AI-Enabled Job Markets & Market Participation:
Jobseekers' "Rational Expectations" About Competition vs "AI Aversion"

Sarah H. Bana, Chapman University
Kevin J. Boudreau, Northeastern University & NBER

Abstract:

As AI matching algorithms become increasingly important in modern job search, we examine how jobseekers respond to learning that job recommendations come from AI systems. While existing research has focused on improving algorithmic matching quality, far less attention has been paid to jobseeker responses. Through a field experiment with 4,562 jobseekers on a labor market platform, we find a series of patterns consistent with jobseekers forming expectations about how AI recommendations change expected payoffs. When AI involvement is disclosed, participation drops by 27.4 percent (2.4 percentage points) compared to when no information about the recommendation source is provided. While this could reflect preferences regarding AI (i.e., AI aversion), our evidence points to a more fundamental mechanism: jobseekers respond strategically to how AI changes the competitive environment. Specifically, jobseekers appear to update their beliefs about match quality, evaluation criteria, and the reach of recommendations when they learn AI is involved. These findings have important consequences for firms and platform designers, as they navigate inviting high quality candidates under different expectations of competition.

Keywords: Human-AI interaction, Workers, Employers, Hiring, Artificial Intelligence, Digital Platforms, Field Experiment

1 Introduction

Over the past decade, matching algorithms have become embedded within virtually every major job platform, silently transforming how employers and jobseekers connect. With millions of Americans relying on these platforms to discover job opportunities, and employers using them to identify potential candidates, there is immense potential to improve these algorithmic systems to better serve both sides of the labor market.

Existing research has largely focused on improving the efficacy of algorithms underlying these search platforms. Studies have developed their own recommender systems (Le Barbanchon, Hensvik, and Rathelot 2023), proposed incorporating additional attributes, integrating time-variant components (Kokkodis and Ipeirotis 2023), and measured the extent of algorithmic biases (Zhang and Kuhn 2024, Wall and Schellmann 2021). While technical improvements can and will continue with advances in large language models (LLMs), far less attention has been paid to how jobseekers respond to algorithmically generated recommendations.

The jobseeker's decision about whether to accept a job recommendation is a multidimensional one - in addition to assessing the match quality associated with the recommendation, they may consider the probability of getting this particular job and their relative comparative advantage at proceeding with one recommendation compared to others. Moreover, algorithmically generated recommendations may operate at different scales: targeting jobseekers like a needle in a haystack, or using the inherent scalability of such technology with greater reach.

In this paper, we measure how jobseekers' respond to algorithmically generated recommendations by conducting a field experiment on a labor market platform. 4,562 jobseekers received an email with a job recommendation, generated through a process that involves both a human and an algorithm. The key experimental variation is the different information about how the recommendation was generated. In the control group, jobseekers were provided no information about how the recommendation was generated. In the AI treatment group, the jobseekers were told that the recommendation was generated through an AI matching algorithm. Finally, in the employer treatment group, the jobseekers were told that the recommendation was generated by an employer.

We find that jobseekers are much less likely to participate and respond to the job recommendation if they are in the AI treatment - a 27.4 percent (2.4 percentage point) reduction in the probability of participating. Scholars have considered this type of response to algorithms "AI aversion." Described by Dietvorst, Simmons, and Masey (2015), this is the idea that people remain resistant to using algorithms, opting to use suggestions made by an inferior human rather than suggestions made by

a superior algorithm. In contexts ranging from investment decisions (Dietvorst, Simmons, and Massey 2015), to purchase decisions (Luo et al. 2019) to performance feedback (Tong et al. 2021), humans have reacted aversely to AI.

This idea that individuals hold inherent preferences against AI that may lead them to opt out of AI mediated labor markets has major implications for the efficacy of these recommender systems, especially when AI disclosure laws are being proposed and implemented around the country. However, the literature largely considers AI aversion to occur when an individual is making a decision that is independent of others' decisions. The job search context fundamentally differs from these individual settings, as candidates must evaluate recommendations within a competitive landscape. A jobseeker must determine whether to apply for a job based on the probability of getting the job, which can also be altered by their expectation of the quality of the recommendation, their expectation of inputs valued in the underlying process, and their expectation of the size of the candidate pool.

We find evidence that jobseekers in the AI treatment respond consistently with different underlying expectations about the probability of getting the job. First, the ex ante best matched jobseekers in the AI treatment appear to be less likely to apply compared to these same jobseekers in the control group, consistent with the expectation that the AI treatment produces lower quality predictions. Second, the jobseekers with more LinkedIn connections, a measure of investment in social capital that is potentially contrary to codified processes like algorithms, are less likely to apply in the AI treatment than in the control group. This is consistent with expectations that the AI treatment alters the evaluation process, changing the distribution of the types of candidates interested in applying. Finally, in the control group, there is a negative relationship between pool size and the probability of applying - jobseekers are more likely to apply for roles with fewer candidates. This relationship disappears in the AI treatment, such that jobseekers no longer factor in pool size when making application decisions. This is consistent with the AI treatment altering expectations of the number of jobseekers who receive the recommendation.

This evidence that jobseekers expect the structure of competition to change under AI matching has important policy implications, particularly given the recent growth of AI disclosure laws. These laws require firms to disclose when they are using AI tools in hiring processes. For example, Illinois' Artificial Intelligence Video Interview Act, passed in August 2019, requires employers to notify applicants before interviews if AI will be used in video analysis. More recently, NYC implemented 'The Artificial Intelligence in Hiring Law,' which mandates disclosure of automated employment decision tools (AEDTs), including background check tools, resume scanners, and candidate ranking programs. As

hiring is increasingly considered a high-risk AI activity, similar legislation is being proposed across multiple jurisdictions.

Our findings align with broader evidence that jobseekers strategically respond to information about competition in labor markets. Fradkin, Bhole and Horton (2021) demonstrate that on Jobs on Facebook, jobseekers actively avoid competition, sometimes opting for less popular positions. This suggests that platform designers, who control how competition-related information is presented, can significantly influence jobseeker behavior.

Importantly, our results indicate that firms' and platform designers' choices around AI disclosure can substantively affect how jobseekers respond to AI-generated recommendations. Unlike traditional AI aversion, which may or may not diminish as jobseekers gain familiarity with AI tools, our measured response stems from expectations about competition. Both Fradkin, Bhole, and Horton (2021) and Gee (2018) demonstrate that firms and platforms can effectively manage these expectations through strategic information provision.

We extend the literature on human-AI interactions (Tong et al. 2021; Luo et al. 2019), showing a large change in behavior in the novel context of job market matching while describing potential mechanisms consistent with this response. Our findings also contribute to understanding the digitization of labor market dynamics, showing how AI generated recommendations alter jobseeker expectations, potentially reshaping market outcomes.

The remainder of this paper is structured as follows - Section 2 describes the prior literature and a guiding characterization of a jobseeker's decision. Section 3 describes the experimental design, sample, and data collection methods. Section 4 presents the analysis and results. Section 5 summarizes and discusses results with implications for platform design and policy.

