Online Appendix

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

Andrey Fradkin

Boston University

Monica Bhole

Meta

John J. Horton

MIT Sloan and NBER

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

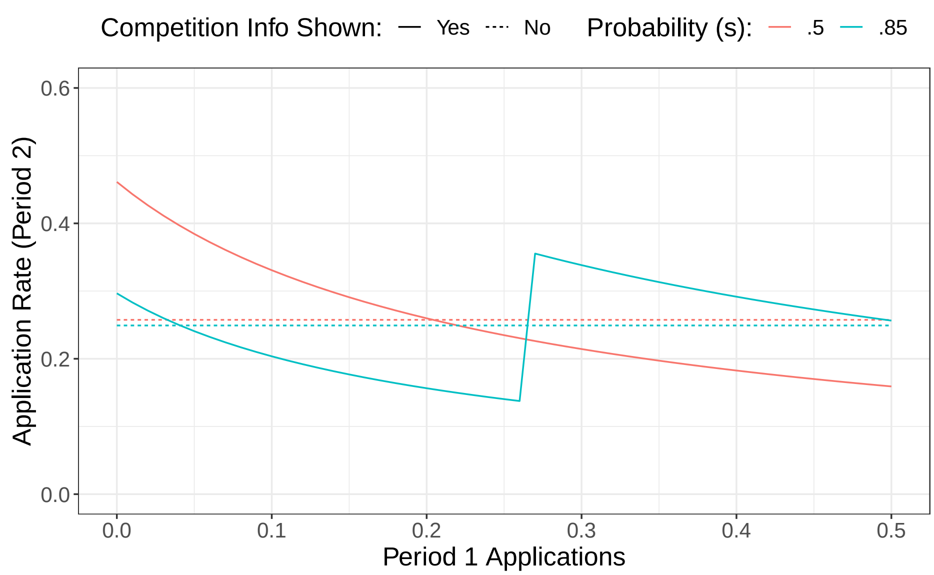

Notes: This figure plots the equilibrium period 2 application rate, as well as the response of an infinitesimal seeker who is given information about period 1 applications. We assume that $F \sim Exp(.5)$, $m = x + .1$, $G \sim U[0,1]$, and $\lambda = .5$.

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

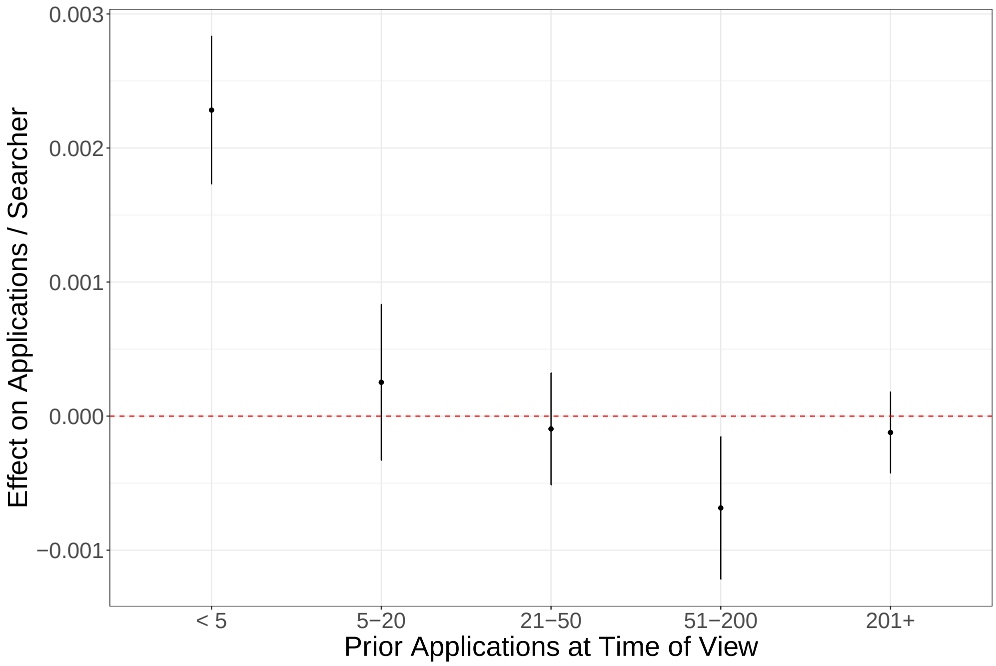

Notes: This figure plots the average effect of the treatment (pooled across all arms) in levels. Each observation is a searcher in an experiment and all treatments that included information about prior applications were pooled.

