jobs that would give a high utility may be hard to obtain. This could be because the employer judges the worker unqualified for the job, because the competition for the job is too high, or because the employer has already interviewed the candidates who will be hired. Given these concerns, job seekers value information about the likelihood of obtaining a job in addition to the quality of the job. Below, we describe a simple model of job search that accounts for these factors.

We set up the model with many simplifications in order to highlight the core tension between competition avoidance and herding. Consider application decisions for a vacancy with mass 1. Suppose there are two periods and that a random mass ($n_t \sim G$) of job seekers consider the vacancy in each period. We assume that the number of seekers in each period is independent of each other. There are two types of vacancies, high types yield a utility of 1 conditional on getting the job, while low types yield a utility of 0, and each of these occur with equal probability.

In period one, all job seekers receive a binary signal of quality, $s \geq .5$, such that the high type occurs with probability s if the signal is received and $1-s$ otherwise. The job seekers know they all receive the same signal and they all know they've arrived in period 1, this is consistent with the status quo platform design where the vacancy age is displayed. There is a cost of applying $c_i \sim F$. Job seekers arriving in period 2 know they've arrived in period 2, and do not receive a signal about vacancy quality, although they know that those in period 1 did.

The employer has the opportunity to make hiring decisions either in the first or in the second period. Variation in when the employer checks may be due to notifications about an application, or a set recruiting schedule. Either of which may cause an employer to evaluate prior to all applications arriving. With an exogenous probability, λ, the employer hires only in period 1, and otherwise the employer hires only in period 2. We posit that the hiring probability for an applicant is $\frac{1}{m(a)}$, where a is the mass of applicants at the time of hiring and $m' > 0$.

We first consider an equilibrium of this model in which searchers do not receive information about the number of applications but know in which period they arrive. This is consistent with the equilibrium in the status quo of JOF, where seekers see the age of a vacancy but not the prior number of applicants. An equilibrium consists of application rates ($r_t = \frac{a_t}{n_t}$), which are independent of the number of arriving seekers.

There are three application rates in equilibrium, r_{1h}, r_{1l}, r_2, characterized by the equations 1, 2, and 3 below. In these equations, expectations are taken over the distributions of the arrival of seekers in period 1 and period 2.

$$r_{1h} = F(sE[\frac{\lambda}{m(n_1 r_{1h})} + \frac{1-\lambda}{m(n_1 r_{1h} + n_2 r_2)}]) \tag{3}$$

$$r_{1l} = F((1-s)E[\frac{\lambda}{m(n_1 r_{1l})} + \frac{1-\lambda}{m(n_1 r_{1l} + n_2 r_2)}]) \tag{4}$$

$$r_2 = F(.5(1-\lambda)E[\frac{1-s}{m(n_1 r_{1l} + n_2 r_2)} + \frac{s}{m(n_1 r_{1h} + n_2 r_2)}]) \tag{5}$$

Note that when there is no informative signal, the application rate in period 1 is greater than the application rate in period 2. This comes simply from the fact that period 1 seekers may be hired at the end of period 1 with probability λ while period 2 seekers cannot. At the same time, the likelihood that an application is matched with the vacancy is greater when the application arrives in period 1. This is due to the fact that the employer can check the application prior to all applications arriving. Both of these model predictions are confirmed in our empirical analysis.

Now suppose that at the equilibrium described above, an infinitesimal individual seeker in period 2 is given information about the number of prior applicants. This corresponds to the case where a small share of the market is treated with additional information. In Figure A.1 we plot the application rate as a function of period 1 applications for such a seeker, using a parametrized version of the above model.

The red lines represent the model outcomes when those in period 1 receive an uninformative signal about vacancy quality. The flat dotted red line represents the application rates when seekers receive no information about period 1 applications. The solid red line plots the function for a seeker who does receive such information. We see that, consistent with competition avoidance, application rates are higher relative to the blue line when period 1 applications are low, and vice versa when period 1 applications are high.

The blue lines represent the equilibrium in which period 1 seekers receive an informative signal of vacancy quality (we assume $s = .85$ for illustrative purposes).[22] Two observations are in order. First, when prior applications are low, the application rate of informed period 2 seekers is higher when period 1 seekers do not receive a signal of quality. This reflects the fact that when there is an informative signal, a low number of applications in period 1 signals a lower quality vacancy. Second, when prior application signal quality, there exists a discontinuity and non-monotonicity where the application rate jumps up as period 1 applications increase. This sudden jump occurs because there exists an application amount that can only occur when period 1 seekers received a positive signal. Period 2 seekers understand this and increase application rates at the discontinuity. Note that the presence of a discontinuity is not necessary. We picked a signal structure that yielded this sharp structure for illustrative purposes. A different signal structure could imply a smooth and non-monotonic curve if herding were strong enough.

The above model maps onto our experiment, in which a small share of users were exposed to information about prior applications. If this information just conveys information about the level of competition, then we should expect to see a curve like the red line, where treatment effects are highest when prior applications low. But if the level of competition strongly conveys information about vacancy quality, then the curve should flatten or even reverse when prior applications are high. Although our model is meant to be illustrative, it is consistent with the empirical results we find in the experiment. Our model also has implications for the effects of information on the age of the vacancy. Job seekers should prefer newer vacancies because there is a higher chance that the employer has not yet made interview and hiring decisions.

We abstract away from the possibility that the information provided about a particular job conveys broader information about the platform as in Tucker and Zhang (2010) and Fong (2019). In Section 8 we offer suggestive evidence to support this assumption. In particular, applications increase just for vacancies for which information is shown rather than for vacancies for which information is not shown.

[22] Note that the red and blue dotted lines are not exactly identical. This is due to the fact that there is some curvature in the application cost and matching functions.