2 Related Literature and Hypothesis Development

2.1 Related Literature

One mechanism by which AI-enabled matching could influence jobseeker behavior relates to inherent preferences or attitudes toward interacting with AI. A number of papers identify AI aversion by comparing responses to the same tool when AI disclosure occurs at different times. These studies randomize whether participants learn about AI involvement at the start of their interaction or only after completing it. When disclosure occurs early, participants proceed through the entire interaction knowing they are engaging with an AI system, allowing researchers to measure how this knowledge affects their behavior. In contrast, late disclosure reveals preferences unconfounded by awareness of AI involvement.

These studies typically interpret negative responses to AI disclosure as reflecting inherent preferences against AI interaction, and occur largely outside of the labor market.

This literature has documented widespread evidence of AI aversion across various contexts. Davenport et al. (2020) demonstrate that despite the efficiencies AI can bring, users often express discomfort and reluctance when engaging with AI systems. For example, in e-commerce, Luo et al. (2019) found that disclosing the use of AI chatbots for customer support reduced purchase rates by approximately 80%. In decision-making contexts, Lee (2018) observed that experimental subjects viewed algorithmic decisions as less trustworthy and fair compared to human decisions. Dietvorst, Simmons, and Massey (2015) documented that people avoid using algorithms after seeing them make mistakes, even when algorithms are demonstrably more accurate overall. In human resources, Tong et al. (2021) found that employees may develop negative perceptions of AI-based performance evaluations once their AI origin is disclosed. In healthcare, Longoni, Bonezzi, and Morewedge (2019) demonstrated that consumers are less likely to utilize medical AI than human providers.

These findings suggest some degree of systemic reluctance to engage with AI across domains. The negative perceptions of AI, including subjective perceptions about its trustworthiness, fairness, and ability to understand human nuances, may deter some jobseekers from participating in AI-driven job matching processes—even in cases in which machine-based intelligence could be superior to human-based methods. On the other hand, Avery et al. (2024) finds that female applicants are more likely to complete an application assessed by an AI recruiter compared to being assessed by a human recruiter. So ex ante, the direction of the preference can be ambiguous. Importantly, this preference is often considered uncorrelated with other elements involved in the process.

More recently, researchers have begun to examine responses to algorithms beyond simple preference-based explanations, particularly in contexts involving strategic decisions. Albright (2024) provides particularly compelling evidence by studying a setting where judges first received only algorithmic predictions about defendant risk, and later received explicit recommendations about bail decisions. By studying the introduction of recommendations while predictions remained constant, Albright shows that judges grant more lenient bail when they can share responsibility for negative outcomes with an algorithmic recommendation system. This demonstrates how responses to algorithms can reflect strategic considerations around liability rather than just preferences.

In another context, Kim et al. (2024) shows that restaurant inspectors' apparent rejection of algorithmic recommendations stems from inspectors attempting to optimize for objectives not captured by the algorithm. In labor markets, Hoffman et al. (2018) examine how the introduction of job testing

affects hiring decisions when managers possess private information, finding that managers who hire against test recommendations consistently make worse hiring decisions.

These papers, while not eliminating the possibility of algorithmic aversion, highlight how responses to AI can arise from strategic considerations and institutional environments rather than purely from preferences. This distinction becomes especially relevant in contexts like job search, where decisions involve strategic interactions and competition.

In the next subsection, we develop a simple framework characterizing the jobseeker's decision process. This framework guides the formulation of testable hypotheses about how information environments and expectations of competition may change in AI-mediated labor markets.

2.2 Conceptual Framework and Hypothesis Development

We now present a conceptual framework for understanding how jobseekers evaluate labor market participation opportunities. This structure allows us to isolate specific channels through which AI-generated recommendations may influence expectations and behavior.

A jobseeker participates in a particular labor market if the expected value of participating exceeds zero. This expected value depends on several factors: the probability of receiving the role, the expected wage, application costs, any utility or disutility associated with AI-generated recommendations, and their outside option. For a risk-neutral jobseeker, we can express this as:

$$EV = p_{Hire} * E[Wage] - \underbrace{c}_{\text{Cost}} - \underbrace{B}_{\text{Preference re AI}} - \underbrace{O}_{\text{Outside Option}} \geq 0$$

While previous literature has focused heavily on the preference element, B, (i.e., AI aversion), we begin by emphasizing the competitive components of this decision, particularly the probability of being hired (p_{Hire}).

First, the probability of being hired is inherently a contest. With one open role, p_{Hire} depends on the jobseeker's perception of their true rank within the applicant pool, how accurately this rank will be assessed during the hiring process, and the number of other jobseekers in the pool. Importantly, AI disclosure may affect each of these elements: First, a jobseeker might believe that AI recommendations are more or less accurate in assessing their true qualifications, potentially changing their perceived relative standing. Second, the jobseeker might anticipate that the evaluation process will differ when AI is involved, placing different weights on various aspects of their qualifications. Third, the jobseeker might expect that AI recommendations reach a different number of potential candidates, altering the expected competition.

This characterization, though not comprehensive, yields testable hypotheses about the effect of AI disclosure on jobseeker participation. Our framework highlights how information about AI involvement might alter strategic considerations related to competition, while also acknowledging the preference-based explanations that have dominated the literature. Together, these perspectives suggest four distinct hypotheses about why jobseekers might respond differently to AI-generated recommendations: three based on how AI disclosure might change expectations about match quality, evaluation processes, and the competitive landscape, and one based on inherent preferences against AI systems. The hypotheses are as follows:

H1: Jobseekers expect AI recommendations to be less accurate in identifying suitable matches compared to non-AI processes.

Jobseekers may believe that AI systems lack the nuanced understanding of human potential and qualifications that human recruiters possess. They might expect AI to rely on simpler pattern-matching rather than holistic evaluation, potentially missing important context or unique attributes that don't fit standardized criteria.

H2: Jobseekers expect hiring processes involving AI recommendations to place greater emphasis on codified inputs rather than tacit qualifications.

Jobseekers might anticipate that AI systems primarily evaluate quantifiable credentials (degrees, years of experience, specific skills) rather than more noncognitive traits like fit, potential, or emotional intelligence. This expectation could lead candidates who have invested in social capital or other difficult-to-codify attributes to perceive lower returns to applying.

H3: Jobseekers expect AI recommendation systems to reach a broader pool of potential candidates than non-AI processes.

Algorithms may serve multiple objectives. Some may focus on identifying rare, exceptional candidates within large populations, while others may prioritize broad reach across potential applicants. The inherent scalability of algorithmic systems may lead jobseekers to expect that AI recommendations are distributed more widely, potentially reaching a substantially larger candidate pool than traditional

human-driven processes. This expectation of increased competition could significantly influence their assessment of the probability of being hired.

Additionally, we add one hypothesis on the preference related explanations.

H4: Jobseekers' reduced participation in AI-mediated job recommendations reflects inherent preferences against interacting with artificial intelligence systems in hiring.

Existing research demonstrates that individuals frequently exhibit aversion to AI systems in decision contexts. This phenomenon could explain reduced engagement with AI-generated job recommendations through fundamental preferences against algorithmic involvement in career decisions, independent of any strategic considerations about competition or evaluation processes. These preferences might stem from concerns about algorithmic bias, privacy implications, or general technological anxiety. This preference-based explanation might suggest participation differences that align with general comfort with technology.

