

Behavioral Responses to Algorithmic Matching: Experimental Evidence from an Online Platform

Sarah Bana and Kevin Boudreau*

May 26, 2023

Abstract

Data-driven algorithmic and AI-based matching systems have become crucial features of modern online labor markets, given the inherent complexity and uncertainty of matching workers to jobs. Whereas considerable prior research focuses on optimal algorithmic design, in this study we investigate whether machine-based automated systems affects alter jobseekers' own behavior and choices. For example, on the one hand, if scalable machine-based recommendation systems affect labor market competition, then rational job seekers should anticipate these changes and alter their choices. If, on the other hand, job seekers' behavior could be affected by idiosyncratic preferences towards AI. If so, the probability of a match will be affected not only by the inherent quality of match, but also by individuals' behavioral response per se to the use of machine-based automated recommendations and decision rules. We report on a field experiment involving 4,534 subjects who received job matching recommendations. Subjects were randomized as to whether they were made aware of the use of algorithmic matching in identifying the match. We find that simply being made aware of the algorithmic source of the recommendation has an overall negative population-wide effect on interest in pursuing the opportunity. However, the effect varies, depending on applicant's disciplinary training. Thus, the experiment demonstrates the existence of systematic behavioral changes in response to the use of algorithmic systems in the context of labor markets. As we discuss herein, the results are consistent with both "algorithm aversion" along with a rational response to how AI shapes labor market competition. Thus, we highlight the complex interplay between human behavior and machine-based recommendations in the labor market. We validate these findings and interpretation with supplementary survey evidence.

Keywords: Algorithmic Matching, Labor Markets, Matching Platforms, Field Experiment

*Bana: Chapman University and Stanford University. Boudreau: Northeastern University and the National Bureau of Economic Research (NBER). We are very grateful to the participants at the Workshop on Information Systems and Economics (WISE) 2022 for helpful comments.

1 Introduction

Labor markets are complex and uncertain. Finding the right person for a job often requires insight into multiple dimensions of skills and aptitudes, behavioral traits, and non-pecuniary preferences of candidates. Apart from the use of human sources (e.g. recruiting agencies and headhunters, personal networks, interviews), labor market matching can be facilitated and enabled by machine-based, data-driven algorithms (Horton, 2017). Advances in cost and power of computation along with advances in research and industrial practice within machine learning and algorithm design (Zhang et al., 2022) have further enabled these automated, algorithm-based methods of matching. The rise of machine-based methods has also been complementary to and given rise to online platform-based labor markets (e.g., Fivver, Task Rabbit, oDesk, Elance, Guru, Freelancer.com).¹ To the extent that rich data sources and powerful algorithmic matching have been automated, these markets now support especially fast-moving spot markets for labor, and markets for increasingly high-skilled positions, where the search for suitable candidates requires discerning among complex profiles with nuanced differences (The Economist, 2023).

For the most part, the rise of automated algorithms and machine-based matching can be understood as a technological advance – an improvement in information-generation or noise-reduction—in especially complex markets. In this sense, to the extent that data or the algorithms exploiting these data improve, there can be increases in efficiency. For this reason, a great bulk of attention in research and practice has been focused on improving this “technology” of matching. However, outside of research on labor markets, researchers in a variety of fields have begun to uncover instances in which humans interacting with machine-based and AI systems simply behave differently than they would otherwise. If the presence and use of automated algorithmic “AI” systems were to generate differences in job-seekers behaviors, the effectiveness of these matching systems will depend on more than just the information generated by these systems. In this paper, we seek and find evidence of the existence of behavioral changes when job seekers are aware of interacting with automated algorithmic market recommendations, even when holding constant the quality of recommendations they receive. We also attempt to interpret the nature and underlying mechanisms of these behavioral changes.

Isolating the behavioral impact of algorithmic recommendations poses a challenge: deploying an algorithmic recommendation system likely changes the quality of the recommendations as well. To overcome this obstacle, we design a field experiment on a job-matching platform. Employers on this platform select a series of keywords that match with the role, which creates pools of candidates based on keywords in their LinkedIn profiles. These candidates are simultaneously, therefore, selected via algorithm, i.e. keyword search, and employer, allowing us to truthfully randomize them into three conditions: an algorithmic condition, an employer condition, and a control condition. Candidates receive an email encouraging them to apply for a role, indicating the recommendation source (algorithm, employer, or undisclosed). The primary outcome variable in our study is the probability of

¹The digitization of labor markets in online platforms and the rise of automated algorithmic matching is also intimately related to the use of online reputation systems (Filippas et al., 2022).

clicking on the link, a measure of potential intention to apply to the role.