Figure A.1: Seeker Responses to Information - Model

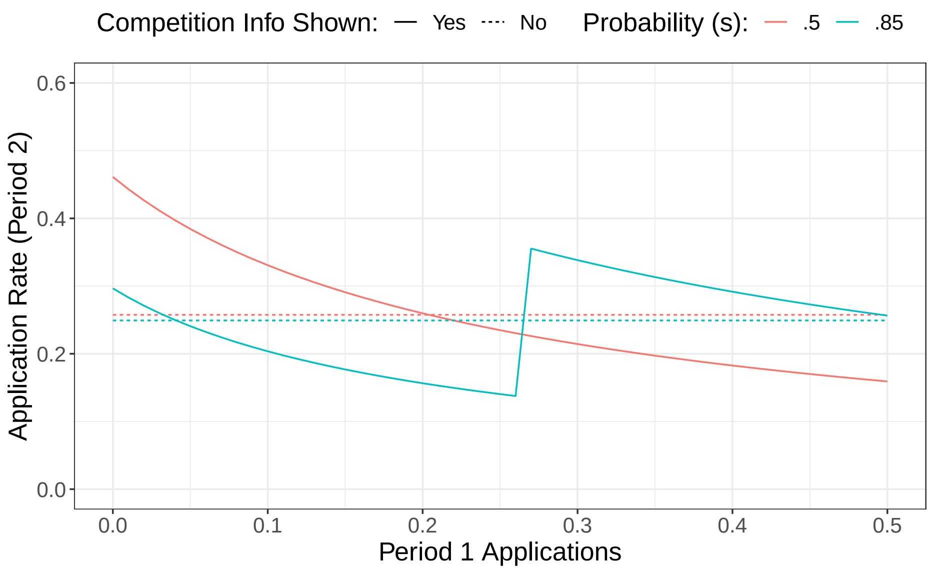

Notes: This figure plots the equilibrium period 2 application rate, as well as the response of an infinitesimal seeker who is given information about period 1 applications. We assume that $F \sim Exp(.5)$, $m = x + .1$, $G \sim U[0,1]$, and $\lambda = .5$.

A.2. Additional Figures and Tables

Table A.1: Control group summary statistics over the three experiments

	25th	Median	75th	Mean	StDEv
Exp I (n = 1,763,735)					
Age	26	33	44	35.95	13.63
US User	NA	NA	NA	0.25	NA
Friends	219	444	873	746.74	896.28
iOS User	NA	NA	NA	0.34	NA
Male	NA	NA	NA	0.59	NA
Applications	0	0	0	0.37	2.08
Detail Views	0	0	2	2.61	9.38
Views	5	17	64	80.33	474.58
Exp II (n = 863,214)					
Age	25	32	42	34.45	13.18
US User	NA	NA	NA	0.21	NA
Friends	203	437	922	785.59	969.06
iOS User	NA	NA	NA	0.26	NA
Male	NA	NA	NA	0.54	NA
Applications	0	0	0	0.45	2.05
Detail Views	0	0	3	3.38	9.74
Views	8	33	105	107.01	261.84
Exp III (n = 3,265,160)					
Age	24	31	42	34.23	13.74
US User	NA	NA	NA	0.20	NA
Friends	173	399	873	747.61	962.37
iOS User	NA	NA	NA	0.29	NA
Male	NA	NA	NA	0.51	NA
Applications	0	0	0	0.26	1.67
Detail Views	0	0	2	2.53	9.45
Views	6	22	85	94.19	266.75

Notes: User characteristics by experiment.

Table A.2: Control group summary statistics over the three experiments
Conditional on at least one application

	25th	Median	75th	Mean	StDEv
Exp I (n = 190,325)					
Age	23	29	38	31.92	12.03
US User	NA	NA	NA	0.16	NA
Friends	253	555	1207	978.40	1,115.92
iOS User	NA	NA	NA	0.18	NA
Male	NA	NA	NA	0.54	NA
Applications	1	2	4	3.39	5.47
Detail Views	3	6	14	11.94	22.32
Views	46	123	304	297.39	1,329.79
Exp II (n = 124,157)					
Age	23	29	38	31.78	11.95
US User	NA	NA	NA	0.19	NA
Friends	246	540	1172	955.94	1,101.93
iOS User	NA	NA	NA	0.21	NA
Male	NA	NA	NA	0.52	NA
Applications	1	2	3	3.11	4.59
Detail Views	3	6	14	12.20	18.85
Views	66	159	356	306.63	492.49
Exp III (n = 277,881)					
Age	23	28	37	31.39	11.91
US User	NA	NA	NA	0.23	NA
Friends	239	534	1165	944.14	1,089.89
iOS User	NA	NA	NA	0.23	NA
Male	NA	NA	NA	0.50	NA
Applications	1	2	3	3.08	4.92
Detail Views	3	7	16	13.53	23.66
Views	86	200	439	382.29	640.78

Notes: User characteristics by experiment.

Table A.3: Applications to vacancies by bin in the control group

	Exp. 1		Exp. 2		Exp. 3	
	Mean	SD	Mean	SD	Mean	SD
0 - 4 App.	0.68	2.25	0.65	1.78	0.53	1.53
5 - 20 App.	0.78	1.82	0.84	1.71	0.88	1.87
21 - 50 App.	0.51	1.18	0.68	1.30	0.69	1.36
51 - 200 App.	0.70	1.43	0.89	1.63	0.90	1.76
200+ App.	0.76	1.74	0.13	0.60	0.03	0.19

Notes: User characteristics by experiment.

Figure A.2: Effects (in levels) of competition information on applications to different status vacancies

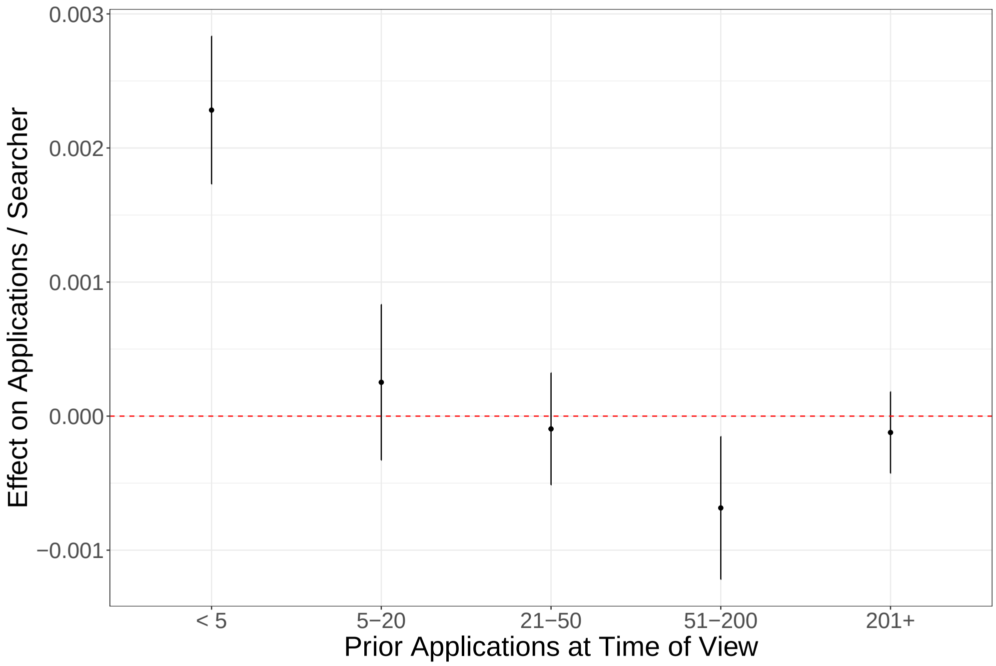

Notes: This figure plots the average effect of the treatment (pooled across all arms) in levels. Each observation is a searcher in an experiment and all treatments that included information about prior applications were pooled.