3 Experimental Design and Data

To measure the effect of AI-enabled matching on market participation, we design and conduct a field experiment in collaboration with a university job platform specializing in short-term, part-time skilled work. Our objective is to isolate the behavioral response to an AI recommendation, separate from any effects related to the quality of the algorithm. The central challenge of our design, therefore, is to maintain the same recommendation quality across jobseekers while systematically varying their understanding of how these recommendations were generated. Additionally, we must adhere to the principle of non-deception, as participants are real jobseekers applying for actual roles in a functioning labor market.

We work with the platform to consider a set of jobs launched over a three week period. The roles include jobs related to cryptocurrency application, full stack engineering, front-end design, data architecture, psychology, marketing strategy, video design, and visual design. Table 1 contains a brief description of the postings.

We work with the employer to identify a set of keywords associated with the ideal applicant - for example, for the full stack engineering role, we included keywords such as "Computer Science,"

"Computer Engineering" that may be on the jobseeker's LinkedIn profile as their major or related experience. We match these keywords with the universe of LinkedIn profiles for this university. A total of 4,562 jobseeker LinkedIn profiles match with the keywords identified by the employer for the various roles.

Each jobseeker that is a match is sent an email with the specific role they are matched to. The recommendation email contains some simple information about the job - a title, a one sentence description, an expected wage, and a link to click through and learn more about the role and/or apply.

The jobseekers are randomized into one of three conditions. The control condition provides no information about the recommendation source. The AI matching treatment includes a line that says "Recommendation source: A.I. matching algorithm using publicly available LinkedIn profiles." The Human matching treatment includes a line that says "Recommendation source: The employer generated a list of candidates by reviewing publicly available LinkedIn profiles."

It should be noted that we made a deliberate effort to design this research and label the machine intelligence recommendation sources to include both "AI" and "algorithm." At this point in history, subjective definitions of machine-based decision-making and artificial intelligence (AI) often vary. Some people use the term "artificial intelligence" broadly to refer to all machine-based decision rules, however simple or complex; while others reserve it for specific subsets of machine learning methods.[1] We both leverage and allow for this current vagueness to refer to the involvement of machines in our simple matching procedure as related to AI, machine intelligence, or algorithmic decision-making.

As mentioned above, the central challenge of our design is to maintain recommendation quality across jobseekers, while also varying disclosure. Importantly, this design meets that challenge as both an algorithm (keyword matching) and an employer (by selecting the keywords) generated the recommendation, making both statements true. Randomization varies only disclosure, and therefore what the jobseekers know about the process.

[1] Princeton computer scientists Narayanan and Kapoor (2024) describe this situation in the first pages of their book: _"Imagine an alternate universe in which people don't have words for different forms of transportation—only the collective noun "vehicle." They use that word to refer to cars, buses, bikes, spacecraft, and all other ways of getting from place A to place B. Conversations in this world are confusing. There are furious debates about whether… vehicles are environmentally friendly, even though… one side of the debate is talking about bikes and the other side is talking about trucks. There is a breakthrough in rocketry, but the media focuses on how vehicles have gotten faster... Now replace the word "vehicle" with "artificial intelligence," and we have a pretty good description of the world we live in."

The balance test across the three conditions on relevant covariates is shown in Table XX.

We also implement a second treatment, varying the expected wage associated with the role. In our context, we are specifically interested in understanding the relationship between the recommendation source and the jobseekers' willingness to participate. The expected wage should also, independently, affect the jobseekers' willingness to participate. Comparing these two objects will give us a measure of the size of the recommendation source effect.

Within this wage treatment, the conditions vary the expected wage or list no wage associated with the role. We design an approach to truthfully vary the expected wage. At this institution, the hourly wage is a band from 16 to 28 dollars, and jobseekers will receive a specific wage depending on their level of education and experience. We present expected wages in $2 increments ($16, $18, $20, $22, $24, $26, or $28), all of which fall within the actual wage band. By varying the expected wage that is included on the initial recommendation email, we have varied the participants' expectation of wage, without changing the true wage associated with the role (which is derived from a formula).

The experiment was conducted in Fall 2022. We use the platform's customer relationship management (CRM) system to collect data from the experiment. Our primary outcome variable is jobseeker participation in response to the recommendation, measured by whether they clicked on the link in the email to access the platform. This measure captures the jobseeker's initial engagement and willingness to collect additional information about the recommended opportunity. We define Market Participation as a binary variable that equals one if the jobseeker clicked on the recommendation link, which occurred for 8.7 percent of the full sample of jobseekers in the experiment.

In addition to our primary outcome, we also track several secondary engagement metrics. The platform records a measure of whether the message was opened, although this is dependent on the users' email client settings and so not perfectly accurate, as well as the number of clicks made to the website. The system received a positive signal of opening from 41.8 percent of the participants, which is relatively high compared to typical email marketing open rates, suggesting our sample consists of engaged jobseekers. The results remain both quantitatively and qualitatively similar when conditioning on a positive indication of email opening. However, we do not condition our main results, as it is possible that a jobseeker looked at the email without receiving an open signal due to technical limitations in email tracking mechanisms or because users can view and click links in email previews without formally opening the message. We discuss the results conditional on opening in Section 5.

We have a number of covariates from the LinkedIn profile - the degree they are pursuing at the institution, the field of this degree, and the projected year of graduation. We combine this with information on STEM fields from the university and the United States Citizenship and Immigration Services to characterize whether jobseekers have STEM backgrounds. The match rate between fields listed on LinkedIn and programs at the institution is 90 percent. We also collect the number of LinkedIn connections, measured at the beginning of the experiment (Fall 2022).

Our jobseekers are a mix of bachelors, masters, doctoral, and professional degree students, with a majority (63 percent) seeking STEM degrees. The mean professional network size is 294 connections. Additional descriptive statistics are provided in Table 1.

4 Empirical Strategy

4.1 Main Effect: AI-Matching Treatment Effect on Jobseeker Participation

How does the AI matching treatment affect the probability that a jobseeker participates in the labor market? Because participants have been randomized into treatment, we can estimate the effect through a simple regression.

We estimate the equation

$$MarketParticipation_{ij} = \beta_0 + \beta_{AI} AITreatment_i + \beta_{Human} HumanTreatment_i + \varepsilon_{ij}$$

In this equation and those that follow, i indexes individual jobseekers and j indexes jobs, as we observe whether each jobseeker participates in the market for each potential job match. In this experiment, to minimize any learning effects, we only set one recommendation to each jobseeker. Our outcome variable, $MarketParticipation_{ij}$, is an indicator equal to 1 if the jobseeker demonstrates any interest or seeks more information by clicking through to the platform. β_0 represents the average participation rate for jobseekers in the control group. β_{AI} represents the effect of the AI treatment compared to the control group. β_{Human} represents the effect of the Human treatment compared to the control group. We use robust standard errors.