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

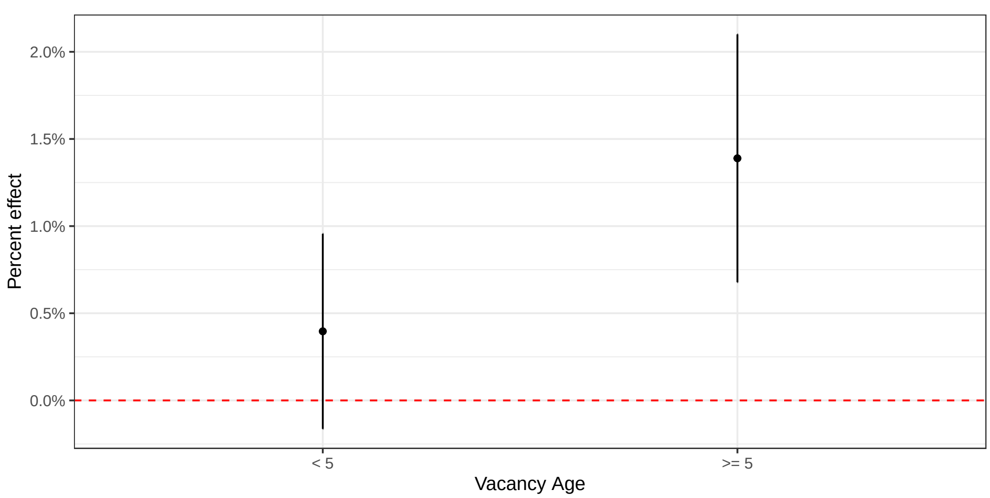

Notes: This figure plots the average effects in percent terms and their 95% confidence intervals. Each observation is a searcher in the experiments who was either in the control group or in the treatment group for which competition information was shown.