Figure A.3: Treatment effects of competition information
Split by age of vacancy (days)

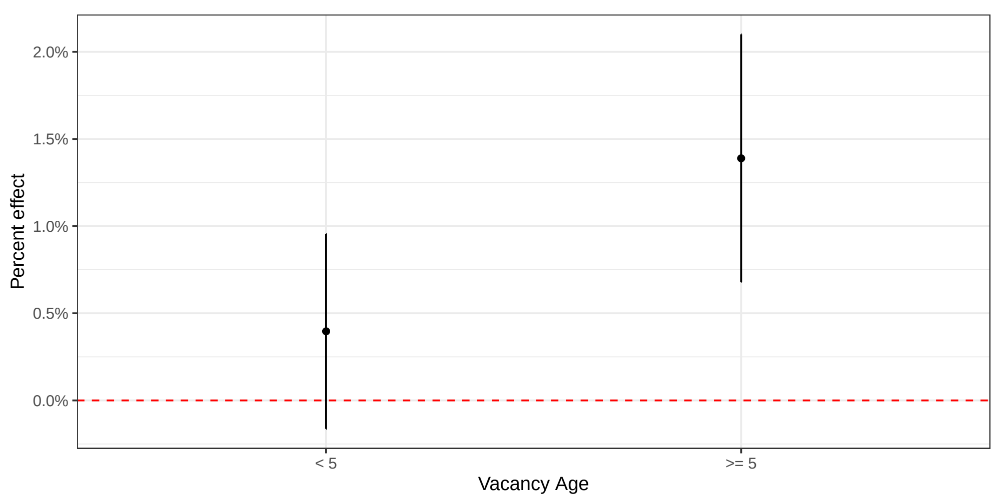

Notes: This figure plots the average effects in percent terms and their 95% confidence intervals. Each observation is a searcher in the experiments who was either in the control group or in the treatment group for which competition information was shown.

A.3. Analysis by Experimental 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 used in the study. Figure A.4 shows how treatment parameters varied across arms and experiments.

The first 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 first 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 first to apply' was displayed.[23] 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 first to apply', '5 - 20 applications', '21 - 50 applications', '50 - 200 applications', '201+ applications'. Treatment arms also 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.

To check that the randomization was properly conducted, we performed a set of balance tests. Figure A.5 displays these tests, where the p-value for the difference in means between each treatment arm and the corresponding control group is displayed for a set of pre-treatment covariates. 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 in each experiment.

In addition to treatments with social information, Experiment 1 also contained arms that varied whether vacancy age was displayed. One of these arms was discussed in Section 5. Two other arms removed vacancy age, but added competition signals (either just 'Be the first to apply' or all competition signals).

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

Figure A.4: Treatment arms relating to only competition information

	Every Tile	Every 3 Tiles	Every 10 Tiles	Other
Only "Be one of the first to apply"	Exp. / Blue 1 / None	Exp. / Blue 2 / All 3 / All	Exp. / Blue 2 / All	*Experiment 3 also had a treatment arm that displayed `Be one of the first to apply' on only the 3rd tile when eligible but displayed other congestion information in grey on every tile when eligible.
All Congestion Signals*	Exp. / Blue 1 / None 3 / First	Exp. / Blue 2 / All 2 / None 2 / First 3 / First	Exp. / Blue 2 / All 2 / None 2 / First	

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 competition 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. Note, three additional arms also varied vacancy age.

Figure A.5: Covariate balance test p-values across experiments

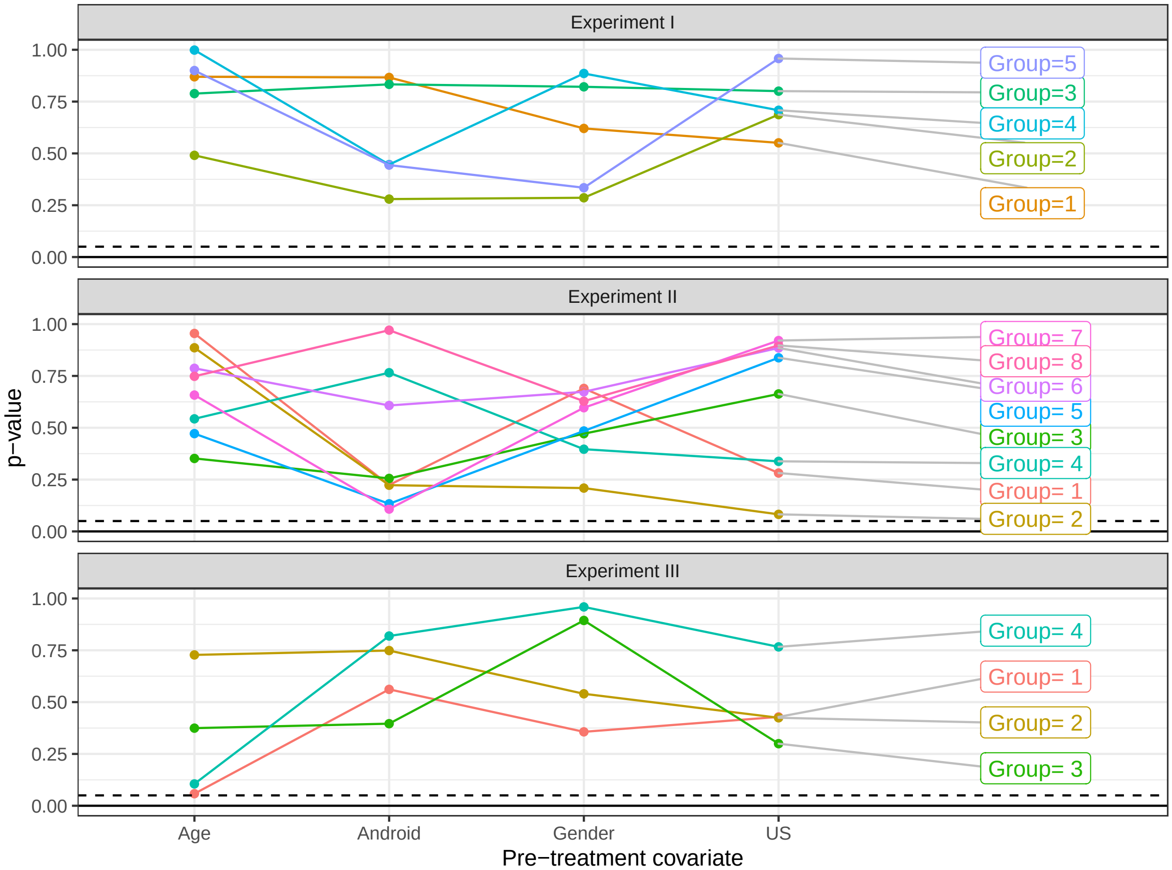

Notes: This figure displays the p-value from a linear regression where the characteristics of each treatment arm were compared with the control arm of the corresponding experiment.