We estimate three additional regressions, progressively adding in the wage treatments, demographic characteristics, and then job fixed effects.

$$\begin{aligned} MarketParticipation_{ij} &= \beta_0 + \beta_{AI} AITreatment_i + \beta_{Human} HumanTreatment_i + \beta_{Wage} Wage_i \\ &+ \beta_{NoWage} NoWage_i + \varepsilon_{ij} \end{aligned}$$

$$MarketParticipation_{ij} = \beta_0 + \beta_{AI} AITreatment_i + \beta_{Human} HumanTreatment_i + \beta_{Wage} Wage_i + \beta_{NoWage} NoWage_i + \delta X_i + \varepsilon_{ij}$$

$$MarketParticipation_{ij} = \beta_0 + \beta_{AI} AITreatment_i + \beta_{Human} HumanTreatment_i + \beta_{Wage} Wage_i + \beta_{NoWage} NoWage_i + \lambda_j + \varepsilon_{ij}$$

X_i, a vector of demographic characteristics, are derived from the user's LinkedIn profiles and include the degree of the jobseeker, the expected year of graduation, and whether they are in a STEM degree program.

All regressions are estimated as linear probability models using robust standard errors. These provide the most straightforward interpretations of the main effects. Probit models are provided in the Appendix and are consistent with the linear probability models.

4.2 Testing Mechanisms Behind AI Treatment Effects

We examine four potential mechanisms driving the effect of AI treatment on market participation, corresponding to our four hypotheses. Three mechanisms relate to how AI disclosure alters jobseekers' perceived probability of being hired (p_{Hire}), while the fourth concerns inherent preferences about AI (reflected in the B term of our conceptual framework). For the first three hypotheses, we estimate separate regressions across our three experimental conditions (Control Group, Human Matching, and AI Matching) with key covariates, comparing how their relationships with market participation differ across conditions. For the fourth hypothesis regarding preferences, we instead compare treatment effects across different subgroups defined by technical background.

First, we operationalize the hypothesis that jobseekers associate AI matching with higher or lower quality compared to the control group. We presume that the expected prediction quality may be different, but not the direction. We test this hypothesis by considering an ex-ante well matched group. If AI leads to relatively low-quality matches, this reduces the incentives for the well-matched. For example, the very top candidate can only be ranked lower with additional noise. On the other hand, it may increase in the incentive for the less well-matched.

On the other hand, if AI leads to relatively high quality matches, then this would yield the opposite prediction. The well-matched individuals would have a greater incentive to apply, and those less-well matched might have less of an incentive to apply.

In our context, for technical jobs, the STEM masters' level candidates in relevant subfields are ex ante well matched candidates. We work with the employer in each of the roles to identify fields that would be well-matched in our context. Then, we generate an indicator variable for being in the related field, and a second indicator variable for being a STEM candidate in the related field. We then estimate the regression separately on the three different conditions - Control Group, Human Matching, and AI Matching to compare how these different groups react to the recommendation source. This analysis is only conducted for STEM jobseekers as this ex-ante well matched group is only identified in this case.

Next, we test hypothesis H2: that AI disclosure signals different evaluation criteria, particularly a shift toward codified inputs over tacit qualifications. In our conceptual framework, this hypothesis relates directly to how jobseekers form expectations about p_{Hire} - their probability of being hired. Specifically, jobseekers with strong professional networks may believe these networks confer advantages in human-led evaluation processes that are diminished in AI-driven processes. Like degree program, these are an ex-ante measure of investment in human-focused processes.

To operationalize this hypothesis, we use the LinkedIn connection counts as our measure of professional network investment. Like degree program, these are an ex-ante measure of investment in human focused processes. We then examine how network size relates to market participation differently across our three experimental conditions by estimating:

$$MarketParticipation_{ij} = \beta_0 + \beta_{Network} NetworkSize_i + \lambda_j + \varepsilon_{ij}$$

We estimate this equation separately for each of our three conditions: Control Group, Human Matching, and AI Matching. If hypothesis H2 is correct, we expect $\beta_{Network}$ to be significantly smaller (or even negative) in the AI Matching condition compared to the Control Group condition. This would indicate that the relationship between professional network investment and participation weakens when AI is involved in the recommendation process, suggesting jobseekers perceive different evaluation criteria when AI is disclosed.

Next, we test hypothesis H3: that receiving an AI recommendation alters the expected size of the candidate pool, directly impacting the perceived probability of hire p_{Hire}). Jobseekers may form different expectations about market competition when AI is involved - either believing AI is highly scalable and thus recommends opportunities to more candidates, or that AI excels at finding specific matches ("needle in a haystack") and thus recommends to fewer candidates.

To test this empirically, we use the actual number of competitors within the "local market" as an approximation of potential competition. We define this local market as all on-campus individuals with LinkedIn keyword matches to the job description. We then examine how local market size relates to market participation across our three experimental conditions by estimating:

$$MarketParticipation_{ij} = \beta_0 + \beta_{Competition} \ln(LocalCompetitors_j) + \lambda_j + \varepsilon_{ij}$$

We estimate this equation separately for each condition: Control Group, Human Matching, and AI Matching. If hypothesis H3 is correct, we expect the coefficient $\beta_{Competition}$ to differ significantly across conditions. In standard conditions (Control and Human Matching), we expect $\beta_{Competition}$ to be negative, as more competition should reduce the perceived probability of hire. However, if AI disclosure alters expectations about competition, this relationship may be attenuated or even reversed in the AI Matching condition, suggesting that jobseekers form different beliefs about competitive dynamics when AI is involved in the matching process.

Finally, we test hypothesis H4: that jobseekers have greater disutility from receiving a recommendation from an algorithm, reflected in the preference term B in our conceptual framework. The information systems literature suggests that individuals with technical backgrounds often have different attitudes toward technology (Agarwal and Prasad 1998; Venkatesh et al. 2003). To test this preference-based explanation, we examine whether STEM and non-STEM applicants respond differently to the AI treatment.

We classify degree programs using the LinkedIn profile field matched to the university's certified instructional program (CIP) code, following the United States Citizenship and Immigration Services categorization of STEM programs. For approximately 10 percent of jobseekers, there is insufficient matching between the LinkedIn field and programs offered at the university; these individuals are excluded from this analysis. We then estimate our main specification separately for STEM and non-STEM subsamples to determine if technical background moderates the effect of the AI recommendation source, which would support preference-based explanations for AI aversion.

5 Results

How does the AI matching treatment affect the probability that a jobseeker participates in the labor market? Table 2 compares the effect of the AI matching treatment and the Human matching treatment to the control group. Column (1) shows that the AI matching treatment decreases the probability that a jobseeker participates by 2.5 percentage points. This represents a 26.3 percent decrease,

as the average market participation rate in the control group is 9.5 percent. The Human matching treatment, on the other hand, shows no statistically significant effect. With a coefficient quite close to 0, this can be interpreted as the jobseekers perceiving no additional information from the human treatment compared to no information about the recommendation source.

In Column (2), we add the orthogonal set of wage treatments. The AI matching treatment coefficient stays significant and virtually identical, with a coefficient of 2.4 percentage points. The coefficient on wage is positive and statistically significant, which is reassuring, as higher expected wages are associated with greater market participation rates. The coefficient on not including wage is statistically significant and positive. This may suggest either that the expected wage is low, or that providing the wage anchors jobseekers to compensation at the expense of other valuable elements of the opportunity.

The coefficient on wage allows us to compare the effect of the AI recommendation source to the wage effect. The AI treatment effect can be quantified as 2.6 (10* 2.4/0.091) times 10 dollar wage reduction. This substantive wage-equivalent effect highlights that jobseekers are responding meaningfully to information about AI involvement in the recommendation process. It also emphasizes the tradeoff that the employer might be making, from a candidate pool perspective, of disclosing the AI recommendation source.