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

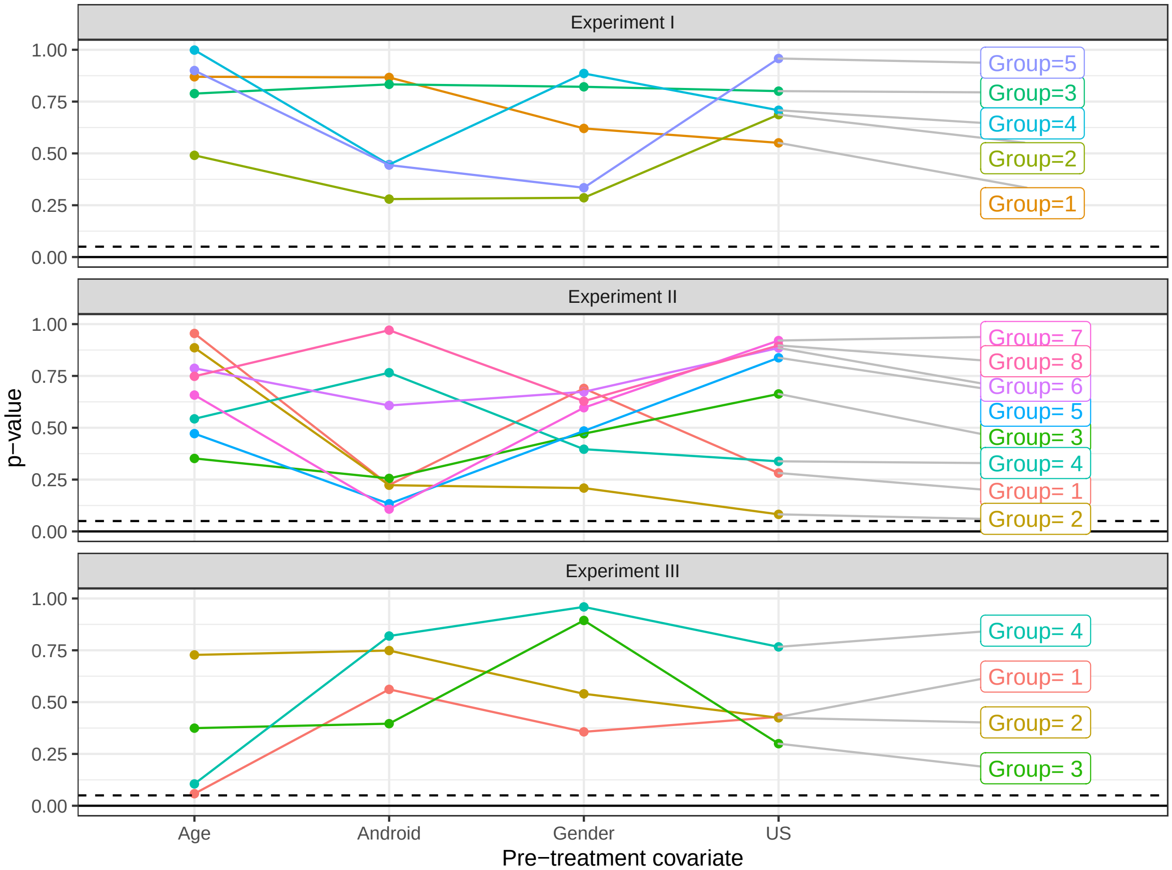

Notes: This figure displays the p-value from a linear regression where the characteristics of each treatment arm were compared with the control arm of the corresponding experiment.