Figure A.6: Treatment Effects for Experiment 1

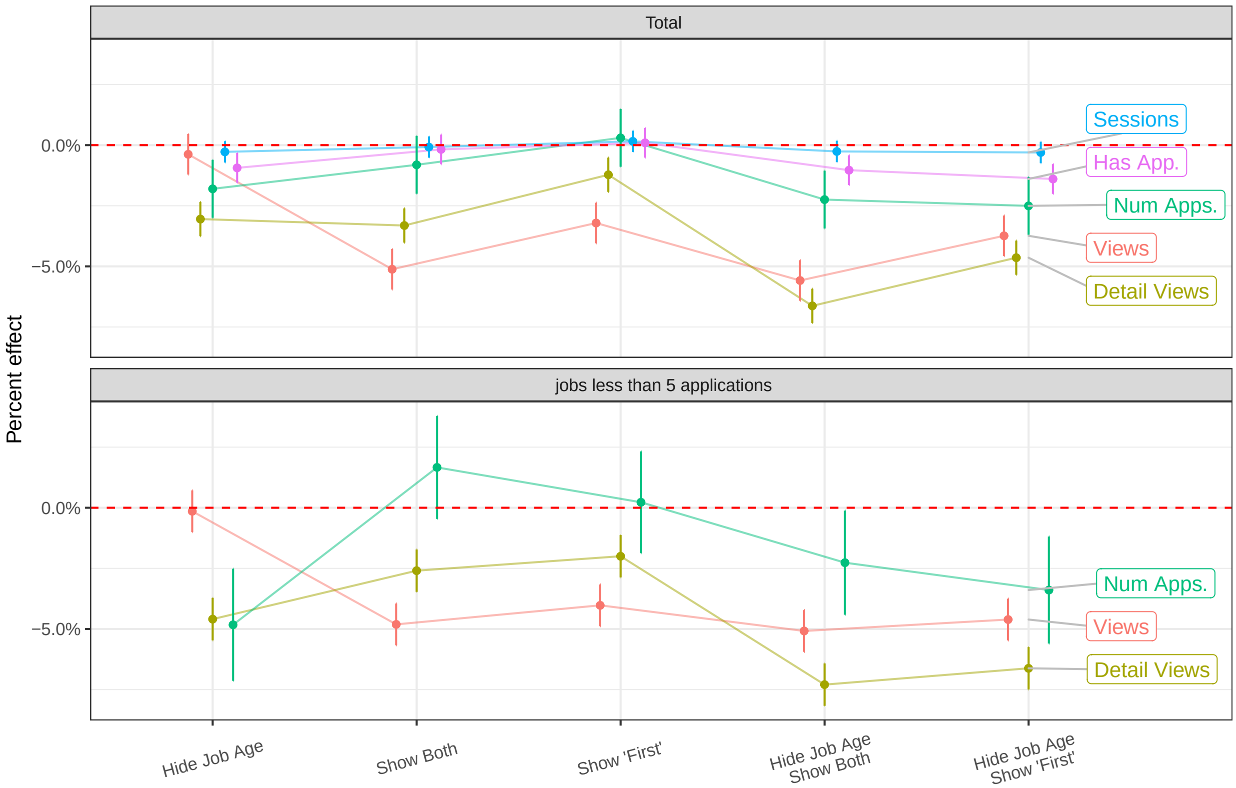

A.3.1. Effects by Treatment Arm Next, we discuss the by-arm treatment effects for each treatment and experiment. We begin with Experiment 1 (Figure A.6). Columns 1, 4, and 5 of of the figure plot the treatment effects where the vacancy age is hidden. Columns 2 - 5 plot treatments where competition information 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 A.7 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 drop 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 A.8 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 A.7: Treatment effects for experiment 2