Column (3) adds jobseeker characteristics of degree and graduation year. The AI Matching coefficient remains stable at 2.4 percentage points. Finally, in Column (4), we add job fixed effects, allowing the participation probability to differ across jobs. The AI Matching coefficient stays at 2.4 percentage points, suggesting that this effect is not driven by heterogeneity across job types or unobserved job characteristics. The consistency of the AI coefficient across all four specifications emphasizes the robustness of our finding that disclosing AI involvement in recommendation significantly reduces market participation. To understand the mechanisms that drive this effect, we now turn to testing our four hypotheses regarding how and why jobseekers respond differently to AI-generated recommendations.

Column (1) of Table 4 regresses Market Participation on an indicator for _Related Fields_ and for those pursuing master's Degrees within these fields, along with a constant. This model is estimated on STEM candidates, so the constant term can be interpreted as representing STEM candidates from fields less directly relevant to the job at hand. The coefficient on _Masters_ within Related Field is positive, statistically significant, and substantial at 0.07 (s.e. = 0.017), or 7 percentage points. This result supports the idea that well-matched candidates are more likely to participate in the market.

The negative coefficient on Related Fields (outside of those pursuing master's degrees) may seem counterintuitive. However, this negative effect is relative to candidates outside of directly related fields, who were selected into the study sample based on keyword matching.

Table 3 shows the coefficients for the STEM candidates, for whom we can derive the closeness of match measure. Columns (4), (5), and (6) split out the regression for the control group (Column 4), the Human matching group (Column 5), and the AI matching group (Column 6). In the control group, the well-matched group is much more likely to apply. This is also true in the Human matching group. The AI matching treatment, however, has a statistically insignificant effect. The well-matched group shows lower relative incentives to participate. Comparing the coefficients in the control group and the AI matching group, we find that these coefficients are statistically different from one another. This is consistent with the expectation that AI matching results in lower-quality matches, changing the incentives of the top ranked candidates.

Table 4 repeats the subsampling exercise using the relevant covariate for ex ante investment in human processes - professional network size.

Table 4 repeats the subsampling exercise using the relevant covariate for ex ante investment in human processes - professional network size. The first three columns subsample on the control group, the next three subsample on the human matching group, and the final three subsample on the AI matching group. Column (1) shows a positive relationship between Market Participation and the number of LinkedIn connections within the control group. (Similar results are found when estimating the model on the entire study population.) However, the size of a professional network—like any personal characteristic—could be correlated with fields or other individual attributes. Therefore, column (2) adds fixed effects for each job posting to account for any potential spurious correlation between the nature of the job and the size of professional networks. Adding these fixed effects has no impact on the estimated relationship with professional network size. Column (3) further adds controls for 50 fine-grained subfields within STEM and includes a Female indicator to account for any demographic effects. Again, we see no impact on the estimated coefficient for professional network size.

Re-estimating these models for those in the Human-Matching treatment (columns 4 to 6) shows a similar positive relationship with the size of professional networks. However, in the AI-Matching treatment (columns 7 to 9), we find no such relationship between network size and Market Participation. A statistical test confirms that the coefficients are statistically different from one another. This means that jobseekers who traditionally were more invested in their professional networks were differentially less likely to participate under the AI-Matching treatment.

This evidence strongly supports hypothesis H2: that AI disclosure signals different evaluation criteria. Specifically, jobseekers with larger professional networks appear to believe that their investment in social capital will yield lower returns (in terms of p_{Hire}) when recommendations come from AI systems compared to human recruiters. This suggests that the negative AI effect is at least partially driven by strategic considerations about how different recommendation sources might value different types of qualifications, rather than simply reflecting a preference against AI.

We next examine hypothesis H3: that receiving an AI recommendation alters jobseekers' expectations about the size of the candidate pool. Specifically, we test whether AI disclosure creates uncertainty about how recommendations are distributed, potentially disrupting the normal relationship between local market competition and participation decisions. This uncertainty could arise because jobseekers lack clear mental models for how AI systems determine which candidates receive recommendations, unlike more transparent human-driven processes.

Table 5 Column (1) regresses Market Participation on the estimated number of competitors in the local market for each position. This is based on the total number of individuals in the defined market, on campus, whose profiles are keyword matched with the position (Section 3), No. Competitors. We find a negative coefficient of -0.017 (s.e. = 0.006) on the log of No. Competitors, consistent with individuals tending to have lower incentives to participate when facing many competitors.

The result is consistent with a negative crowding effect, but we should be cautious in interpreting this result, as variation in numbers of competitors exploits differences of No. Competitors across job postings. Job posting differences may differ in ways that could be correlated with the number of competitors for a given position in the local market. Further, we cannot add job posting fixed effects, as job fixed effects are clearly collinear with the number of competitors.[2] We can, however, control for jobseeker characteristics, which will be highly related to the nature of the job. This includes such characteristics as education level, cohort, and whether the candidate is from a STEM field. Adding these controls, as in model (2), does not perturb the estimated relationship between Market Participation and No. Competitors. Neither does adding other market-related variables such as Wage in model (3) influence the estimated relationship. Therefore, the negative relationship between Market Participation and No. Competitors at least be somewhat robust.

[2] We can, however, control for jobseeker characteristics, which will be highly related to the nature of the job. This includes such characteristics as education level, cohort, and whether the candidate is from a STEM field. Adding these controls, as in model (2), does not perturb the estimated relationship between Market Participation and No. Competitors.

Stratifying estimates across treatments, reveals significant differences in this link across treatments. Columns (4) to (6) re-estimate the model separately for the control group, the Human-Matching treatment, and the AI-Matching treatment. For the AI-Matching treatment, the relationship between Market Participation and No. Competitors is insignificant, suggesting that jobseekers in the AI-Matching treatment do not perceive the size of the local market as directly linked to the competition they will face. Additionally, it is statistically different from the coefficient on No. Competitor in the Control group. This result is consistent with AI-Matching jobseekers expecting that the number of competitors they face is unrelated to the size of the local market.

Finally, we test hypothesis H4: that jobseekers experience greater disutility when receiving recommendations from AI systems, reflecting inherent preferences above and beyond strategic considerations. The technology acceptance literature suggests that individuals with technical backgrounds typically have more positive attitudes and greater trust in new technologies (Agarwal and Prasad 1998; Venkatesh et al. 2003). If the AI effect is primarily driven by preference-based aversion rather than strategic beliefs about hiring probabilities, we would expect to see stronger negative responses among non-technical jobseekers.

Table 6 presents our analysis comparing STEM and non-STEM candidates. Models (1) through (3) estimate treatment effects for non-STEM fields, while models (4) through (6) do so for STEM fields. Model (1) reveals a large negative coefficient on AI Matching for non-STEM jobseekers of -0.035 (s.e. = 0.18). This 3.5 percentage point reduction represents a substantial 35% drop in participation relative to the 10 percent baseline participation in the control group. Notably, the Human-Matching treatment effect remains statistically indistinguishable from zero. Column (2) adds individual covariates for education level and graduation cohort, showing that market participation in non-STEM fields is more responsive to wage and less related to graduation year or education level. Column (3) adds job fixed effects with minimal impact on the estimates.