Theoretically, we identify two primary reasons why candidates may respond differently to AI-based algorithms. The first is preferences: candidates may inherently like or dislike machine generated recommendations. This phenomenon, often referred to in the literature as algorithm aversion, suggests people will avoid algorithms they recognize to be imperfect (Dietvorst et al., 2015). While other research has found evidence of these preferences in consumer contexts, such as chatbots (Luo et al., 2019), there is no similar research in the labor market context.

Beyond a simple preference based explanation, candidates may rationally perceive AI as a technology with certain features, and therefore change their behavior due to this different technology being employed. The first well-known feature of AI is scale - AI systems ingest large volumes of data. Differently from human recruiters, an AI-based recommendation could arrive instantaneously, and for more candidates. Thus, a candidate may perceive greater competition in the case of an AI based recommendation. A candidate may also perceive AI as producing higher or lower precision recommendations. This is distinct from the preference channel as it speaks to the perceived returns from pursuing the recommendation. We call this a quality effect.

We find that there is a substantive negative effect generated by the disclosure of AI-based matching systems, above and beyond the information and predictions generated. Candidates are 28 percent (8.2 percentage points) less likely to click an advertisement with the disclosure that the source of the recommendation is an AI. Moreover, this effect is significantly different from the effect of disclosing that the source is an employer, which is virtually identical to disclosing no source of recommendation.

This main effect can be consistent with three potential explanations: algorithmic aversion, candidates believing that the AI produces lower quality recommendations, and candidates rationally expecting their chances to be lower given the greater scaling capacity of AI. We estimate the main effect across different subgroups to differentiate between these explanations. We find that STEM candidates display a much smaller difference between the AI source and no source treatments. This finding can be consistent with algorithmic aversion, with STEM candidates displaying less aversion due to greater understanding of the technology.

We also find that STEM candidates are more likely to click given the Employer recommendation, evidence that they may understand that employer resources scale more slowly than algorithmic resources. Finally, we compare candidates in fields that each employer views as a good match for the role with candidates in less aligned fields. The high quality candidates are just as likely to be dissuaded by the AI recommendation source as the other candidates, suggesting our findings are not consistent with candidates believing that AI produces lower quality recommendations.

We plan to supplement the field experiment with a survey of candidates, measuring their perceptions with regards to the number of applicants invited for the role to distinguish between competing explanations.

2 Conceptual Background and Hypothesis Development

In the past decade, AI-based matching systems have proliferated labor market platforms. While [Horton \(2017\)](#) considers algorithmic recommendations rare in 2016, more recent work by [Fuller et al. \(2021\)](#) finds that 63 percent of all employers surveyed use a Recruitment Management System (RMS) and over 90 percent of those employers used their RMS to initially filter or rank potential candidates. This functionality also exists on a wide variety of job search platforms. While important work has measured the effect of algorithmic recommendations on the fill rate and other consequential labor market outcomes, we are unaware of work that differentiates between the valuable information provided by the recommendation itself and the behavioral response to such a recommendation.

While algorithms have been deployed on labor market platforms, candidates may not necessarily be aware of the pervasiveness of these tools. However, there are new calls for greater disclosure, for example proposing that autonomous agents reveal themselves upon request ([Potts et al., 2021](#)). This has manifested in the form of legislation: Illinois passed HB 2557, The Artificial Intelligence Video Interview Act, which requires employers who use artificial intelligence to analyze video interviews to provide notice ([Waltz et al., 2019](#)). Will candidates become less likely to apply to jobs that contain these disclaimers? The behavioral effects might be a potential barrier to the progress that AI-based systems may bring. We hypothesize that

Hypothesis 1 (H1 Overall Effect): *Candidates randomized into the algorithm treatment are less likely to click on the link to the role.*