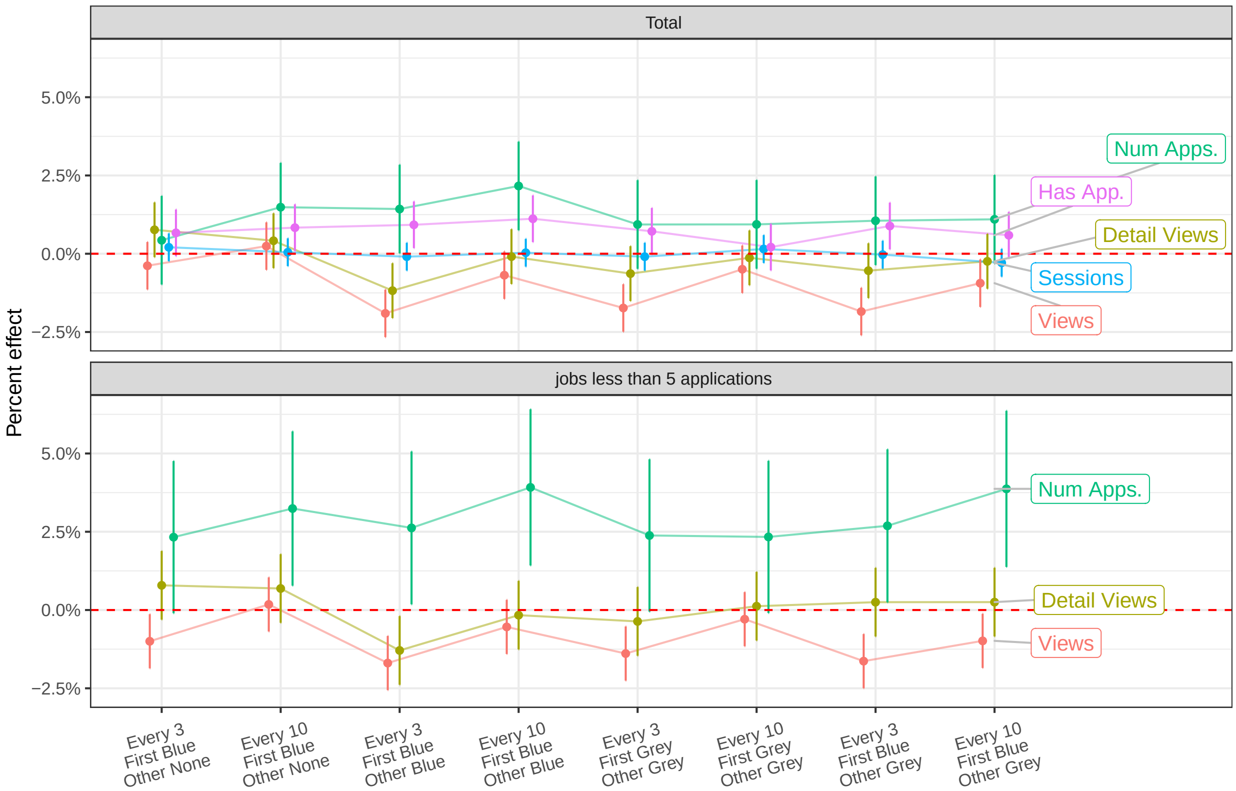

Table A.4: Treatment effects pooled across all three experiments

	Num. App. (1)	Has App. (2)	Detail Views (3)	Views (4)	Sessions (5)
Treatment	0.0019**	0.0003**	-0.0592***	-3.611***	-0.0079**
	(0.0009)	(0.0001)	(0.0044)	(0.1463)	(0.0034)
Mean of Y:	0.332	0.105	2.727	92.452	4.161
R^2	0.002	0.007	0.002	0.001	0.008
Observations	29,375,231	29,375,231	29,375,231	29,375,231	29,375,231
Experiment fixed effects	✓	✓	✓	✓	✓

Notes: This table plots the effects of the competition signal treatment pooled across experiment. ‘Num. App’ refers to the number of applications, ‘Has App.’ refers to whether a searcher has any application at all, ‘Detail Views’ are clicks onto a vacancy, ‘Views’ are views in the search list, and ‘Sessions’ are distinct visits to JOF. ** $p < 0.05$; *** $p < 0.01$.

Figure A.8: Treatment effects for experiment 3

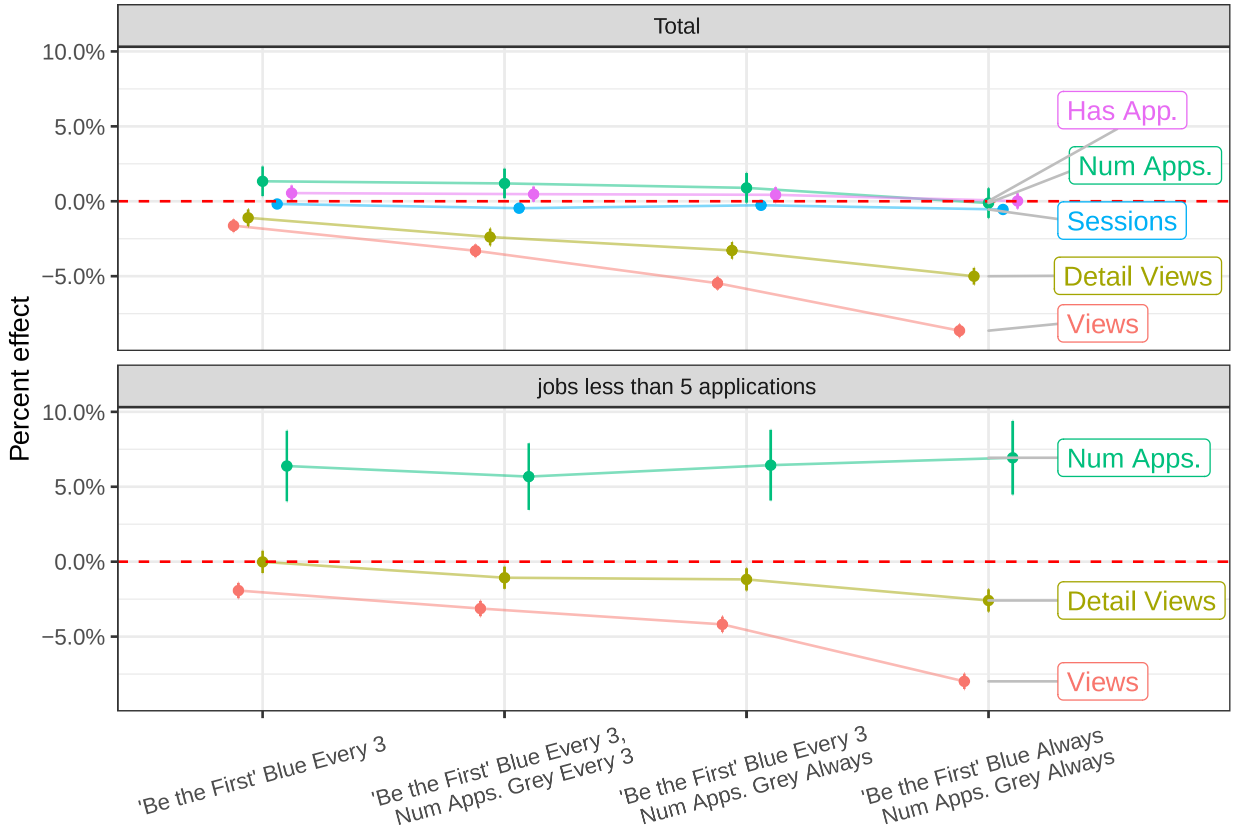

Table A.5: Treatment effects pooled across all three experiments - by application type

	0 - 4 App. (1)	5 - 20 App. (2)	21 - 50 App. (3)	51 - 200 App. (4)	201+ App. (5)
Treatment	0.0023***	0.0003	-9.53×10^{-5}	-0.0007**	-0.0001
	(0.0003)	(0.0003)	(0.0002)	(0.0003)	(0.0002)
Mean of Y:	0.065	0.089	0.069	0.089	0.021
R^2	0.001	0.001	0.002	0.002	0.010
Observations	29,375,231	29,375,231	29,375,231	29,375,231	29,375,231
Experiment fixed effects	✓	✓	✓	✓	✓

Notes: Effects of the competition signal treatment pooled across the experiments. Each column refers to application to vacancies with a given number of prior applications. ** $p < 0.05$; *** $p < 0.01$.

A.4. 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.

A.4.1. Differences in treatment As explained in Section A.3, 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 A.9. 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 experiments 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.

A.4.2. 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 A.1 reports summary statistics for user characteristics for the three experiments. 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 the number of vacancies. Figure A.10 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 A.11. 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.