In contrast, columns (4) through (6) show that STEM candidates exhibit a small and statistically insignificant response to AI disclosure. This stark difference between STEM and non-STEM jobseekers provides compelling evidence for hypothesis H4, suggesting that preference-based explanations play a significant role in the overall AI effect. The pattern aligns with prior literature on technology acceptance, indicating that technical backgrounds may mitigate negative responses to algorithmic involvement in hiring processes.

6 Discussion and Conclusion

How does the disclosure of AI-based job matching affect jobseekers' willingness to participate in the labor market? This question is of paramount importance as AI systems increasingly permeate various aspects of recruitment and hiring processes, potentially reshaping the dynamics of job search and labor market participation (Agarwal and Prasad 1998; Venkatesh et al. 2003).

Despite the rapid adoption of AI in labor markets, little was previously known about its behavioral impacts on jobseekers. Our study contributes to filling this gap by describing two different ways in which AI-generated recommendations may affect jobseekers' market participation: an approach built on strategic considerations about hiring probabilities and an approach based on inherent preferences regarding technology.

Our findings reveal that disclosing AI involvement in job matching significantly reduces market participation, with differential effects across jobseeker populations. The AI-Matching treatment led to a 26% reduction in participation compared to both human matching and control groups. Importantly, our results provide compelling evidence that both proposed mechanisms contribute to this effect.

On one hand, jobseekers appear to form strategic beliefs about how AI recommendations affect their probability of being hired, particularly regarding how AI systems evaluate qualifications and distribute opportunities. On the other hand, we observe substantial preference-based responses, with technically-oriented candidates showing markedly less negative reactions to AI disclosure than their non-technical counterparts. Especially concerning, we found that ex-ante well-matched candidates were less likely to participate when informed about AI involvement, suggesting that disclosure may deter the very candidates employers most wish to attract, potentially exacerbating skills mismatch issues in labor markets .

Our conceptual framework, combined with three mechanism tests, highlight how AI recommendations could change the jobseekers' beliefs about the probability of hire, beyond simple AI aversion. First, more qualified candidates being less likely to apply appears consistent with the candidates' belief that AI recommendations are of lower quality. High-ranked candidates are dissuaded. Second, we find that candidates with stronger networks opt out, consistent with signals that AI evaluation processes might shift emphasis from relationship-based qualifications toward more codified attributes, leading those with investments in professional networks to perceive their comparative advantage in human-led processes. Third, we find that application rates are a function of the candidate pool size in the control group and the human matching treatment, but are unrelated related in the AI matching condition. This

suggests that jobseekers have uncertainty about how widely AI distributes recommendations, disconnecting local market competition from their participation decisions.

This work has substantive implications for firms and platforms implementing AI in their hiring processes. Despite the decreasing cost of generating recommendations, jobseekers demonstrate considerable sophistication in interpreting signals about evaluation processes. Beyond avoiding competition (Fradkin, Bhole, and Horton 2021), jobseekers form strategic beliefs about their comparative advantages in different selection mechanisms, which significantly influence their participation decisions.

Firms and platforms can respond by strategically addressing these concerns: providing transparency about algorithm quality, clarifying what inputs drive recommendations, and potentially signaling match strength, similar to how displaying applicant counts has proven effective (Gee 2016). Some platforms, such as ZipRecruiter, have already begun providing explicit signals about match quality. While algorithm interpretability is one possible approach, and has been successful in other contexts (Lebovitz et al. 2022), further research remains in the value of providing these types of explanations and the behavioral responses to such explanations in job recommendations.

Our study is not without limitations. First, it represents a single snapshot in time, not allowing us to examine how jobseekers' responses to AI disclosure might evolve with repeated exposure or as they update beliefs about algorithmic quality. Second, the research was conducted prior to the widespread adoption of generative AI tools like ChatGPT, which may have shifted public perceptions of AI capabilities. This timing difference could mean our findings underestimate acceptance in the current environment, though it might also have exacerbated differences between technically-oriented users and others. Finally, we intentionally framed the AI system generically, leaving open questions about how specific AI system descriptions, capabilities, or implementation details might moderate the effects we observed. Future research examining these temporal dynamics and contextual factors would provide valuable extensions to our work.

Given the crucial role of labor market matching in overall economic efficiency, taking greatest advantage of technical and information-processing potential of AI should also be a concern for policymakers and public authorities. Whether, how, and to what extent AI- and digitally-enabled labor markets could result in inefficiencies and unintended consequences are questions deserving of further research, building on existing work on the economics of digitization (Goldfarb and Tucker 2019), and the impact of AI on labor markets (Brynjolfsson, Mitchell, and Rock 2018; Acemoglu and Restrepo 2022). To mitigate the negative effects of AI aversion and heightened competition concerns, for

example, policymakers could promote transparency in AI-driven hiring by mandating the disclosure of AI's role in job matching processes, as exemplified by initiatives like the AI Disclosure Act of 2023. Public education campaigns could also be launched to demystify AI and build trust among jobseekers, where and if appropriate.

Our study is not without limitations. First, it represents a single snapshot in time, not allowing us to examine how jobseekers' responses to AI disclosure might evolve with repeated exposure or as they update beliefs about algorithmic quality. Second, the research was conducted prior to the widespread adoption of generative AI tools like ChatGPT, which may have shifted public perceptions of AI capabilities. This timing difference could mean our findings underestimate acceptance in the current environment, though it might also have exacerbated differences between technically-oriented users and others. Third, while we observe behavior consistent with strategic responses to AI disclosure, we were unable to directly measure jobseekers' beliefs about AI systems, requiring us to infer perceptions from behavioral responses rather than capturing explicit reasoning. Finally, we intentionally framed the AI system generically, leaving open questions about how specific AI system descriptions, capabilities, or implementation details might moderate the effects we observed. Future research examining these temporal dynamics, specific AI technologies (such as different types of language models), and directly measuring jobseekers' beliefs about AI evaluation processes would provide valuable extensions to our work.

References

Acemoglu, Daron, and Pascual Restrepo. 2022. "Tasks, Automation, and the Rise in U.S. Wage Inequality." *Econometrica* 90 (5): 1973–2016. https://doi.org/10.3982/ECTA19815.

Agarwal, Ritu, and Jayesh Prasad. 1998. "A Conceptual and Operational Definition of Personal Innovativeness in the Domain of Information Technology." *Information Systems Research* 9 (2): 204–15. https://doi.org/10.1287/isre.9.2.204.

Alex Albright. 2024. "The Hidden Effects of Algorithmic Recommendations." Institute Working Paper 104. https://www.minneapolisfed.org/research/institute-working-papers/the-hidden-effects-of-algorithmic-recommendations.

Autor, David H. 2001. "Wiring the Labor Market." *Journal of Economic Perspectives* 15 (1): 25–40. https://doi.org/10.1257/jep.15.1.25.

Avery, Mallory, Andreas Leibbrandt, and Joseph Vecci. 2024. "Does Artificial Intelligence Help or Hurt Gender Diversity? Evidence from Two Field Experiments on Recruitment in Tech." *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4764343.

Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. First Edition. New York London: W. W. Norton & Company.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. "What Can Machines Learn and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings* 108 (May):43–47. https://doi.org/10.1257/pandp.20181019.

