

# AI-Enabled Job Markets & Market Participation: A Field Experiment on how AI Shapes Jobseekers' Expectations of Competition

Sarah Bana  
Chapman University

Kevin J. Boudreau\*  
Northeastern & NBER

## Abstract

Artificial intelligence is increasingly mediating how jobseekers are matched with employers, raising fundamental questions about how its use affects participation in labor markets. We examine whether AI-based matching alters jobseeker behavior—specifically, their willingness to participate. Drawing on a field experiment with 4,562 jobseekers randomly assigned to disclosure conditions, we find that participation was about one-quarter lower when AI use was disclosed than in either a human-matching treatment or a control group with no source specified. Participation responses varied systematically, consistent with jobseekers forming expectations about how AI affects (i) predicted match quality, (ii) the types of inputs and information used, and (iii) the scale of competition. These relationships were strongest among jobseekers with greater familiarity with AI, proxied by STEM backgrounds. Our findings highlight how AI disclosure can reshape both the level and composition of participation, as AI alters expectations of payoffs, with implications for platform design, policy, and the governance. More broadly, the results underscore the importance of distinguishing structural payoff effects—which may persist over time—from subjective attitudinal reactions to AI, which may be more transient.

**Keywords:** Artificial Intelligence (AI), Human–AI Interaction, Digital Platforms, Job Markets, AI Aversion, Field Experiment, Platform Governance.

**JEL Codes:** J01, J20, M15, O3, L86, C93

\* For valuable comments and suggestions that helped shape this study, we thank seminar participants at the NBER Summer Institute program for Digital Economics and Artificial Intelligence, Wharton's AI and the Future of Work conference, USC's Artificial Intelligence in Management (AIM) Conference, HEC Paris AI and Entrepreneurship workshop, ISB's AI and Strategy Consortium, the NBER Productivity & Innovation Seminar, the Workshop on Information Systems and Economics, Symposium on Statistical Challenges in eCommerce Research. We are grateful for excellent research assistance from Aneesh Arunjunai Saravanan, Conor Mackey, Ethan Mandel, Jake Dwyer, Lokhi Nalam, Shruti Srinivasen, Yogesh Arjun Kumar, Anand Hegde, Adam Ma, and Steven Abbott. This research was determined to be exempt by the Northeastern University Institutional Review Board (NUIRB180419). We gratefully acknowledge financial support from the Kauffman Foundation and Northeastern University's D'Amore-McKim School of Business. Correspondence: k.boudreau@northeastern.edu. All errors are our own.

# 1 INTRODUCTION

Job markets have now been reshaped by successive waves of digitization and platforming over recent decades, gradually altering how employers and jobseekers connect, compete, and evaluate one another since the turn of the century (Autor 2001; Horton 2017; Goldfarb & Tucker 2019). In the past decade, in particular, advances in datafication and machine-based intelligence have introduced a new stage in this evolution, where identifying and matching relevant workers with jobs is increasingly mediated by algorithmic and AI-driven systems (Brynjolfsson et al., 2018). These tools—now used in at least some part of recruitment by over 90% of large U.S. firms (Fuller & Raman, 2021)—range from keyword filters and ranking algorithms to more advanced systems leveraging embeddings, large language models, and text generation (The Economist 2023), reflecting employer efforts to address the complexity and costs of recruitment (Oyer & Schaefer 2011; Tambe et al. 2019). Thus far, much research attention has been devoted to improving the technical performance and information processing of AI-based matching systems.<sup>1</sup> Far less is known about whether and how their use alters the behavior of job market participants. In this paper, we investigate whether the use and disclosure of AI-based matching affects a most basic behavioral choice of jobseekers: whether to participate in a job market at all. We argue that AI use and disclosure may change participation decisions to the extent that they alter jobseekers’ expectations about the likely payoffs to pursuing a position and the competition they will face.

Although little research has examined these behavioral effects in job markets directly, a growing body of work in information systems, behavioral science, and human–AI interaction shows that people systematically respond differently when interacting with AI systems rather than human intelligence in consequential decisions. These responses—often termed AI “aversion” or “appreciation” in the literature depending on the sign of the effect found—have been observed in a wide variety of settings from forecasting, performance evaluations, medical decisions, and hiring (Section 2.2). Prior work proposes many possible explanations for these patterns, which can be broadly grouped into (i) subjective attitudinal reactions (e.g., perceptions of fairness, legitimacy, empathy, or bias) and (ii) payoff-based expectations (anticipations of how AI use will affect the benefits or outcomes of participation). In this study, we focus on payoff-based mechanisms, as the most direct and parsimonious explanations for participation behavior. Because these mechanisms rest on

---

<sup>1</sup> For recent advances in algorithmic job matching, see studies that develop novel recommender systems (Le Barbanchon, Hensvik, and Rathelot 2023), propose incorporating richer attribute sets and time-varying components (Kokkodis and Ipeirotis 2023), and examine algorithmic biases in matching outcomes (Zhang and Kuhn 2024; Wall and Schellmann 2021).

informed expectations, they might also generate enduring rather than transient responses to AI in job markets (Turel and Kalhan, 2023).