Figure A.6: Treatment Effects for Experiment 1

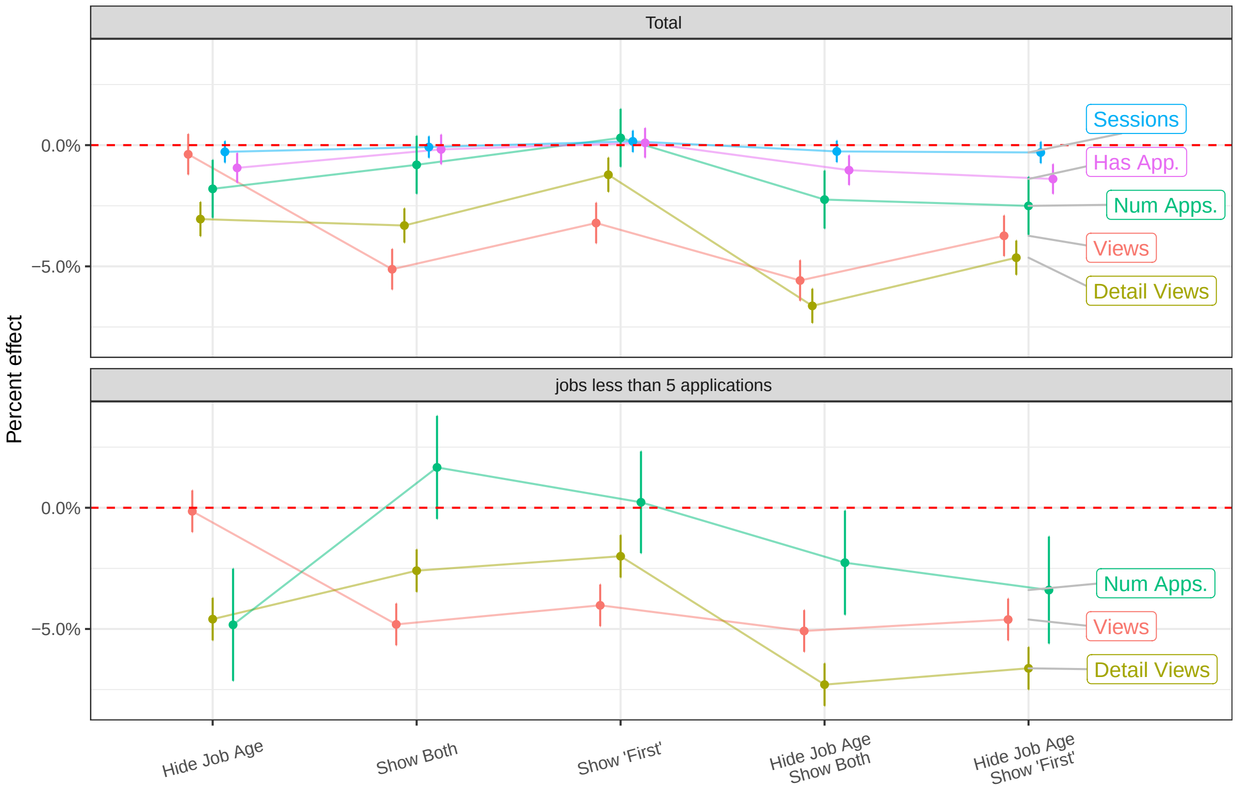

A.3.1. Effects by Treatment Arm Next, we discuss the by-arm treatment effects for each treatment and experiment. We begin with Experiment 1 (Figure A.6). Columns 1, 4, and 5 of of the figure plot the treatment effects where the vacancy age is hidden. Columns 2 - 5 plot treatments where competition information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