Davenport, Thomas, Abhijit Guha, Dhruv Grewal, and Timna Bressgott. 2020. "How Artificial Intelligence Will Change the Future of Marketing." *Journal of the Academy of Marketing Science* 48 (1): 24–42. https://doi.org/10.1007/s11747-019-00696-0.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey. 2015. "Algorithm Aversion: People Erroneously Avoid Algorithms after Seeing Them Err." *Journal of Experimental Psychology: General* 144 (1): 114–26. https://doi.org/10.1037/xge0000033.

Fradkin, Andrey, Monica Bhole, and John Horton. 2021. "Competition Avoidance vs Herding in Job Search: Evidence from Large-Scale Field Experiments on an Online Job Board." https://andreyfradkin.com/assets/jobs_authors.pdf.

Gee, Laura K. 2018. "The More You Know: Information Effects on Job Application Rates in a Large Field Experiment." *Management Science*, April, mnsc.2017.2994. https://doi.org/10.1287/mnsc.2017.2994.

Goldfarb, Avi, and Catherine Tucker. 2019. "Digital Economics." *Journal of Economic Literature* 57 (1): 3–43. https://doi.org/10.1257/jel.20171452.

Hoffman, Mitchell, Lisa B Kahn, and Danielle Li. 2018. "Discretion in Hiring*." *The Quarterly Journal of Economics* 133 (2): 765–800. https://doi.org/10.1093/qje/qjx042.

Kim, Hyunjin, Edward L. Glaeser, Andrew Hillis, Scott Duke Kominers, and Michael Luca. 2024. "Decision Authority and the Returns to Algorithms." *Strategic Management Journal* 45 (4): 619–48. https://doi.org/10.1002/smj.3569.

Lee, Min Kyung. 2018. "Understanding Perception of Algorithmic Decisions: Fairness, Trust, and Emotion in Response to Algorithmic Management." *Big Data & Society* 5 (1): 2053951718756684. https://doi.org/10.1177/2053951718756684.

Longoni, Chiara, Andrea Bonezzi, and Carey K Morewedge. 2019. "Resistance to Medical Artificial Intelligence." *Journal of Consumer Research* 46 (4): 629–50. https://doi.org/10.1093/jcr/ucz013.

Luo, Xueming, Siliang Tong, Zheng Fang, and Zhe Qu. 2019. "Frontiers: Machines vs. Humans: The Impact of Artificial Intelligence Chatbot Disclosure on Customer Purchases." *Marketing Science*, September, mksc.2019.1192. https://doi.org/10.1287/mksc.2019.1192.

Narayanan, Arvind, and Sayash Kapoor. 2024. *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*. Princeton Oxford: Princeton University Press.

Tong, Siliang, Nan Jia, Xueming Luo, and Zheng Fang. 2021. "The Janus Face of Artificial Intelligence Feedback: Deployment versus Disclosure Effects on Employee Performance." *Strategic Management Journal* 42 (9): 1600–1631. https://doi.org/10.1002/smj.3322.

Venkatesh, Morris, Davis, and Davis. 2003. "User Acceptance of Information Technology: Toward a Unified View." *MIS Quarterly* 27 (3): 425. https://doi.org/10.2307/30036540.

Tables

Table 1 Descriptive Statistics

Variable	Mean	Std. Dev.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
(1) *Market Participation*	0.09	0.28													
(2) *AI Matching*	0.33	0.47	-0.04												
(3) *Human Matching*	0.34	0.47	0.02	-0.51											
(4) *Control Group*	0.31	0.46	0.02	-0.50	-0.50										
(5) *Wage*	0.22	0.04	0.03	0.00	0.00	0.00									
(6) *Bachelors*	0.73	0.44	-0.03	-0.02	0.00	0.02	-0.05								
(7) *Masters*	0.14	0.35	0.09	0.03	0.00	-0.02	0.04	-0.61							
(8) *Doctoral*	0.03	0.17	-0.04	-0.01	0.01	0.01	0.02	-0.30	-0.07						
(9) *Graduation Year*	2024	0.622	0.00	0.00	-0.04	0.03	-0.02	0.16	-0.27	0.10					
(10) *STEM*	0.63	0.48	0.01	-0.02	0.01	0.04	0.01	-0.14	0.22	0.00	-0.01				
(11) *Related Field*	0.39	0.49	0.03	-0.01	0.04	-0.01	0.05	-0.13	0.21	-0.03	-0.03	0.21			
(12) *Masters in Related STEM Field*	0.09	0.29	0.06	0.03	0.00	-0.01	0.04	-0.50	0.78	-0.06	-0.22	0.22	0.39		
(13) *Female*	0.44	0.50	0.01	0.00	-0.02	0.06	-0.06	0.04	-0.05	-0.01	0.01	-0.09	-0.08	-0.05	
(14) *Professional Network Size*	2.94	4.17	0.05	0.00	-0.01	0.01	0.01	-0.17	0.24	0.08	-0.16	-0.04	-0.03	0.19	0.00

Table 2 OLS Estimates of the AI-Matching Treatment Effect on Market Participation

	Market Participation			
Dep. Var.:				
Model:	(1)	(2)	(3)	(4)
Experimentally-Assigned Treatments				
AI Matching	-0.025**	-0.024**	-0.024**	-0.024**
	(0.010)	(0.010)	(0.010)	(0.010)
Human Matching	0.000	0.000	-0.001	-0.001
	(0.011)	(0.011)	(0.011)	(0.011)
Wage		0.091*	0.126**	0.118**
		(0.055)	(0.057)	(0.057)
No Wage Shown		0.024*	0.030**	0.026*
		(0.014)	(0.014)	(0.014)
Jobseeker Characteristics				
Bachelors			0.020*	0.019
			(0.012)	(0.012)
Masters			0.096***	0.094***
			(0.018)	(0.018)
Doctoral			-0.027	-0.024
			(0.019)	(0.019)
2024 Grad			0.023**	0.024***
			(0.009)	(0.009)
2025 Grad			0.011	0.010
			(0.018)	(0.017)
2026 Grad			-0.051***	-0.048***
			(0.016)	(0.018)
Job Posting Characteristics				
Job Posting FE's				Y
Constant	0.095***	0.075***	0.028	
	(0.01)	(0.01)	(0.02)	
R^2	0.002	0.003	0.014	0.019
No. Obs.	4562	4562	4235	4235

Notes: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; *Wage* is in units of $100.

Table 3 Differences in the Relationship between Market Participation and Closeness of Match in the AI-Matching Treatment (for STEM = 1)

Dep. Var.:	*Market Participation*					
Subset/Treatment:	*STEM = 1*			*Control Group*	*Human Matching*	*AI Matching*
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Ex Ante Closely-Matched Subgroups						
Master's in Related		0.070***	0.062***	0.066**	0.123***	0.003
STEM Field		(0.017)	(0.018)	(0.033)	(0.033)	(0.029)
Related Field	0.001	-0.020*	-0.020	-0.026	-0.032	0.000
	(0.011)	(0.012)	(0.013)	(0.023)	(0.024)	(0.022)
Individual Characteristics						
2024 Grad			0.014	0.033	0.024	-0.014
			(0.012)	(0.021)	(0.022)	(0.020)
2025 Grad			-0.011	-0.046	-0.002	0.022
			(0.024)	(0.041)	(0.046)	(0.038)
2026 Grad			-0.080	-0.054	-0.065	-0.086
			(0.099)	(0.163)	(0.210)	(0.152)
Job Posting Characteristics						
Job Posting FE's			Y	Y	Y	Y
Constant	0.086***	0.086***				
	(0.008)	(0.008)				
R^2	0.000	0.006	0.015	0.029	0.026	0.008
No. Obs.	2659	2659	2659	877	905	877

Notes: * p < 0.05, ** p < 0.01, *** p < 0.001; using only data for participants from STEM fields; OLS coefficient estimates reported; robust standard errors reported in parentheses.