If any such behavioral responses to AI were to emerge in job markets, the consequences would be particularly important because the question of who chooses to participate and compete is vital to the functioning of these markets (Mortensen 1986; Pissarides 2000). Participation matters not only in terms of the number of jobseekers but also in terms of which jobseekers enter: decisions to participate involve costly investments in exploring opportunities, signaling interest, and competing for roles. If highly qualified or best-matched candidates disproportionately opt out, matches may be less efficient and equitable. Conversely, if disclosure encourages broader participation by underrepresented groups, markets may become thicker and more inclusive. Because labor markets are the primary mechanism by which human capital is allocated across firms and industries, distortions in participation can cascade through the economy, shaping productivity, innovation, and growth far beyond the labor market itself.

It is equally important to make progress in discerning the underlying reasons or mechanisms driving any change in job market participation arising from the use and disclosure of AI in matching. Structural shifts in competition that alter jobseekers' expectations of payoffs would reflect enduring shifts in equilibrium behavior, perhaps persisting unless the design of AI-based systems is fundamentally changed. By contrast, differences rooted in society's current subjective attitudes could plausibly evolve or dissipate over time as knowledge, experience, and understanding of AI's role in labor markets diffuse.

In this initial exploration of behavioral responses to AI use and disclosure in job markets, we focus on whether disclosure alters market participation choices by shifting jobseekers' expected payoffs from participating. Section 3 develops the full theoretical framework; here we briefly outline it. We consider three channels: jobseekers may expect AI systems to differ in match quality relative to human judgment; to rely on different inputs to prediction, more heavily on codified data and statistical inference rather than soft information, intuition, and human networks; and to lead to a different scale of competition, either broadening or narrowing the pool of candidates that jobseekers will face as competitors. Because such reasoning should be sharper among those with greater knowledge of AI, we analyze heterogeneity by technical background, using STEM training as a proxy. By introducing a randomized wage treatment, we also benchmark the effect of AI disclosure against a tangible monetary incentive to assess its practical significance.

To investigate these ideas, we ran a large-scale randomized field experiment on a live university-based job platform for high-skilled project roles. Our design held the underlying match process constant while randomly varying whether jobseekers were told their recommendation came from an AI system, from a human, or a control group with no source specified. We achieved this experimental variation while—crucially—maintaining full truthfulness and avoiding any deception of subjects. Our experiment captures the decisions of 4,562 jobseekers and measures how disclosure affects participation. We then conduct a series of tests exploiting within-sample variation to assess evidence regarding the three hypothesized channels: perceived match quality, reliance on codified versus network-based information, and expectations about the scale of competition. Because we expect these mechanisms to be sharper among jobseekers with greater understanding of how AI operates, we also analyze heterogeneity by technical background, using STEM training as a rough proxy for such knowledge. We attend to the possibility that STEM could potentially be correlated with other issues in our research design and econometric checks. By also introducing a randomized wage treatment, we benchmark the behavioral effect of AI disclosure against a tangible monetary incentive, enabling us to assess its relative magnitude and practical significance.

Our empirical analysis begins by estimating the overall treatment effect of AI disclosure. Consistent with much of the prior work on human–AI interaction, market participation is reduced by AI disclosure, reducing participation by about one-quarter compared to both the Human and control conditions (a 2.4 percentage-point drop from a 9.53 % baseline, or about  $7.5\times$  the impact of a \$10 wage reduction).

Our empirical analysis then turns to testing our main theoretical predictions, where we find patterns broadly consistent with shifting payoff expectations, especially among jobseekers with STEM backgrounds who are more likely to form reasoned beliefs about how AI operates. The reported patterns are consistent with informed jobseekers expecting lower match quality, diminished returns to human-based information, and a greater number of competitors beyond just the local market under the AI-Matching treatment. The responses of those with technical training backgrounds align more closely with our empirical hypotheses than those of non-STEM participants, consistent with greater knowledge of machine-based prediction shaping expectations of payoffs. Importantly, however, there is a larger overall negative effect of AI-Matching on market participation in the case of non-technical jobseekers than for those with technical backgrounds—suggesting that additional mechanisms beyond those emphasized here in this study also merit future work.

These findings have important implications for the design and governance of AI-enabled job markets. Even when disclosure does not reduce the total number of jobseekers, it can alter which candidates choose to participate, shifting the composition of the applicant pool in ways that affect efficiency and equity (Mortensen 1986; Pissarides 2000). Because some responses are consistent with rational expectations about structural differences in competition created by AI, these effects may persist or even strengthen over time, unlike attitudinal aversion which may eventually dissipate as knowledge and experience grows over coming years among jobseekers.

Our primary contribution is to extend the human–AI interaction literature into the domain of job markets, showing first that large AI-Matching treatment effects exist that affect both levels and composition of those willing to pursue and compete for positions. We also show how use and disclosure of AI reshapes jobseekers’ participation and competition through payoff-based mechanisms in particular (Dietvorst et al., 2015; Logg et al., 2019; Castelo et al., 2019; Longoni et al., 2019). We attempt to go further in documenting discriminating tests to support the existence of this especially important category of mechanisms, building on prior work that has emphasized disclosure and transparency in human–AI interaction (Luo et al. 2019; Tong et al., 2021). To situate this contribution within information systems, we also build on recent integrative reviews in MISQ and related outlets that highlight algorithm aversion and the socio-technical design of AI systems (Burton, et al., 2020; Diederich et al., 2022). In doing so, we highlight implications for the growing research tradition on the digitization of labor markets (Horton 2017; Kuhn & Mansour 2014; Goldfarb & Tucker 2019), and for platform governance and algorithmic management (Lee 2018; Burton et al. 2020).