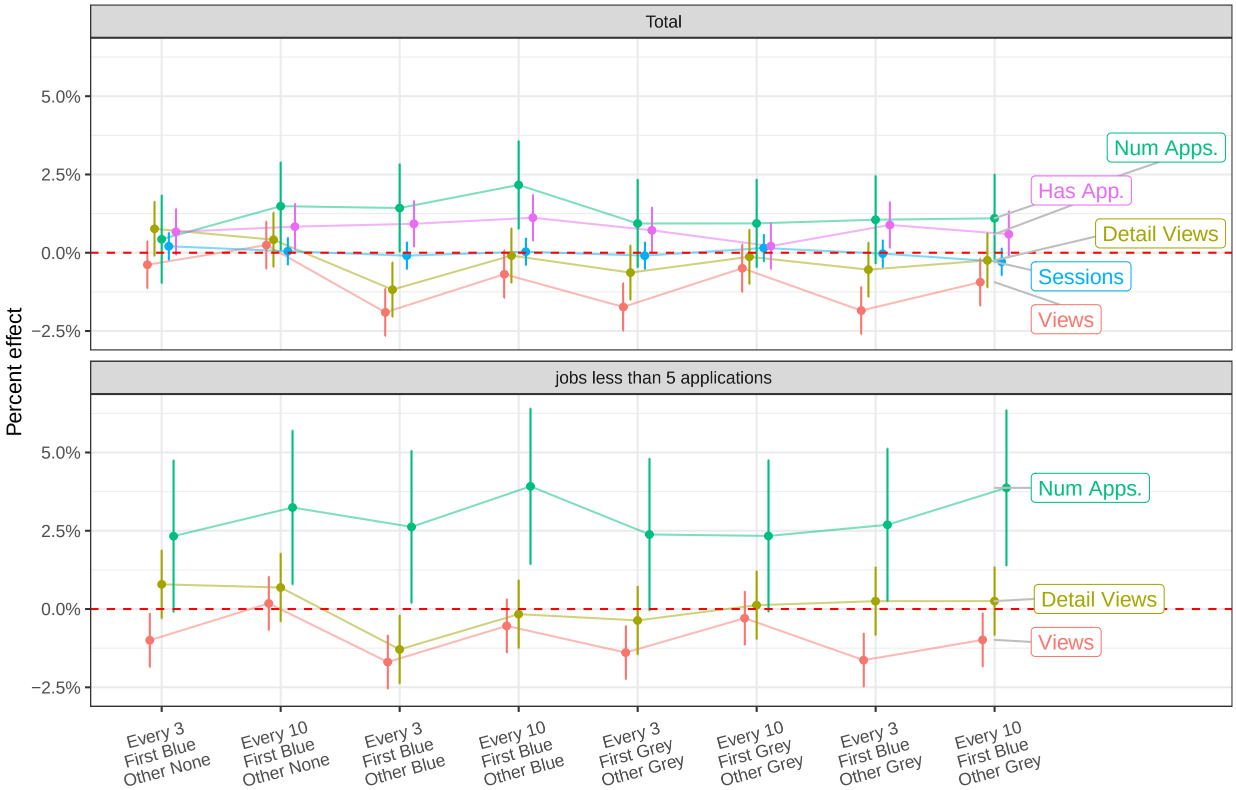

Table A.4: Treatment effects pooled across all three experiments

	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

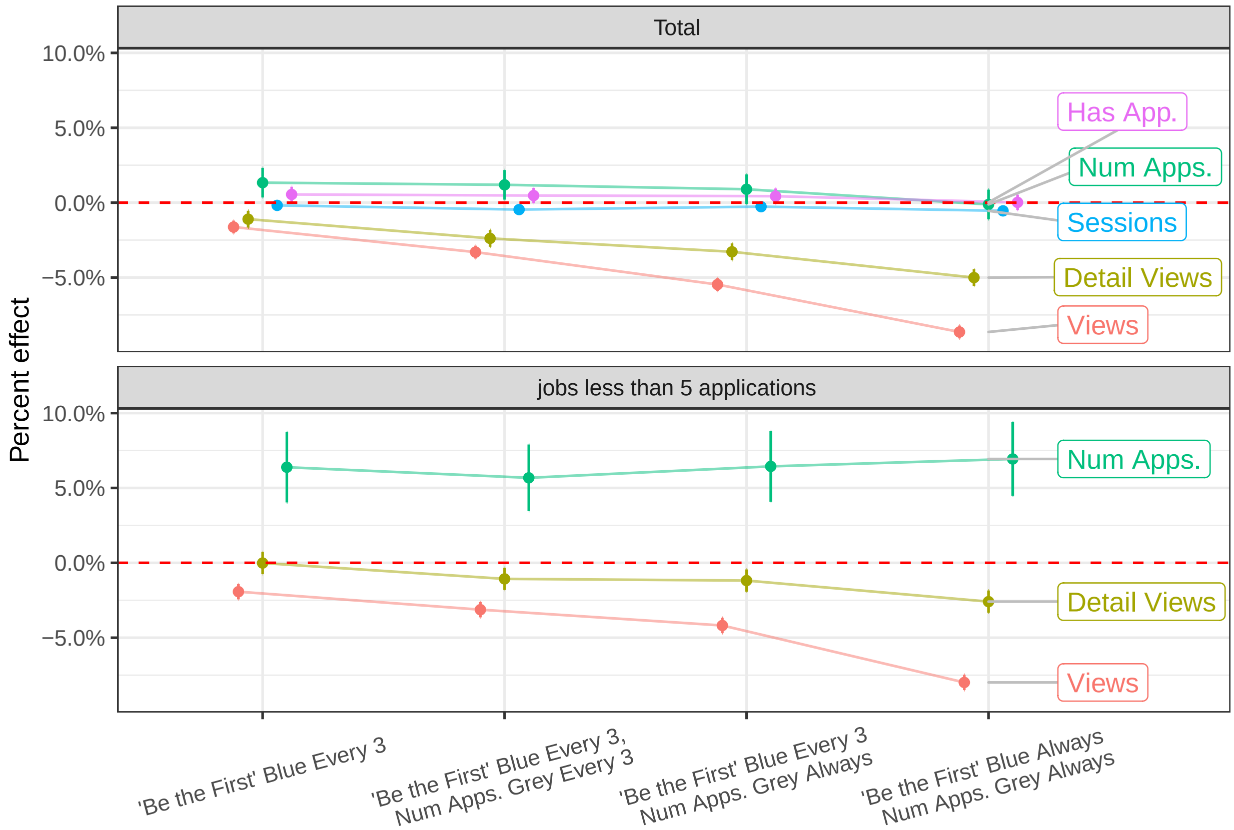

Table A.5: Treatment effects pooled across all three experiments - by application type