Figure A.9: Effects of the same treatment across experiments

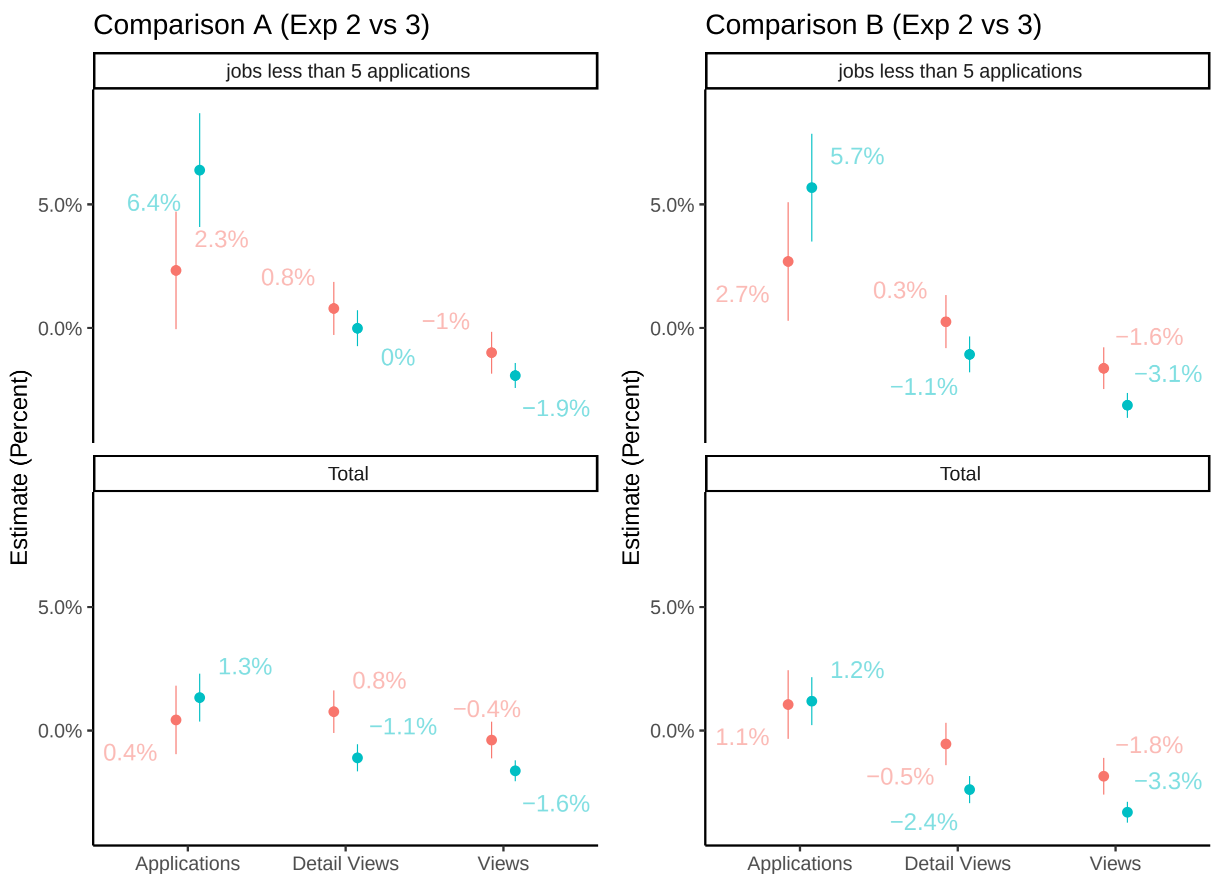

We also investigate whether there is heterogeneity based on the skill requirements of the vacancy. To do this, we use a skill requirement classification of vacancies into low, medium, and high skill that is available in the data. We consider the first vacancy exposed to each user and condition on the subset of those which had fewer than 5 applications at the time of view. We then estimate a linear regression separately for when the vacancy was one of each of the three levels. Table A.6 displays the results. We detect positive and similarly sized treatment effects for each vacancy type.

Figure A.10: Evolution of market tightness over time

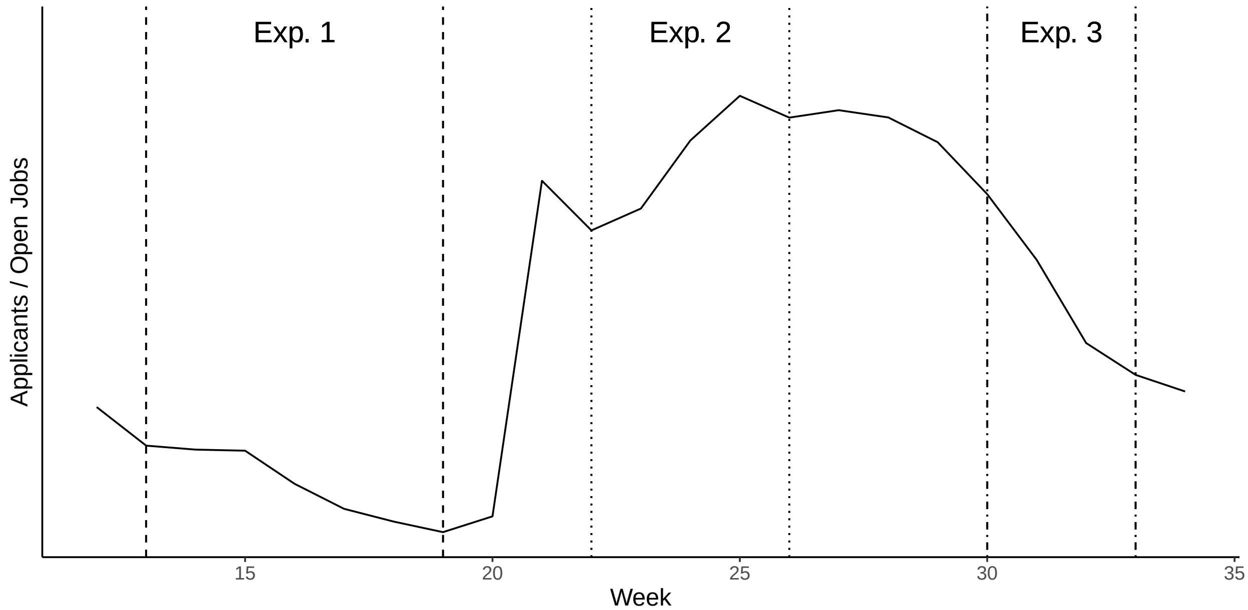

Figure A.11: Heterogeneous treatment effects - Applications to under-subscribed vacancies