The remainder of this paper is structured as follows. Section 2 describes the prior literature. Section 3 develops the hypotheses with a guiding characterization of a jobseeker's decision. Section 4 describes the experimental design and sample. Section 5 presents the analysis and results. Section 6 summarizes and discusses results with implications for platform design and policy and contributions to the literature.

## **2 RELATED LITERATURE: DIGITIZATION AND HUMAN–AI INTERACTION**

Our focus in this study is on a most basic behavioral choice in job markets: whether jobseekers decide to participate when AI is disclosed in the matching process. Direct evidence on our question is sparse, but two adjacent literatures help situate our contribution and clarify what is known versus unknown on this question. First, work on the digital transformation of labor markets documents how digitization and platforms have reshaped matching and introduced algorithmic intermediation.

Second, the newer human–AI interaction literature shows how people’s behavior systematically changes when they learn AI is involved in consequential decisions. Both point to our question, but neither has yet addressed it directly.

## 2.1 Digital Transformation of Labor Markets

Over the past two decades, research in information systems, management, and economics has developed a substantial body of work examining how the digitization of markets reshapes economic exchange (Goldfarb & Tucker 2019). Digital technologies have reduced frictions by lowering search and transaction costs, increasing the speed and precision of matching, and enabling continuous market participation regardless of geography (Bakos 1997; Brynjolfsson & Smith 2000; Autor 2001). The rise of online platforms has also transformed the informational infrastructure of markets: reputation systems, ratings, and reviews have made previously tacit quality signals visible and portable across transactions (Cabral & Hortaçsu 2010; Dellarocas 2003), while advances in data collection and analytics have “datafied” participant attributes and behaviors, enabling algorithmic intermediation at scale (Mayer-Schönberger & Cukier 2013; Brynjolfsson et al., 2018). In labor markets specifically, digitization has, for example, expanded the geographic scope of competition, shifted matching from local, network-based processes to global, platform-mediated systems (Horton 2017), and created new governance regimes in which platforms control visibility, ranking, and access through proprietary algorithms (Kellogg et al. 2020). These shifts have redefined the structure of competition and the ways that jobseekers and employers discover and evaluate one another. While algorithmic matching is now a common feature of digital labor platforms, the specific role of machine-based intelligence (AI/ML systems)—and how its use is disclosed to participants—remains a distinct and comparatively underexplored topic within this broader digitization literature.

A handful of studies speak most closely to our setting, though none address disclosure or jobseeker participation. Horton (2017) showed that algorithmic candidate recommendations increased fill rates by 20% in an online labor market. Hoffman et al. (2018) found that algorithmic pre-employment test scores provided useful signals—managers who overrode them made worse hires. More recently, Wiles et al. (2023) showed that AI-assisted résumé writing improved jobseeker outcomes, and Wiles and Horton (2024) found that AI-generated job postings increased employer posting but not matches. Together, these studies demonstrate that automation can alter outcomes, but they stop short of examining how disclosure of AI affects jobseeker participation decisions—the focus of our study.

## 2.2 Human-AI Interaction: AI Aversion, Appreciation, and Behavioral Responses

Although digitization research does not yet address jobseeker behavior under AI disclosure, a separate literature on human–AI interaction provides strong motivation. This work consistently reports differences in how people respond to machine- versus human-based judgments, most often findings AI “aversion” and sometimes AI “appreciation.” Several excellent surveys are provided by Burton et al. (2020), Mahmud et al. (2020), Jussupow et al. (2023), and Diederich et al. (2022).<sup>2</sup>

### 2.2.1 Subjective Attitude Explanations

A wide range of studies document lowered willingness to engage with AI-based systems, with proposed explanations that can broadly be grouped as related to subjective attitudes, perceptions, and preferences. For example, several studies interpret results in terms of perceived legitimacy and fairness. Tong et al. (2021) found that employees rated performance reviews more negatively once they learned they came from an AI system; the authors interpreted this as reduced perceived legitimacy. Bansak et al. (2021) reported that citizens were less supportive of algorithmic tools for refugee resettlement compared to human decision-making, despite similar outcomes, interpreting this as linked to fairness and accountability concerns. Green and Chen (2019) examined judges’ use of algorithmic risk assessment scores in bail and sentencing and found that disclosure influenced judicial decisions: judges anchored on the scores even while expressing doubts about fairness. They interpreted this as legitimacy concerns coexisting with partial reliance, producing both aversion and deference. Arnold et al. (2018) showed that judges imposed systematically harsher bail on Black and Hispanic defendants. Although this study did not involve AI directly, the evidence of human bias has been interpreted as motivating interest in algorithmic interventions as potentially fairer alternatives, underscoring how legitimacy concerns shape attitudes toward AI.