There are at least two broad views suggesting the possibility that the use of AI-based matching systems could affect the behavior of job-seekers, holding constant any information about matching generated by these systems. A first comes from the existing evidence and literature on human behavioral responses while interacting with AI (automated algorithms, machine-based systems, robots). Notably, the existing literature on changes in behavior in response to AI –from a variety of contexts outside of labor markets. For example, [Luo et al. \(2019\)](#) finds that the disclosure of chatbot identity reduces purchase rates by close to 80 percent. This effect is driven by humans perceiving disclosed bots as less knowledgeable and less empathetic, a subjective human perception rather than the objective competence of the chatbot. This response is consistent with [Hidalgo et al. \(2021\)](#), who finds that humans judge humans and machines in similar scenarios differently. For the most part, this existing evidence and interpretation suggests changes of behavior rooted in preferences (often aversion) for working with machine-based systems, in some cases perhaps moderated by the background and training of the subject in question. Any such per-se preference or aversions to AI might play some role in AI-managed labor markets, perhaps creating greater hesitation or greater eagerness to pursue AI-identified matches and opportunities.

Because algorithmic aversion appears to be more common than algorithm appreciation – see, e.g. [Altintas et al. \(2023\)](#) for a more extensive list of relevant research – we hypothesize that

Hypothesis 2 (H2 Preference Effect): *Non-STEM candidates are less familiar with algorithms, and display more aversion if in the algorithm treatment.*

A second and separate view from that above examines how AI is a novel technology for producing relevant information and might impact the rational choices and behaviors of job seekers. Here, we outline two different potential effects: a scale effect and a quality effect. Job seekers searching and competing for jobs might form different expectations about their likelihood of pursuing certain positions if the recommendation is generated by an algorithm rather than humans. AI-based matching, which draws on large data sets, inexhaustible algorithms and automation, may well be more scalable and instantaneous than human-based matching and recommendations. It is plausible then that job seekers receiving AI-based recommendations could expect to face greater competition than would be the case with human recommenders.

Hypothesis 3 (H3 Scale Effect): *Because AI matching is more scalable than employer matching, candidates in the algorithm treatment are less likely to click the link than candidates in the employer treatment.*

This scale effect may differ from expectations or perceptions of the quality of the recommendations. If job seekers believe AI generates superior information and predictions, it could lead superior candidates to believe there are higher returns to actively pursuing recommended positions. This view is potentially consistent with the findings of [Liel and Zalmanson \(2022\)](#), who measure gig economy workers' propensity to conform to algorithmic recommendations. If, on the other hand, if the AI is believed to be inferior in making predictions of matches, it could dissuade superior candidates from pursuing a recommended position relative to what would be the case with a recommendation generated by human intelligence.

While it is challenging to measure this quality effect because it requires information about the candidates' beliefs, not captured here in our current study, we might expect that candidates in fields that the employers view as good matches for the role should be less susceptible to perceptions about algorithms. This quality related hypothesis is outlined as follows:

Hypothesis 4 (H4 Quality Effect): *High quality candidates will identify themselves as such, and therefore are be less sensitive to algorithmic treatment.*

With these four hypotheses in mind, we describe the research design.

3 Field Experiment Research Design

To make progress in better understanding the existence and nature of possible changes in behavior of job seekers in response to AI-based matching systems, we designed and deployed a field experiment on a job-matching platform. The gist of the research design was to identify pools of candidates suitable for a series of jobs and to then randomize whether candidates were informed that the recommendations they receive were generated by an algorithm, or a human manager, in comparison with a control group in which neither was stated. In using this approach, we can then effectively observe whether and how individuals diverged in their likelihood of pursuing an opportunity depending on

the source of the recommendation, while holding constant the quality or information-value of a recommendation. (A central challenge in this design, discussed below, is to implement this idea in a way that avoids any deception of or misrepresentation to subjects.) The remainder of this section elaborates on the details of the research design.

3.1 Field Experimental Context & Study Population

We conducted this study on a university-based labor market platform at a top-40 ranked US university. Job seekers on this platform therefore include undergraduate and graduate students. Table 1 demonstrates that 73 percent of students are pursuing their bachelors’ degrees, 14 percent are pursuing masters’ degrees, and the remainder are pursuing doctoral or professional (J.D.) degrees. This study involved 4,534 job seekers from a range of disciplinary fields of education. The nature of work on this platform is typical of short term project or “gig” work that is now common on online platforms, with a particular focus on skilled tasks associated with disciplinary postsecondary training in fields such as engineering, psychology, media, computer science, data science, and business analysis tasks.