	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

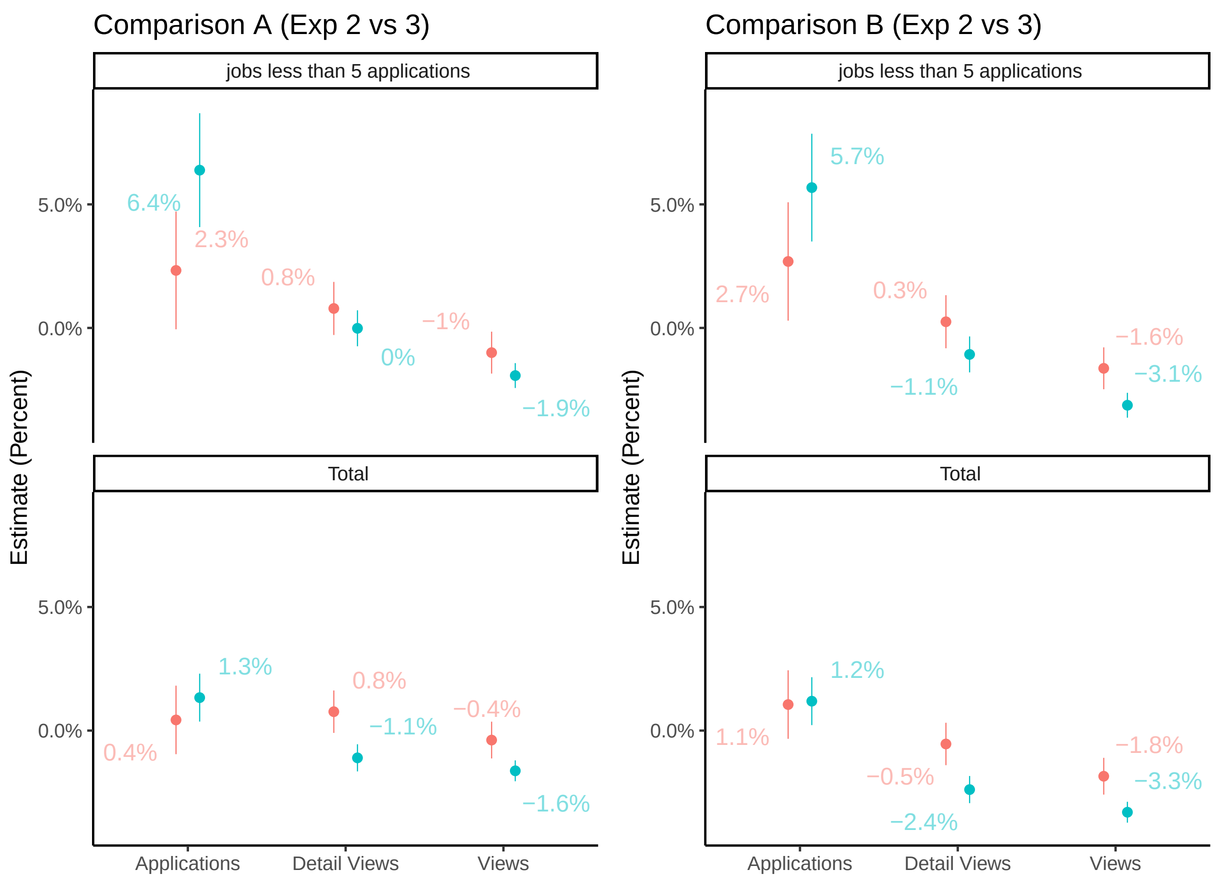

We also investigate whether there is heterogeneity based on the skill requirements of the vacancy. To do this, we use a skill requirement classification of vacancies into low, medium, and high skill that is available in the data. We consider the first vacancy exposed to each user and condition on the 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