Another recurring theme is sensitivity to errors. Dietvorst et al. (2015) showed in forecasting tasks that people abandoned algorithms after seeing them make mistakes, even when the algorithms were still more accurate than humans. They interpreted this as heightened sensitivity to algorithmic errors relative to human ones. Similarly, Prahla and van Swol (2017) found that participants penalized algorithmic forecasters more heavily than human forecasters after errors, again suggesting asymmetry in tolerance that reflects attitudinal bias.

---

<sup>2</sup> Note that beyond questions of positive or negative valence towards machine-based intelligence, several studies have begun to make progress in characterizing interactions between machine-based and human intelligence, along with rich related organizational questions (e.g., Fügener et al., 2021; Kim et al., 2024; Stelmaszak et al., 2025).

Perceptions of social presence and moral sensitivity also matter. Longoni et al. (2019) found that patients preferred human over AI medical providers, interpreting this as concerns about loss of individualized care. By contrast, Bojd et al. (2024) found that patients in stigmatized medical contexts were more willing to engage with AI providers, reasoning that the absence of social judgment and greater anonymity made AI more attractive. Starke et al. (2022), in a review of IS and management research, emphasized that opacity and unfairness perceptions are central drivers of algorithm aversion. Haesevoets et al. (2021), reviewing evidence from moral and social contexts, argued that deficits in empathy and social understanding likely underlie resistance to AI when contextual nuance is valued. Yeomans et al. (2019) added nuance by showing that algorithmic advice was more readily accepted in structured, quantifiable tasks, reinforcing that the context strongly conditions attitudinal responses.

### **2.2.2 Payoff Based Explanations**

While much of this literature emphasizes subjective attitudes, we see growing evidence that expectations of payoffs also matter. It is this second line of explanation that we develop and test in the present study.

For example, Bojd et al. (2024) found that patients were more willing to engage with AI providers in stigmatized medical contexts than with human providers, interpreting this as evidence that participants expected AI to reduce reputational or social costs, thereby increasing the net payoff of engagement. Castelo et al. (2019) ran a series of experiments comparing algorithmic and human decision-makers and found that people preferred algorithms for objective, rule-based judgments (e.g., statistical tasks) but humans for socio-emotional ones (e.g., dating advice), consistent with participants updating expectations about which decision-maker would yield better outcomes in a given domain. Logg et al. (2019) similarly found that people sometimes preferred algorithmic over human advice in forecasting ( “algorithm appreciation” ), but that appreciation diminished once errors were revealed—suggesting that expectations are revised in light of performance information. Sinclair-Desgagné (2024) provided a formal decision-theoretic model that reconciles such patterns, showing that willingness to engage with AI depends on expected error rates. Albright (2024) studied judges in Kentucky (2011-2013) and found they were more likely to follow algorithmic recommendations when those recommendations were more lenient, interpreting this as judges expecting that deference to AI could reduce reputational risk—again a payoff-based logic. Closest to our context, Avery et al. (2024) report in a working paper that female job applicants were more likely to complete applications when assessed

by AI recruiters than by humans, interpreting this as evidence that women expected less systematic prejudice from AI evaluators than from humans.<sup>3</sup>

### 2.3 Summary and the Current Study

Taken together, the human–AI interaction literature demonstrates effects across diverse contexts, with explanations that fall into two broad categories: subjective attitudes and rational expectations of payoffs. Whereas prior studies have largely documented broad patterns of “aversion” or “appreciation,” our contribution is to provide discriminating evidence of the specific mechanisms—especially payoff-based ones—that shape participation in job markets.

This study offers some of the first field-experimental evidence on how AI use and disclosure influence jobseekers’ participation. We focus on how jobseekers anticipate AI use and disclosure will affect their payoffs, as a first and simplest basis for developing hypotheses. Beyond documenting overall aversion or appreciation, we seek evidence of the particular mechanisms at work in labor markets, beginning with the most straightforward: payoff-based expectations.

## 3 THEORY & HYPOTHESIS DEVELOPMENT

As discussed above, our starting premise is that AI use and disclosure may alter jobseekers’ expectations of payoffs from participating. Our hypotheses focus exclusively on payoff-based channels. We do not attempt to adjudicate attitudinal explanations here, except to note them as potential alternative interpretations within our empirical analysis.

In a nutshell, we claim that—holding the nature of jobseekers, jobs, and matches constant—the effect of AI use and disclosure per se will be to shift jobseekers’ expectations of (i) the quality of match, as in H1, below; (ii) the sources of information and inputs used to generate match predictions, as in H2; and (iii) the scale of competition, as in H3. Our empirical hypotheses are devised to address the possibility that the direction of effects as regards (i) and (iii) are ambiguous. Further, important to note, these mechanisms presuppose that jobseekers are sufficiently knowledgeable or experienced with AI to form such expectations, which we address in H4.