3.2 Experimental Protocol

Students were randomized into one of three conditions about the source of recommendation: no information, an employer generated list, or an A.I. matching algorithm. Example emails are available in the appendix of the manuscript. The actual selection was made through the employer generating relevant keywords to match to publicly available LinkedIn profiles, which allowed both statements to be true. For example, for the job related to designing a data model for an algorithm design competition platform, students with “Information Technology” or “Data Analytics Engineering” in their profiles were those recommended the role. A total of eight jobs were included in this iteration of the experiment. Students were also randomized to the wage offered in the role. The true wage on this platform is determined by a formula at the institution level. The estimated wages are all in the range of this formula. The students were randomized into no statement about wage in the email or an estimated hourly wage ranging between \$16 and \$28. These two treatments allow us to estimate the size of the click rate for the recommendation source relative to wage.

4 Empirical Approach

To estimate the effect of the recommendation source on the click rate, we run the following regression at the subject level:

$$Click_i = \beta_0 + \beta_1 Employer_i + \beta_2 AI_i + \lambda_{WageBins} + \varepsilon_i \quad (1)$$

where $Click_i$ is a binary indicator for whether the student clicked on the posting, a measure of

interest, $Employer_i$ is an indicator for whether the student was randomized into the employer source condition, AI_i is an indicator for whether the student was randomized into the AI recommendation source condition, and $\lambda_{WageBins}$ are indicators for whether the student was randomized into each of the estimated wage conditions (\$16, \$18, \$20, \$22, \$24, \$26, or \$28). The omitted categories are no recommendation source and no estimated wage. Robust standard errors are reported.

5 Results

The main specification, described above, is estimated in Figure 1. Panel (a) describes the effect on clicks, and panel (b) describes the effect on clicks conditional on opens.

The effect of the AI treatment is 2.3 percentage points (significant at the 1 percent level). For context, the unconditional probability of click is .082 or 8.2 percentage points. This is a 28 percent decrease in probability of clicking the advertisement with the disclosure that the source of the recommendation is the AI. This effect size is similar conditional on opening the email: the AI recommendation source decreases the probability of clicking by 6.1 percentage points. Given the unconditional probability of clicking is 20 percent, this amounts to an approximately 30 percent decrease. This is strong support of our first hypothesis, an overall behavioral effect.

On the other hand, there is an insignificant effect of the Employer treatment on the probability of clicking the advertisement. In the unconditional specification, the effect of the employer treatment is 0.000993, consistent with job seekers assuming that the employer has developed a recommendation if no additional information is given.

This main effect regression also allows us to test Hypothesis 3, the scale effect. If the scale effect exists, then candidates should respond differently between the employer treatment and the algorithm treatment. This is what we find: a t-test comparing the employer treatment and the algorithm treatment yields a p-value of 0.0126 in the case of the unconditional regression, and a p-value of 0.0242 in the case of the regression conditional on opens. This is consistent with candidates recognizing that a recommendation by employers is a stronger signal than a recommendation by an algorithm.

We then separately estimate the main regression on the eight job postings involved in this experiment in Table 2. The AI treatment causes the biggest significant decreases for the Video posting and the Strategy posting unconditionally, and for the Strategy posting conditionally. For some other postings, the effect size is negative and economically meaningful, but not significant because of the few participants to which the posting is targeted.

A behavioral response to algorithmic matching may differ by student characteristics. Specifically, if a student is more technically oriented and understanding of algorithmic matching, then they may have less of a negative response to advertisements coming from algorithms. In contrast, a student who might be less technically oriented may view the algorithm with more suspicion. We investigate this possibility by connecting the field of studies listed on LinkedIn to the university's instructional program codes, which determine whether the field of study is considered science,

technology, engineering and math (STEM). This is a well-defined designation and serves as the basis for immigration extensions ([Citizenship and Services, 2022](#)).