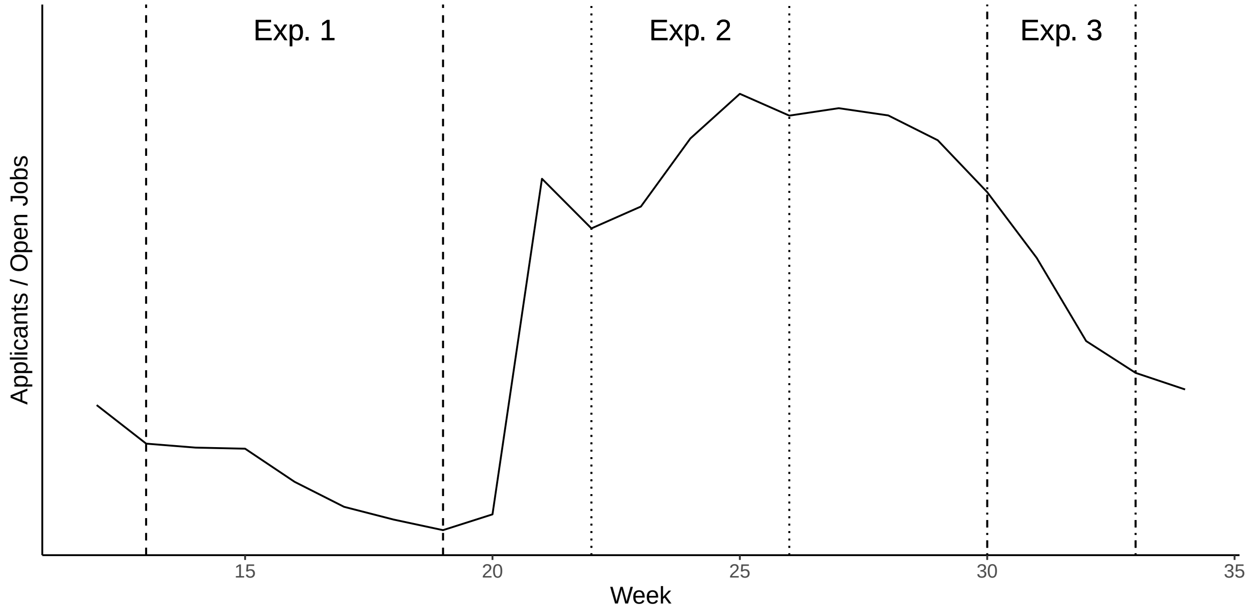

Figure A.11: Heterogeneous treatment effects - Applications to under-subscribed vacancies

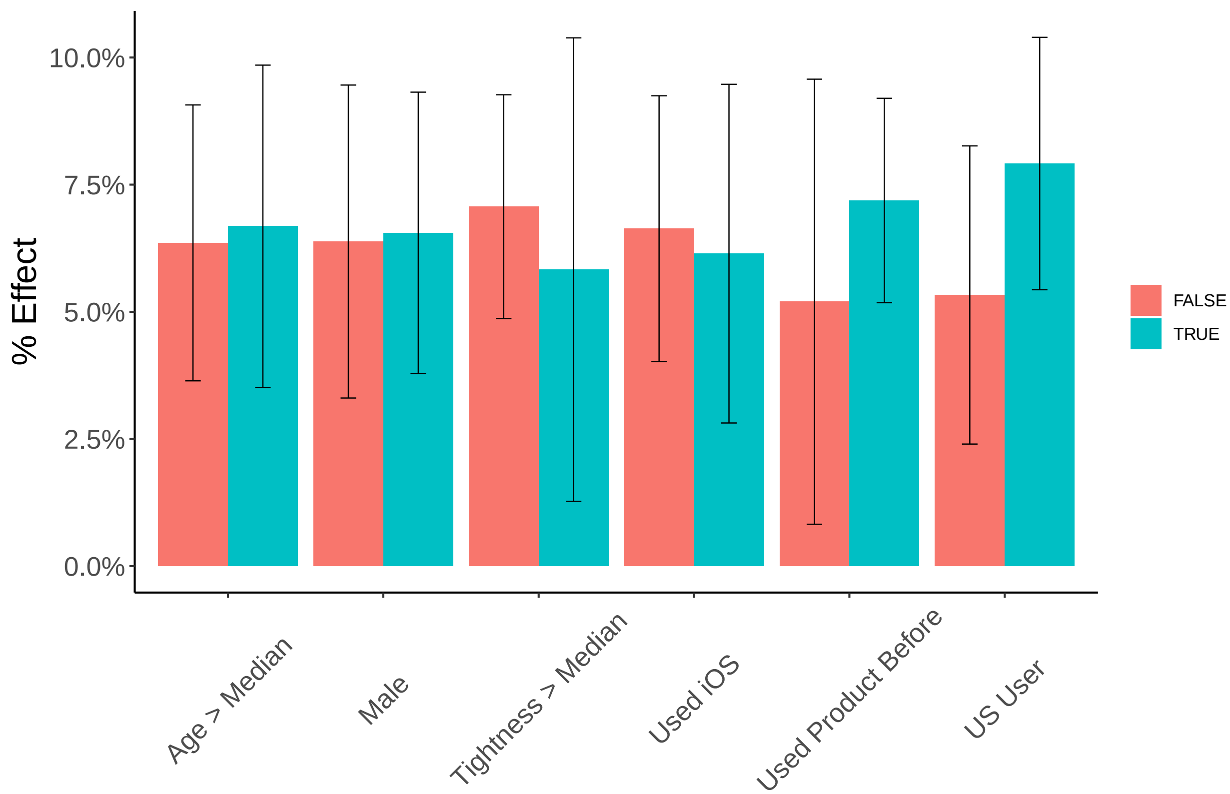

Table A.6: Treatment effects to under-subscribed vacancies — by skill requirement