---

<sup>3</sup> There have been few published recruitment-related studies on behavioral or attitudinal responses to AI and most have relied on non-randomly sampled, self-reported survey evidence, or otherwise do not relate to our question. Horodyski (2023), surveying 552 self-selected respondents recruited through social networks, found concerns about lack of human touch, possible inaccuracy, and insufficient transparency. Lee (2018), using survey data from ride-share drivers, similarly documented perceptions of algorithmic management as unfair and untrustworthy. Luo, Zhang, and Mu (2025) report a survey-based experiment based on hypothetical scenarios, with respondents reporting that disclosure of ‘robo-interviews’ reduced stated willingness to apply

To help ground and motivate our characterization of jobseekers’ participation decisions, we use a simple illustrative expected-payoff framework. A jobseeker participates when the anticipated benefits of doing so outweigh the costs:

$$E[p \cdot W - c] \geq \delta$$

Here,  $p$  is the perceived probability of being hired,  $W$  is the expected value of the position; and  $c$  is the expected cost of pursuing and competing for the role, and  $\delta$  is the threshold outside option value of a jobseeker. Our interest here is in whether using and disclosing that a match was generated by AI rather than human predictions changes the likelihood that this inequality holds. In the remainder of this section, we consider three channels through which AI might alter the competitive environment and bear on this condition for participating.

### 3.1 AI & Expectations of Match Quality

The first channel we consider is how AI disclosure may shape jobseekers’ expectations about match quality—that is, how accurately they expect the system will recognize their true fit for a role. Because perceived match quality feeds directly into  $p$ —with higher-quality predictions increasing the probability of getting the job—and may also influence  $W$  when fit is well predicted, changes in expected accuracy have a direct path to influencing participation through the payoff condition.

A key challenge is that expectations regarding the relative accuracy of AI versus human predictions may vary by context. Variation could stem from the type of model (e.g., large language models versus domain-specific matching algorithms), the maturity of the technology in a given setting, or the breadth and quality of the training data available. For example, jobseekers might view a well-established, narrowly focused matching tool trained on extensive domain-specific data as highly accurate, while seeing a newly deployed general-purpose model as more experimental or error-prone. In some contexts, AI may be expected to deliver more objective or consistent evaluations; in others, applicants may fear it will overlook important indicators of fit, particularly those that are qualitative or difficult to quantify. These beliefs may also shift over time as models improve, become more widely understood, or face new concerns such as vulnerability to gaming or brittleness in less common scenarios.

Whether AI is expected to generate better or worse predictions now or in the future remains hotly debated; despite accelerating adoption and great hopes, there are also high-profile examples of disappointing performance in recruitment contexts (Schellmann, 2023). Whatever the direction of expectations, candidates who are strong ex-ante matches for a given job should be the most sensitive

to them. If AI is expected to be more accurate than human judgment, they stand to gain disproportionately; if AI is expected to be less accurate, they have the most to lose. In the first case, they may anticipate a higher probability of being correctly identified as a top fit, making participation especially attractive. In the second, they may fear their quality will be overlooked, lowering their perceived chances of success and discouraging participation. **In** short, the relative effect of AI disclosure—positive or negative—should be amplified among the strongest ex-ante matches. We summarize these points in the following hypothesis:

**Hypothesis 1 (H1):** If jobseekers expect AI-based recommendations to differ in match quality from non-AI processes, then ex-ante well-matched candidates will be differentially affected in their participation decisions, with their participation either being more positively or more negatively affected, depending on whether AI is expected to be more or less accurate.

### 3.2 AI & Expectations of the Information and Inputs Used to Generate Predictions

A second channel through which AI disclosure may shift participation is by changing expectations about how candidates will be evaluated—specifically, the types of information that serve as inputs to human versus machine information processing systems. If jobseekers expect that AI-based systems will discount or fail to recognize such assets, they may perceive lower returns to participation, perhaps reducing the expectation of successfully securing the position,  $p$ , or reducing the expected value term  $W$ , or increasing costs necessary to secure a position given they cannot use existing assets so readily.

For example, in human-based decisions and matching, one of the most important and well-documented inputs is information transmitted through personal and professional networks. Human recruiters and managers process candidates through network-mediated channels where existing employees provide referrals, endorsements, and informal assessments of match quality (Granovetter 1995; Autor 2001). These referral networks allow firms to access information about candidate quality that is difficult to observe through standard application processes—for example, in instances referred workers have lower turnover rates, fewer safety incidents, and higher innovation outputs, even when their observable qualifications match those of non-referred workers (Burks et al. 2015). The network structure itself becomes a critical information channel where approximately 50% of U.S. jobs are found through informal networks, and about 70% of firms maintain formal employee referral programs that systematically produce higher-quality matches (Burks et al. 2015). Human evaluation systems leverage these network-based information flows where the credibility and relevance of candidate information depends not merely on credentials but on the social relationships and

professional connections through which recommendations flow, allowing firms to identify candidates with superior fit and performance potential that would be undetectable through standard screening processes (Stein 2002; Gibbons & Waldman 2004).

By contrast, algorithmic systems tend to rely more heavily on codified inputs—such as structured résumé data, job titles, education history, keyword matches, and standardized skill profiles—and translate these into probabilistic match predictions using formal models (Baker 1992; Brynjolfsson & McAfee 2014; Tambe et al. 2019). More recent systems can incorporate unstructured data using techniques like embeddings, LLM-powered résumé parsing, sentiment analysis in asynchronous video interviews, or automated evaluation of portfolios and code samples. Yet even these systems ultimately remain dependent on codified data sources, rather than interpersonal and, sometimes, tacit context.