We measure the Jaro-Winkler distance between the university’s instructional programs and the LinkedIn field names, and are able to categorize 90 percent of the LinkedIn fields of study with a similarity score of 0.85 or above. The results are displayed in Table 3. Column (1) is the original estimation, and Column (2) demonstrates that the effect size is quite similar between the matched sample and the original. Column (3) and Column (4) split the sample, and estimate the effect students in STEM and non-STEM programs, respectively. Column (4) shows that the non-STEM students, despite being a smaller portion of the sample, are driving the effect. A non-STEM student is 4.79 percentage points less likely to click the AI source. Recall the original effect is 2.33 percentage points, so this is almost twice that size. This result supports hypothesis two, on preferences. The Non-STEM candidates are more likely to be deterred by the algorithm treatment, suggesting aversion.

Column (5) is a fully interacted model. STEM students are 3.25 percentage points (though not significant) more likely to click a link when they see the algorithm source than non-STEM students.

We continue to probe into differences in student characteristics in Figure 2, where we subsample by year in program, start year and end year at the institution, and degree type. The effect of the AI treatment are strongest for those who started at the institution in 2020, are finishing in 2024, and are in their fourth year at the institution. In Figure 3, we subsample by field of study. Because field of study is a write-in field on LinkedIn, we use a hierarchical clustering algorithm to define the clusters. We limit clusters analyzed to those with at least 100 students. We find that less technical fields like psychology, information systems, and business administration, have more negative effects of the AI treatment, while more technical fields like behavioral neuroscience, computer science, and mechanical engineering have little effects of the AI treatment, supporting our STEM analysis described above.

Candidates may perceive the employer and the AI as producing recommendations of differing quality. One interpretation of our main effect could be that applicants see the AI as producing a lower quality recommendation. This would make the marginal candidate less likely to apply. The treatment should not affect high quality candidates: high quality candidates should have the same probability of applying in both treatments, but lower quality candidates may click with lower probability in the AI treatment. Importantly, this could be distinct from the case that all candidates having an aversion to the AI, uncorrelated with the quality of the match for the role.

We test the fourth hypothesis by asking the employer to characterize fields of study that are high quality matches for the role.² We then analyze the effects for the high quality candidates and for all other candidates.

The results are displayed in Figure 4. The effects of the AI and Employer treatments look very similar for high quality candidates and all other candidates. This appears to be inconsistent with

²The employer conducts this exercise for each posting, and the number of fields and percent of candidates that fall into the high quality category differ by posting.

viewing the AI recommendations as lower quality, as even high quality candidates are less likely to click these postings. We reject hypothesis four.

6 Survey

Thus far, we have been testing perception of candidates indirectly. However, perceptions of scale and quality may be directly tested through a survey of the candidates targeted for these messages.

7 Conclusion

Literature on algorithmic recommendations has focused largely on the technology and not on the behavioral consequences of such recommendations. This paper separates out the behavioral effect from the effect of the technology by conducting a field experiment on an online labor market platform that allows us to randomize participants into three conditions. We find that candidates are far less likely to click on a job posting when it is disclosed that the recommendation came from an algorithm. We identify three potential channels that could be consistent with this behavior: preferences, knowledge of algorithms scaling more widely than humans, and perceptions of algorithm quality. Our results are consistent with algorithm aversion and knowledge of scaling, but not differential perception of algorithm quality.

While algorithms have substantial promise for improving matching in labor markets, disclosure of such technologies yields a behavioral response. This is important as it suggests that the mere presence of algorithmic recommendations might not be enough to fully capitalize on their potential benefits. Instead, understanding and addressing the underlying factors driving these behavioral responses, such as algorithm aversion and knowledge of scaling, is crucial for optimizing the adoption and effectiveness of AI-driven matching systems in labor markets.

References

- ALTINTAS, O., A. A. SEIDMANN, AND B. GU (2023): “The Effect of Interpretable Artificial Intelligence on Repeated Managerial Decision-Making Under Uncertainty,” *SSRN*, available at SSRN: <https://ssrn.com/abstract=4331145> or <http://dx.doi.org/10.2139/ssrn.4331145>.
- CITIZENSHIP, U. AND I. SERVICES (2022): “Optional Practical Training Extension for STEM Students (STEM OPT),” <https://www.uscis.gov/working-in-the-united-states/students-and-exchange-visitors/optional-practical-training-extension-for-stem-students-stem-opt>, accessed: 2022-10-01.
- DIETVORST, B. J., J. P. SIMMONS, AND C. MASSEY (2015): “Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err,” *Journal of Experimental Psychology: General*, 144, 114–126.