	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

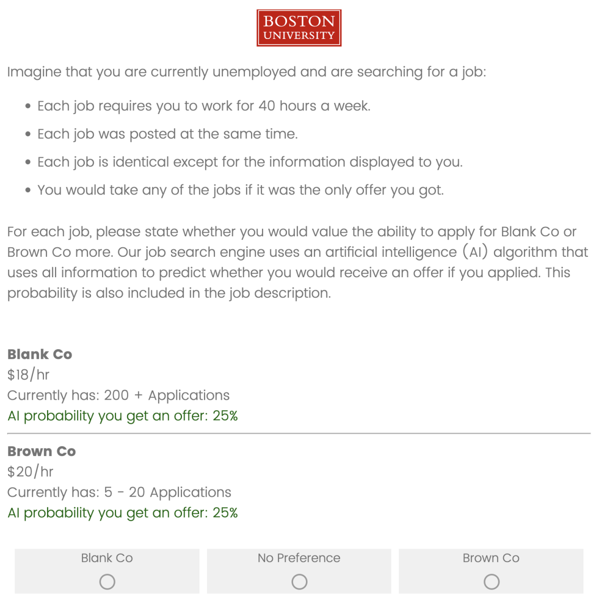

(a) Comparison Question 1

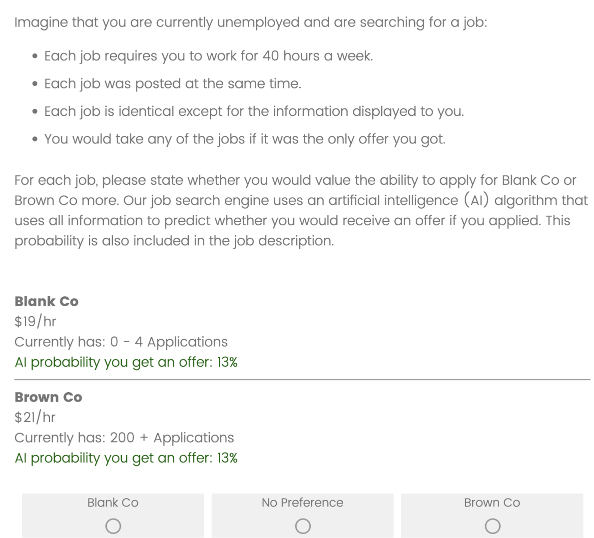

(b) Comparison Question 2

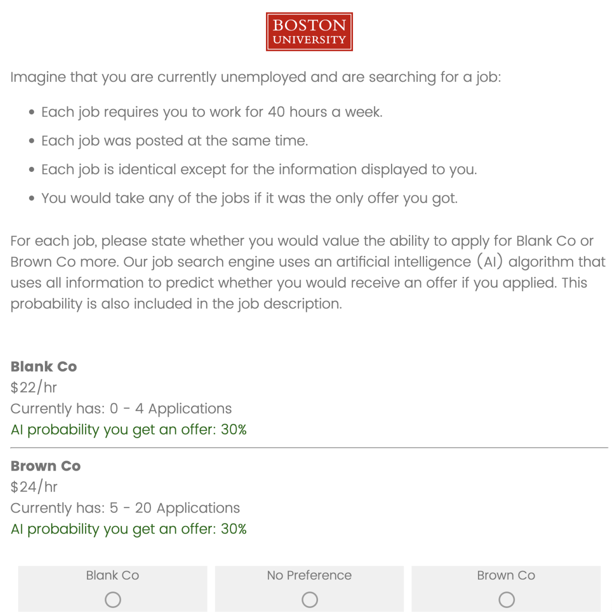

(c) Comparison Question 3

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