Unlike the match-quality channel considered in H1, where AI disclosure could plausibly encourage or discourage participation depending on perceived relative accuracy, here we expect the direction of the effect to be more consistently negative for those whose strengths lie in human-centered inputs. If jobseekers expect that AI-based systems will discount or fail to recognize such assets, they may perceive lower returns to participation in AI-mediated processes. In other words, AI disclosure may differentially discourage participation among candidates whose strengths are aligned with human-centered evaluation systems. The following hypothesis summarizes these ideas:

**Hypothesis 2 (H2):** Jobseekers who have invested in assets oriented toward human-based evaluation—such as professional networks or interpersonal experience—will be differentially less likely to participate when matching is mediated by AI.

### 3.3 AI & Expectations about the Scale of Competition

The third channel we consider is how AI disclosure may change jobseekers' expectations about the scale of competition they will face should they choose to pursue a position. Such expectations influence  $p$ , the perceived probability of being hired, by altering the number of competitors they expect they need to beat out. Jobseekers typically weigh not only their own qualifications but also the perceived size and quality of the competition when deciding whether to pursue a position (Gee 2019; Fradkin, Bhole, & Horton 2025). From an economic perspective, many job markets—especially for high-skilled or desirable roles—can be conceptualized as tournaments, where outcomes depend not only on absolute performance but also on relative standing of candidates (Lazear & Rosen 1981; Waldman 1984).

AI-based matching systems can plausibly shift these competition expectations in opposite ways. On one hand, jobseekers may associate AI with automation and scalable outreach, leading them to expect broader dissemination of job recommendations and larger applicant pools. Unlike human evaluators, who are constrained by time and attention, AI systems can efficiently process thousands of applications at minimal cost (Horton 2017; Brynjolfsson & McAfee 2014). If AI increases the perceived number of competitors,  $p$  may fall, reducing the expected payoff from participation. On the other hand, jobseekers may believe that AI enables more targeted and selective matching, narrowing the pool of relevant candidates. Advances in machine learning allow for fine-grained inferences about candidate fit, potentially leading to more precision-focused recommendations (Tambe et al. 2019; Cowgill et al. 2021). In this case,  $p$  may increase—not because the market is smaller in absolute terms, but because the jobseeker expects to face fewer, better-matched competitors.

Although these mechanisms point in opposite directions, both imply that disclosure of AI-based matching should reduce the salience of local labor market size as a competition cue. If AI expands reach, jobseekers may perceive the local market of competitors as less relevant, since employers can draw from a wider pool. If AI enables tighter targeting, local market size may again matter less, as competition is refined to a smaller set of algorithmically identified candidates. In either case, participation decisions may become less responsive to the size of the local labor market, representing a shift in how competition is structured and how  $p$  is formed in the participation decision.

**Hypothesis 3 (H3):** Jobseekers expect AI-based recommendations to alter the scale of competition in ways that make participation less sensitive to the size (number of competitors) within the local labor market.

### 3.4 The Role of Knowledge and Familiarity

The three hypothesized channels above each rest on the idea that jobseekers respond to AI disclosure by forming expectations about how it alters the payoff structure of market participation—whether through changes in match quality, the types of inputs considered, or the scale of competition. All of these mechanisms presuppose that jobseekers are sufficiently knowledgeable or experienced with AI to form such expectations.

**Hypothesis 4 (H4):** The participation effects in H1–H3 will be stronger among jobseekers with greater knowledge or experience of AI, who can form more informed expectations about its impact on payoffs.

Where relevant knowledge and experience of AI and its effects may be limited, the predictions in H1–H3 cannot be expected to necessarily operate as cleanly or as outlined above. Although our emphasis in this study is to investigate whether market participation patterns conform to jobseekers

forming expectations of changing payoffs with AI use and disclosure, we speculate that there could also be greater scope for the ‘subjective attitudinal’ mechanisms described in Section 2.2 to play a role.

## **4 FIELD EXPERIMENTAL RESEARCH DESIGN**

To study whether AI-based recommendations have any systematic effect on jobseeker participation choices and to test the three hypothesized channels (Section 3), we conducted a randomized field experiment on a live job platform. The design was shaped by several methodological requirements: (i) introduce clean, interpretable variation across treatment arms—comparing recommendations labeled as AI-generated, human-generated, or unspecified—while holding constant the underlying matching process; (ii) ensure that the jobseeker population, job postings, and match quality were identical across groups so that any behavioral differences could be attributed solely to disclosure; (iii) preserve truthfulness and avoid deception in labeling recommendations; (iv) allow within-sample variation and measures in characteristics to test our hypotheses; (v) secure a large and diverse sample for statistical power; and (vi) operate in a context where matching was non-trivial and jobseekers had meaningful stakes, ensuring external validity.