- FILIPPAS, A., J. J. HORTON, AND J. M. GOLDEN (2022): “Reputation Inflation,” *Marketing Science*, 41, 305–317.
- FULLER, J., M. RAMAN, E. SAGE-GAVIN, K. HINES, ET AL. (2021): “Hidden Workers: Untapped Talent,” .
- HIDALGO, C. A., D. ORGHIAN, J. ALBO CANALS, F. DE ALMEIDA, AND N. MARTÍN (2021): *How Humans Judge Machines*, MIT Press.
- HORTON, J. J. (2017): “The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment,” *Journal of Labor Economics*, 35, 345–385.
- LIEL, Y. AND L. ZALMANSON (2022): “Turning Off Your Better Judgment -Conformity to Algorithmic Recommendations,” .
- LUO, X., S. TONG, Z. FANG, AND Z. QU (2019): “Frontiers: Machines vs. Humans: The Impact of Artificial Intelligence Chatbot Disclosure on Customer Purchases,” *Marketing Science*, 38, 937–947.
- POTTS, C., M.-F. CUÉLLAR, J. DEGEN, M. C. FRANK, N. D. GOODMAN, T. ICARD, AND D. SADIGH (2021): “Shibboleth: The Rule for Artificial Agents,” <https://hai.stanford.edu/news/shibboleth-rule-artificial-agents>, [Accessed: 20 April 2023].
- THE ECONOMIST (2023): “How technology is redrawing the boundaries of the firm,” Accessed: 2023-05-22.
- WALTZ, D., M. DIRAGO, AND J. RONALD I. RAETHER (2019): “Illinois Employers Must Comply with Artificial Intelligence Video Interview Act,” <https://www.shrm.org/resourcesandtools/legal-and-compliance/state-and-local-updates/pages/illinois-artificial-intelligence-video-interview-act.aspx>, [Accessed: 20 April 2023].
- ZHANG, D., N. MASLEJ, E. BRYNJOLFSSON, J. ETCHEMENDY, T. LYONS, J. MANYIKA, H. NGO, J. C. NIEBLES, M. SELLITTO, E. SAKHAEI, Y. SHOHAM, J. CLARK, AND R. PERRAULT (2022): “The AI Index 2022 Annual Report,” Tech. rep., AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University.

Figures and Tables

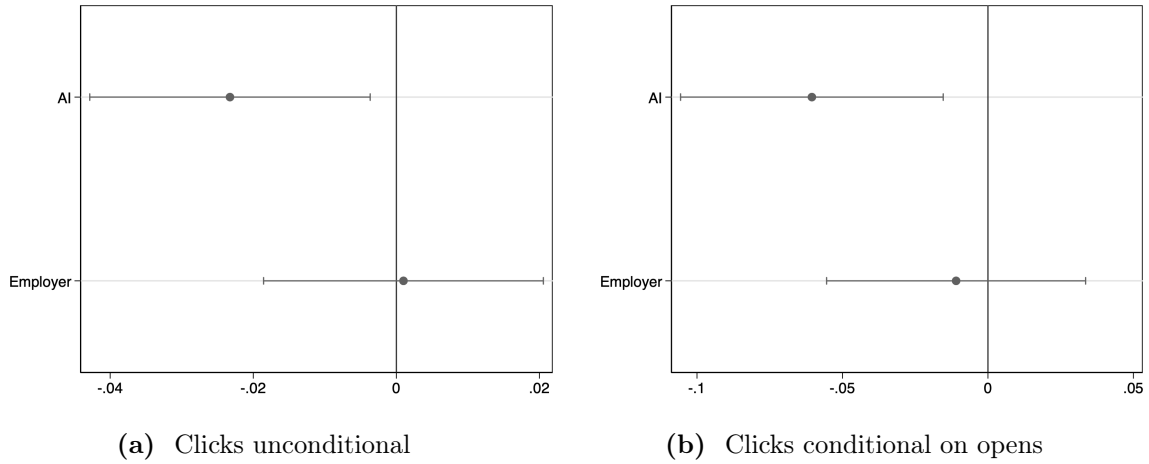


Figure 1: AI and Employer Effects on Clicks

Notes: The effect of Source: AI and Source: Employer treatments on a binary variable for Clicks. The omitted category is no information about the source of the recommendation. Panel (a) is the unconditional effect, and Panel (b) is the effect conditional on opens. 95 percent confidence intervals are displayed.