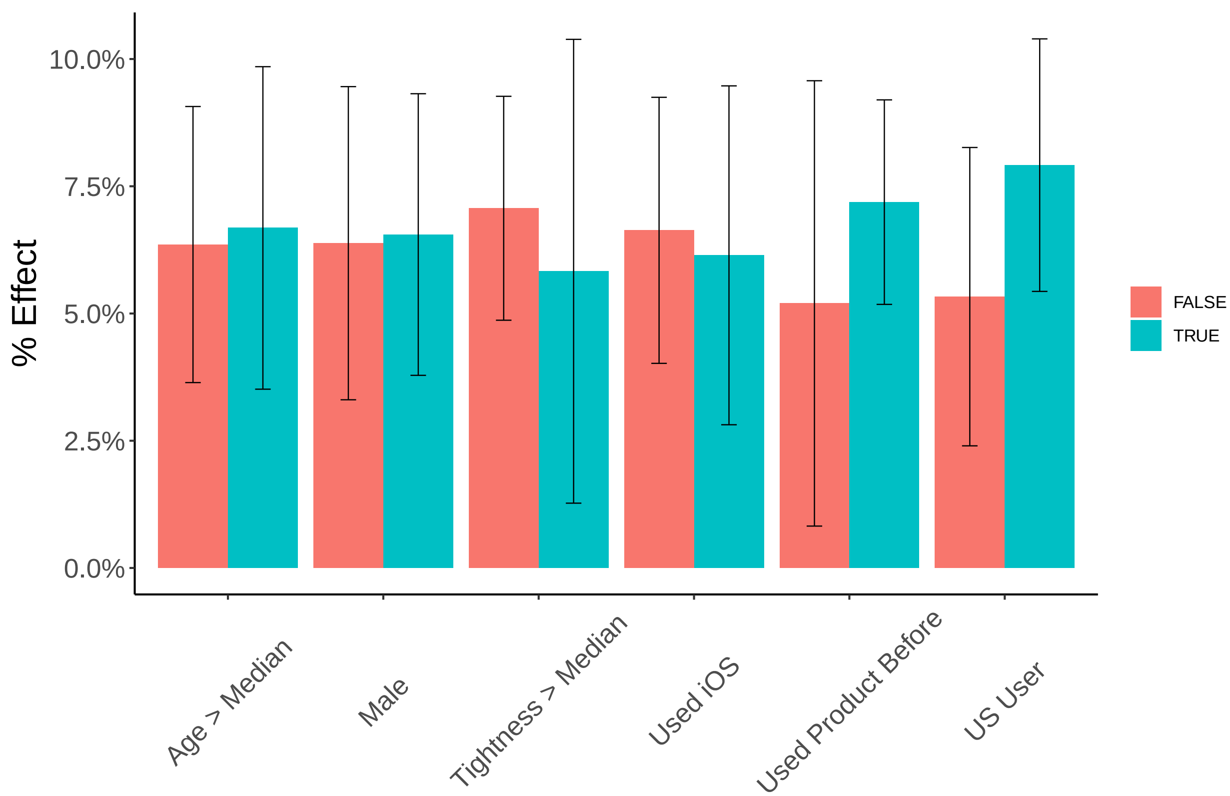

Table A.6: Treatment effects to under-subscribed vacancies — by skill requirement

	Has Application — < 5 Prior Apps		
	Low Skill	Medium Skill	High Skill
	(1)	(2)	(3)
Constant	0.0046***	0.0059***	0.0065***
	(0.0003)	(0.0003)	(0.0006)
Treatment	0.0009***	0.0012***	0.0014**
	(0.0003)	(0.0003)	(0.0007)
R^2	0.000	0.000	0.000
Observations	333,883	397,558	89,811

Notes: This tables displays the effects of information about prior applicants on application probabilities. Each observation is a seeker and the first vacancy they see in the list. For all regressions, just the vacancies that have < 5 applications are included. The three columns further limit the sample to vacancies that require either low, medium, or high skills. ** $p < 0.05$; *** $p < 0.01$.

A.5. Survey Choice Experiment

The pre-registered survey choice experiment consists of the following questions.[24] The first module asks about the employment status, age, and gender of a respondent, whether the respondent is actively looking forward, and an attention check. The survey then consists of three comparisons with two jobs each. Figure A.12 displays the three choice scenarios for one realization of the random draws. Each choice is between two companies, Blank Co and Brown Co, which differ in their wages, number of current applications, and an AI probability that the respondent gets an offer.

There are four elements of the survey that are randomized. First, some participants see information about an artificial intelligence (AI) probability that they receive an offer for a job while others do not. Whether this information is shown is randomized at the respondent level. We added an AI condition to see if we could make the effects of the number of applications less important by holding the probability of job offer constant.

There are three additional randomizations, one for each choice scenario. In each choice scenario, whether Blank Co or Brown Co has the lower number of applications is randomized at a question by respondent level.

There are then several post-choice scenario questions. For each choice in which a respondent answers that they either prefer Blank Co or Brown Co, the respondent is asked to explain their choice in a text box. Note that no such question is asked when the response is 'No Preference'. After the open text responses, we finish the survey by asking whether the participant responded randomly and whether the participant has feedback about the survey.

We now describe additional analysis details that we mentioned in our pre-registration. The experimental sample contained 1189 respondents, of which 592 were in the condition without an AI probability displayed. We investigate the experiment through regression analysis, displayed in Table A.7. The outcome in all of the specifications is a variable that is coded as -1 if Blank Co was chosen, 0 if there was no preference, and 1 if Brown Co was chosen. Column 1 displays the baseline regression with standard errors clustered at the participant level and shows that Brown Co is chosen more often when it has fewer applications.

Next, we consider the effect of information about the AI-predicted probability of an offer. Column 2 displays results with an interaction between the main treatment (lower applications) with whether the AI probability was shown. The coefficient on the interaction is negative, demonstrating that information about prior applications has less of an effect when the probability of an offer is known. However, there is still some effect of the information even in the AI condition. We can reject the null of no effect in the AI group with a Wald Test ($p < 3.4e\text{-}31$).

Lastly, we consider heterogeneous treatment effects. Column 3 displays the effect of the treatment separately for each comparison. We find that for each question, respondents prefer vacancies with fewer prior applications. Columns 4 and 5 estimate heterogeneous effects by gender and whether the respondent searched for a job in the past year. We find that there are no statistically significant differences in responses by gender, but that there are differences by whether the respondent searched for a job. In particular, those who searched for a job have a stronger preference for vacancies with fewer prior applicants than those who did not search for a job in the past year.

[24] The experiment was determined to be exempt from the IRB by MIT's Committee on the Use of Humans as Experimental Subjects. The pre-registration for the experiment is available here: `https://www.socialscienceregistry.org/trials/9344`.

Figure A.12: Survey choice questions

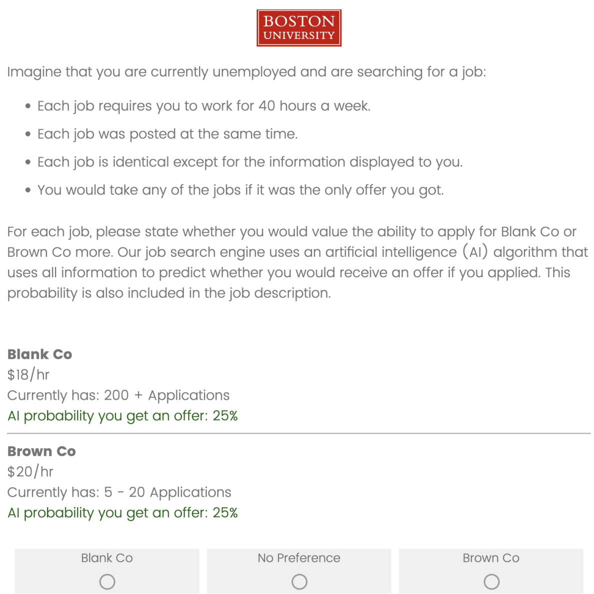
BOSTON UNIVERSITY

Imagine that you are currently unemployed and are searching for a job:

- Each job requires you to work for 40 hours a week.
- Each job was posted at the same time.
- Each job is identical except for the information displayed to you.
- You would take any of the jobs if it was the only offer you got.