We partnered with a university-based platform for high-skilled, short-term, and part-time project roles. Rather than vary the underlying prediction methods, we created a hybrid recommendation process incorporating both human and machine steps, then randomized whether the human or AI source was disclosed to participants. This allowed us to truthfully describe the same match as either AI- or human-generated, preserving experimental integrity and avoiding deception. Using detailed data on jobseekers and roles, along with rich within-sample variation, we test whether AI disclosure systematically alters participation and whether these effects align with our three hypothesized mechanisms.

### **4.1 Experimental Protocol & Treatments**

Our field experiment took place on a university-based job platform focused on short-term, high-skilled project roles. This setting ensured a non-trivial matching environment, where jobseekers possessed substantial technical or professional training, and roles demanded meaningful, specialized capabilities. The employer launched a fixed set of job postings—ranging from cryptocurrency applications and full-stack engineering to data architecture, marketing strategy, video production, and psychology-related projects—over a concentrated three-week period.

#### **4.1.1 Platform Identification of Relevant Candidates & Study Population**

To identify relevant candidates, an approach was devised that incorporated both human and AI inputs. Hiring managers and platform staff worked together to define job-specific keywords that reflected appropriate fields of study, technical skills, and prior experience. This first step in the process ensured that the matching was grounded in substantive human input. A machine-based simple Boolean match algorithm was then encoded to generate matches between these keywords and the public LinkedIn profiles of all university-affiliated jobseekers. The result was a hybrid matching procedure: while the execution was algorithmic, the matching logic was shaped by human inputs. This design allowed us to truthfully describe the same recommendation as either AI- or human-generated, depending on the treatment condition.

This matching method yielded a broad, clearly defined sample of 4,562 jobseekers who met basic relevance criteria for at least one job. We intentionally used a simpler, less discriminating process than the platform's normal recommendation engine to: (i) preserve truthfulness in both AI and human labels, (ii) generate a large pool of relevant candidates (anyone with matching education or expertise), (iii) create greater variation in match quality to exploit in analysis, and (iv) observe the candidate pool (risk set) prior to treatment.

Crucially, this approach made the entire matched candidate pool observable ex-ante. Among the 4,562 jobseekers, 2,769 (62.8%) are classified as coming from a STEM. This classification is central to our hypotheses and guides all stratified analyses. Finally, another advantage of this way of defining all potentially relevant competitors within this "local" closed job market is that our approach also provides a reasonable estimate of the total number of potential competitors who could reasonably pursue a given position within a local market. We will use this feature of our data in testing H3 within Section 5.4.

#### **4.1.2 Random Assignment to Recommendation Source Disclosures**

Each jobseeker received a single job recommendation email, based on the keyword-based match. Each email included a job title, a brief one-line description, an expected wage (see below), and a link to the platform where further information would appear on the position and where it would be possible to apply.

The jobseekers were randomized into one of three conditions related to prediction source. 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."

To accommodate varied interpretations of machine-based decision-making, we deliberately included both “AI” and “algorithm” in the AI treatment label. At the time of the experiment, rule-based systems, natural language processing, and machine learning all coexisted alongside emerging large language models. Our framing and hypotheses were therefore intended to generalize beyond the specific state of technology.<sup>4</sup> Using both “AI” and “algorithms” was therefore intended as a means of embracing the range of meanings and expectations that subjects might hold.

#### **4.1.3 Random Assignment to Wage**

In addition to the disclosure treatment, we implemented a second orthogonal wage treatment, in which the expected wage was randomly drawn from a range of \$16 to \$28 per hour, in \$2 increments, while a control condition received no wage information. It was possible to randomly assign wage given the particular institutional details of our context. This was feasible given institutional norms, where job postings often listed an illustrative range but the final rate was determined based on a candidate’s education level and program of study. Manipulating wage expectations simultaneously to disclosure allows benchmarking of magnitudes of effect sizes.

## **4.2 Data Set & Variables**

Our dataset combines information from the platform's customer relationship management (CRM) system, LinkedIn profiles, and educational data matched to standardized field classifications. The data set is a cross-section of all individuals who were part of the risk-set, i.e., those people who were keyword-matched to the job postings and who then received match recommendations. This results in a cross-section of 4,562 observations, divided evenly among the AI-Matching treatment, Human-Matching treatment, and the control group. Main variable descriptive statistics appear in Table 1.

---

<sup>4</sup> Computer Science researchers, 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.”

**Table 1 Means and Standard Deviations**

	<i>Mean</i>	<i>Standard Deviation</i>
<i>Market Participation</i>	0.087	0.282
AI Matching	0.334	0.472
Human Matching	0.337	0.473
Wage	0.225	0.418
Wage Missing	0.174	0.099
STEM	0.628	0.483
Masters	0.140	0.347
Doctoral	0.031	0.174
Ex-Ante Well-Matched [ <i>Master's in Rel. Fields</i> ]	0.091	0.287
<i>Prof'l Network Size</i> [No. LinkedIn Connections]	2.949	4.173
<i>log No. Local Competitors</i>	6.728	0.738
Female	0.466	0.499
<b>Grad Year</b>	<b>2023.512</b>	<b>0.621</b>

Notes. Number of observations = 4,562.