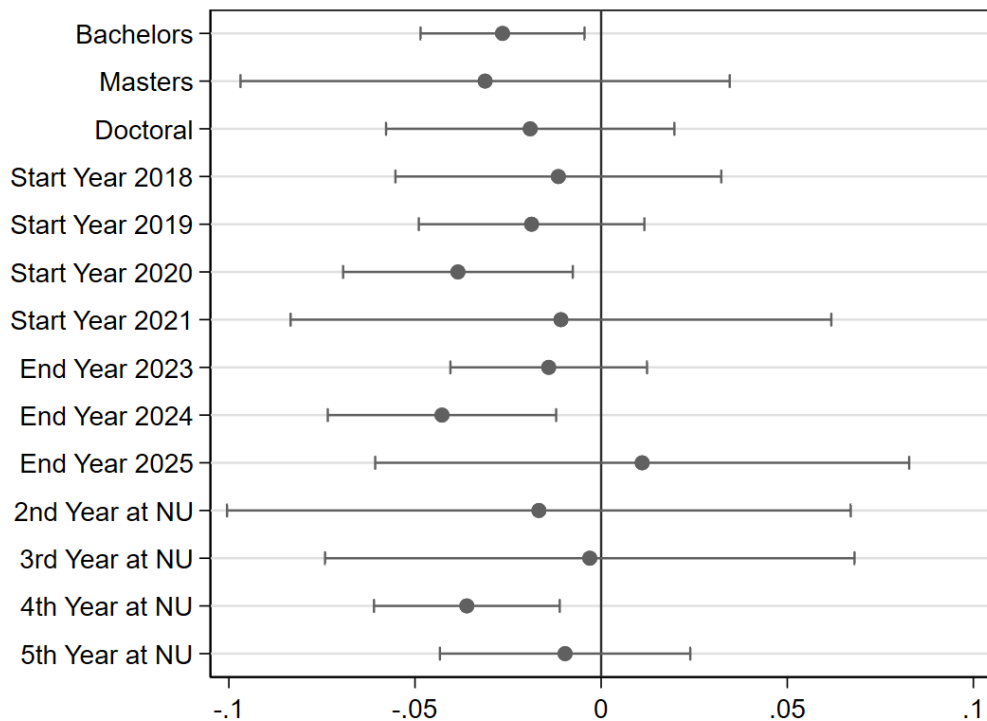


Figure 2: Heterogeneity by Student Characteristics

Notes: The Source: AI effect, sub-sampled by the student characteristics above. 95 percent confidence intervals are displayed.