Table 4 Differences in the Relationship between Market Participation and Professional Network Size in the AI-Matching Treatment

Dep. Var.:	*Market Participation*								
Treatment:	*Control Group*			*Human Matching*			*AI Matching*		
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Jobseeker Characteristics									
Professional Network Size	0.007***	0.006***	0.005**	0.007***	0.006**	0.006**	0.002	0.002	0.001
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)
Female			0.020			0.009			0.023
			(0.021)			(0.022)			(0.019)
50 STEM Subfield FE's			Y		Y				
Job Posting Characteristics									
Job Posting FE's		Y	Y		Y	Y			
Constant	0.069***			0.080***			0.067***		
	(0.011)			(0.012)			(0.010)		
R^2	0.012	0.028	0.104	0.008	0.016	0.060	0.001	0.007	0.082
No. Obs.	877	877	877	905	905	905	877	877	877

Notes: * p < 0.05, ** p < 0.01, *** p < 0.001; OLS coefficient estimates reported; robust standard errors reported in parentheses.

Table 5 Differences in the Relationship between Market Participation and Local Market Size in the AI-Matching Treatment

Dep. Var.:	*Market Participation*					
Subset:	*Entire Study Population*			*Control Group*	*Human Matching*	*AI Matching*
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Market Structural Variables						
ln No. Competitors	-0.017***	-0.017**	-0.017**	-0.017**	-0.041***	-0.007
	(0.006)	(0.008)	(0.008)	(0.008)	(0.012)	(0.013)
Wage			0.027**	0.027**	0.042*	0.016
			(0.013)	(0.013)	(0.023)	(0.023)
Null Wage			0.123**	0.123**	0.280***	0.056
			(0.056)	(0.056)	(0.097)	(0.097)
Individual Characteristics						
STEM		0.001	0.001	0.001		
		(0.010)	(0.010)	(0.010)		
Bachelors		0.031**	0.031**	0.031**	0.044*	0.003
		(0.014)	(0.014)	(0.014)	(0.025)	(0.024)
Masters		0.105***	0.104***	0.104***	0.120***	0.113***
		(0.018)	(0.018)	(0.018)	(0.031)	(0.030)
Doctoral		-0.011	-0.011	-0.011	-0.014	-0.015
		(0.029)	(0.029)	(0.029)	(0.051)	(0.047)
2024 Grad		0.022**	0.022**	0.022**	0.034**	0.039**
		(0.010)	(0.010)	(0.010)	(0.017)	(0.017)
2025 Grad		0.006	0.006	0.006	-0.004	0.021
		(0.019)	(0.019)	(0.019)	(0.033)	(0.036)
2026 Grad		-0.052	-0.055	-0.055	-0.016	-0.062
		(0.089)	(0.089)	(0.089)	(0.122)	(0.204)
Constant	0.085***	0.037**	0.009	0.009	-0.029	0.044
	(0.004)	(0.015)	(0.019)	(0.019)	(0.032)	(0.032)
R^2	0.002	0.013	0.014	(0.01)	0.027	0.018
No. Obs.	4562	4562	4562	1431	1381	1431

Notes: * p < 0.05, ** p < 0.01, *** p < 0.001; OLS coefficient estimates reported; robust standard errors reported in parentheses.

Table 6 Differences in AI-Matching Treatment Effects across STEM and non-STEM Jobseekers

Dep. Var.:	*Market Participation*					
Subset:	*STEM = 0*			*STEM = 1*		
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Experimentally-Assigned Treatments						
AI Matching	-0.035**	-0.036**	-0.036**	-0.015	-0.017	-0.016
	(0.018)	(0.018)	(0.018)	(0.013)	(0.013)	(0.013)
Human Matching	-0.017	-0.016	-0.016	0.011	0.010	0.010
	(0.018)	(0.019)	(0.019)	(0.013)	(0.014)	(0.014)
Wage		0.155*	0.158**		0.081	0.072
		(0.080)	(0.080)		(0.096)	(0.096)
No Wage Shown		0.028	0.023		0.020	0.017
		(0.020)	(0.020)		(0.024)	(0.024)
Jobseeker Characteristics						
Bachelors		0.034*	0.032*		0.035**	0.032**
		(0.018)	(0.017)		(0.016)	(0.016)
Masters		0.068	0.070		0.114***	0.104***
		(0.043)	(0.043)		(0.022)	(0.023)
Doctoral		-0.033	-0.025		0.001	0.001
		(0.027)	(0.028)		(0.027)	(0.028)
2024 Grad		0.020	0.019		0.023**	0.023**
		(0.015)	(0.015)		(0.011)	(0.011)
2025 Grad		0.033	0.028		-0.005	0.000
		(0.033)	(0.033)		(0.020)	(0.020)
2026 Grad		-0.027	-0.008		-0.060***	-0.063***
		(0.025)	(0.028)		(0.020)	(0.022)
Job Posting Characteristics						
Job Posting FE's			Y			Y
Constant	0.100***	0.03		0.088***	0.02	
	(0.013)	(0.026)		(0.009)	(0.026)	
R^2	0.003	0.013	0.021	0.001	0.017	0.022
No. Obs.	1439	1439	1439	2659	2659	2659

Notes: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; *Wage* is in units of $100.

Appendix

Table A1 Follow-on Survey Responses Means and Standard Deviations by Treatment

Question	Response Range	Treatment		
		AI Matching	Human Matching	Control Group
Q1 Feelings: Did you experience any negative or positive feeling (even slightly?) about the above email text?	1 ("Negative Feeling") to 5 ("Positive Feeling")	2.800 (0.645)	3.393 (0.956)	3.038 (0.824)
Q2 Interest: *Do you think you'd be interested in this job?*	1 ("Definitely not interested") to 5 ("Definitely interested")	2.680 (1.108)	3.500 (1.106)	3.231 (1.177)
Q3. No. Competitors: *Guesstimate how many other people were targeted and received this same*	"Write down a number" [Reported in 1000s]	1.756 (3.264)	1.044 (1.414)	2.744 (3.227)
Q4 Expected Match Quality: From the email text, above, do you think your skills are well matched to this position? (assuming you were available and	1 ("Definitely not a good skills match") to 5 ("Definitely a good skills match")	3.280 (1.100)	3.750 (1.295)	3.320 (1.145)
Q5. Expected Chances: If you pursued this job, how would you have rated your chances of getting it?	1 ("No chance") to 5 ("Great chance")	3.080 (0.954)	3.250 (0.967)	3.000 (1.131)
	No. Obs.	25	28	26

Notes: Standard deviation reported in parentheses.