For each job, please state whether you would value the ability to apply for Blank Co or Brown Co more. Our job search engine uses an artificial intelligence (AI) algorithm that uses all information to predict whether you would receive an offer if you applied. This probability is also included in the job description.

Blank Co
$18/hr
Currently has: 200 + Applications
AI probability you get an offer: 25%

Brown Co
$20/hr
Currently has: 5 - 20 Applications
AI probability you get an offer: 25%

Blank Co | No Preference | Brown Co

(a) Comparison Question 1

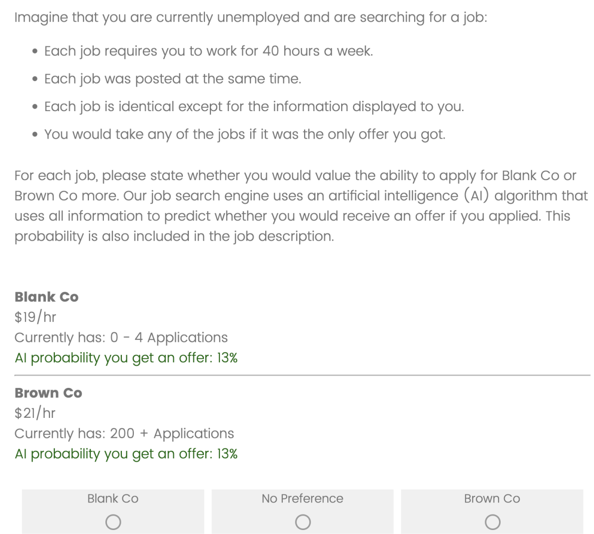
Imagine that you are currently unemployed and are searching for a job:

- Each job requires you to work for 40 hours a week.
- Each job was posted at the same time.
- Each job is identical except for the information displayed to you.
- You would take any of the jobs if it was the only offer you got.

For each job, please state whether you would value the ability to apply for Blank Co or Brown Co more. Our job search engine uses an artificial intelligence (AI) algorithm that uses all information to predict whether you would receive an offer if you applied. This probability is also included in the job description.

Blank Co
$19/hr
Currently has: 0 - 4 Applications
AI probability you get an offer: 13%

Brown Co
$21/hr
Currently has: 200 + Applications
AI probability you get an offer: 13%

Blank Co | No Preference | Brown Co

(b) Comparison Question 2

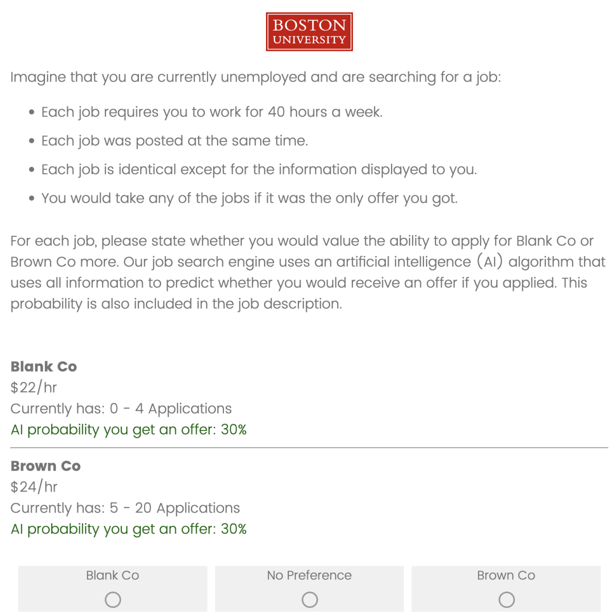
BOSTON UNIVERSITY

Imagine that you are currently unemployed and are searching for a job:

- Each job requires you to work for 40 hours a week.
- Each job was posted at the same time.
- Each job is identical except for the information displayed to you.
- You would take any of the jobs if it was the only offer you got.

For each job, please state whether you would value the ability to apply for Blank Co or Brown Co more. Our job search engine uses an artificial intelligence (AI) algorithm that uses all information to predict whether you would receive an offer if you applied. This probability is also included in the job description.

Blank Co
$22/hr
Currently has: 0 - 4 Applications
AI probability you get an offer: 30%

Brown Co
$24/hr
Currently has: 5 - 20 Applications
AI probability you get an offer: 30%

Blank Co | No Preference | Brown Co

(c) Comparison Question 3

Notes: Survey choice questions. Note that whether the higher application count was displayed for Blank Co or Brown Co was randomized at a question by respondent level. Whether the line about the AI probability was shown was randomized at a respondent level.

Table A.7: Survey Regressions

	Choice (-1 (Blank), 0, 1 (Brown))				
	(1)	(2)	(3)	(4)	(5)
Constant	-0.0238	-0.0238		-0.0136	0.0357
	(0.0364)	(0.0363)		(0.0501)	(0.0576)
Fewer Applications (Brown)	0.8704***	0.8704***		0.8469***	0.7494***
	(0.0422)	(0.0422)		(0.0587)	(0.0692)
AI Probability		0.4051***			
		(0.0517)			
Fewer Applications (Brown) × AI Probability		-0.3736***			
		(0.0583)			
Fewer Applications (Brown) × Question = 1			0.9806***		
			(0.0585)		
Fewer Applications (Brown) × Question = 2			1.140***		
			(0.0610)		
Fewer Applications (Brown) × Question = 3			0.4925***		
			(0.0558)		
Male				-0.0203	
				(0.0727)	
Fewer Applications (Brown) × Male				0.0450	
				(0.0844)	
Searched for Job					-0.1012
					(0.0742)
Fewer Applications (Brown) × Searched for Job					0.2076**
					(0.0867)
Observations	1,776	3,567	1,776	1,776	1,776
R^2	0.25896	0.21245	0.31535	0.25913	0.26254
Within R^2			0.29417		
Sample	No AI	All	No AI	No AI	No AI
Question fixed effects			x		

*Notes: The outcome for all regressions is a variable that is coded as -1 if Blank Co was chosen, 0 if there was no preference, and 1 if Brown Co was chosen. ** $p < 0.05$; *** $p < 0.01$.*

Figure A.13: Distribution of responses
set of choices with AI probability displayed