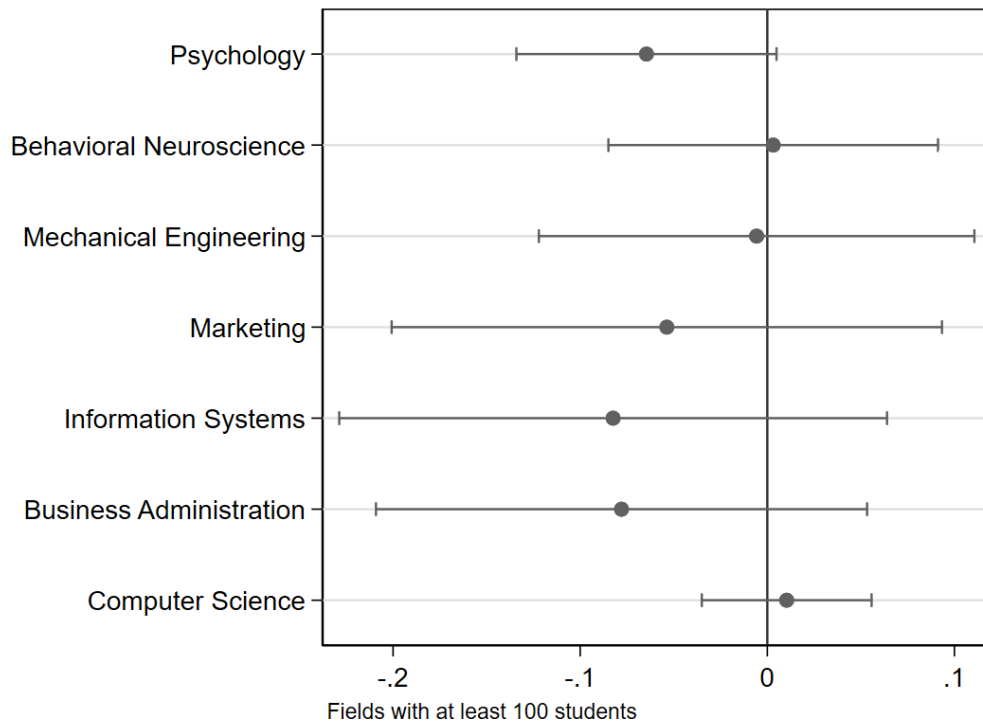


Figure 3: Heterogeneity by Field

Notes: The Source: AI effect, sub-sampled by field. Only fields with at least 100 students in the experiment. 95 percent confidence intervals are displayed.




Figure 4: Comparing High Quality Candidates to All Other Candidates

Notes: The Source: AI effect, sub-sampled by high quality candidates and other candidates. Fields with good matches for each job were identified by the employers, and create the sample of high quality candidates. 95 percent confidence intervals are displayed.

Table 1: Summary Statistics

	(1)
Source: Employer	0.338 (0.473)
Source: AI	0.333 (0.471)
No Wage Listed	0.225 (0.417)
Wage: \$16	0.0359 (0.186)
Wage: \$18	0.123 (0.328)
Wage: \$20	0.125 (0.331)
Wage: \$22	0.135 (0.342)
Wage: \$24	0.133 (0.340)
Wage: \$26	0.131 (0.338)
Wage: \$28	0.0922 (0.289)
Non-Pecuniary Benefit	0.128 (0.334)
Bachelors	0.731 (0.444)
Masters	0.139 (0.346)
Doctoral	0.0308 (0.173)
J.D.	0.00177 (0.0421)
Start Year	2019.5 (0.895)
End Year	2023.5 (0.619)
Year in Program	4.009 (0.903)
Click (Binary)	0.0820 (0.274)
Open (Binary)	0.410 (0.492)
<i>N</i>	4511

mean coefficients; sd in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: By Job

	(1) Full Sample	(2) Psychology	(3) Full Stack	(4) Crypto	(5) Video	(6) Marketing	(7) Strategy	(8) VisualDesign	(9) Kaggle
Source: Employer	0.000993 (0.0104)	-0.00135 (0.0166)	0.0300* (0.0177)	-0.0617 (0.0751)	0.0400 (0.0444)	-0.0178 (0.0461)	-0.0517 (0.0319)	-0.0137 (0.0326)	-0.0642 (0.116)
Source: AI	-0.0233** (0.00976)	-0.0138 (0.0155)	0.00375 (0.0168)	-0.0483 (0.0836)	-0.0573* (0.0303)	-0.0554 (0.0421)	-0.0652** (0.0307)	-0.0315 (0.0324)	-0.0818 (0.109)
Observations	4534	1140	1589	67	248	334	628	469	59

Notes: Regression estimates for the effect of Source: Employer and Source: AI on the binary variable Clicks. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: STEM

	(1) Original	(2) Matched CIP Sample	(3) Only STEM	(4) Non STEM	(5) Interacted
Source: Employer	0.000993 (0.00997)	0.00180 (0.0105)	0.0182 (0.0132)	-0.0277 (0.0175)	-0.0277 (0.0186)
Source: AI	-0.0233** (0.01000)	-0.0271** (0.0105)	-0.0153 (0.0133)	-0.0479*** (0.0174)	-0.0479*** (0.0175)
STEM					-0.0251 (0.0168)
STEM x Source: AI					0.0325 (0.0216)
STEM x Source: Employer					0.0459** (0.0231)
Observations	4534	4150	2649	1501	4150

Notes: Regression estimates for the effect of Source: Employer and Source: AI on the binary variable Clicks. Column (1) is the main effect. Column (2) is the sample for which a field of study described on LinkedIn can be matched to an instructional program at the university, described in detail in Section 5. Column (3) and (4) subsample STEM and Non STEM candidates, respectively. Column (5) interacts the main effect for STEM and non-STEM candidates. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$