As a main dependent variable, we require a measure of an individual's choice to engage in search activity—to make some minimum effort to gain further information about opportunities after receiving a recommended match. Here, we benefit from the platform's ability to track whether a personalized recommendation resulted in that person clicking through to the platform to learn more about the recommended opportunity. Therefore, we define an indicator variable, *Market Participation*, as a positive indication that an individual clicked through to the platform from the initial email description to further review the job posting.

Market participation occurred for 8.7% of recipients across all treatment conditions. This is substantially higher than the 2–5% typical click-through rate for typical email campaigns (Campaign Monitor, 2025; Mailchimp, 2025; GetResponse 2025). This higher level likely reflects the targeted population, the exclusivity of the opportunity channel, and the credibility of the university sender. Importantly, these specialized project roles are not publicly advertised elsewhere, making the recommendation email the primary discovery mechanism and the platform the exclusive application

channel. Thus, the click-through decision represents not just interest but a necessary gateway to application-jobseekers who do not click through effectively forfeit the opportunity entirely. This makes our measure particularly consequential compared to contexts where alternative search channels exist.<sup>5</sup>

Our main explanatory variables are the experimental treatments, represented by indicator variables for *AI-Matching* and *Human-Matching*. The control group is implicitly estimated as the baseline level in the models that follow. Several other variables were used as regressors or as a basis for stratifying analyses in the data. In particular, we analyze heterogeneity by STEM versus non-STEM background. The construction of this measure, its technical details, and validation checks are described separately in Section 4.3.

Table 2 presents the balance of variables across AI, Human, and Control groups, suggesting a broadly successful balance of characteristics across treatments. To also assure that no small differences resulted in skewed results, all results were verified to be insensitive to the inclusion of any of these variables.

---

<sup>5</sup> We also observe tracking of email opening as a secondary metric, with 41.8% of recipients generating a positive opening signal in the platform's CRM system. This fraction is statistically equal across treatments. We do not condition our primary analysis on email opening for two reasons. First and most important, modern email clients allow message preview without triggering opening indicators, making this metric an imperfect measure of actual message exposure. Second, our experimental treatments may influence the decision to open emails themselves, making conditioning on this intermediate outcome inappropriate for causal inference. Nonetheless, robustness checks using participation conditional on observed email opening yield consistent results.

**Table 2 Balance of Characteristics Across Treatments and Control Group**

	<i>Control Group</i>	<i>AI Matching</i>	<i>Human Matching</i>
<i>Wage</i>	0.219 (0.011)	0.227 (0.011)	0.230 (0.011)
<i>Wage Missing</i>	0.174 (0.003)	0.172 (0.003)	0.175 (0.002)
<i>STEM</i>	0.632 (0.013)	0.619 (0.013)	0.632 (0.013)
<i>Masters</i>	0.124 (0.009)	0.155 (0.009)	0.140 (0.009)
<i>Doctoral</i>	0.032 (0.005)	0.029 (0.004)	0.033 (0.005)
<i>Ex-Ante Well-Matched [Master's in Rel. Fields]</i>	0.081 (0.007)	0.101 (0.008)	0.089 (0.007)
<i>Prof'l Network Size [No. Linkedin Connections]</i>	3.013 (0.118)	2.914 (0.102)	2.922 (0.100)
<i>log No. Local Competitors</i>	6.703 (0.020)	6.741 (0.019)	6.740 (0.019)
<i>Female</i>	0.481 (0.013)	0.464 (0.013)	0.453 (0.013)
<i>Grad Year</i>	2023.539 (0.016)	2023.518 (0.016)	2023.481 (0.015)

Notes. Number of observations = 1500, 1525, and 1537. Standard errors reported in parentheses.

### 4.3 Proxy for Variation in Knowledge & Understanding of Machine-based Prediction

Another measure deserving expanded discussion relates to heterogeneity in familiarity with AI and machine-based systems, as all main hypotheses (H1–H3) are tested conditional on this dimension (H4). We therefore require a measure of knowledge or technical familiarity.

We proxy this by classifying jobseekers into STEM versus non-STEM backgrounds, based on LinkedIn education records and standardized program codes following U.S. Citizenship and Immigration Services (2022). This classification is intended to capture differences in technical training that plausibly affect understanding of machine-based prediction and recommendations.

Of the 4,562 subjects, we are able to make a positive match between their degree studies and the STEM classifications in 4,412 cases, meaning that our stratified analyses of STEM and non-STEM subgroups will use 150 (3.3%) fewer observations.

A concern with this proxy is that it might capture not only differences in technical familiarity but also differences in underlying labor markets. We address this in two ways. First, because our keyword-based matching spanned a wide variety of fields, both STEM and non-STEM candidates appear within the applicant pool for most jobs. This limits the risk that results are driven solely by job composition. Second, for the handful of postings that are clearly more STEM-oriented than others, we test whether results are being driven by the STEM nature of the job itself rather than the STEM background of the candidate.

## 5 RESULTS

Building on the theoretical arguments in Section 3 and the experimental design in Section 4, we now turn to the evidence. We first report baseline treatment effects, then evaluate our main hypotheses. Subsections address H1, H2, and H3 individually, while H4 is assessed within each of these discussions.

### 5.1 Baseline Estimates: AI-Matching Treatment Effect

Here, we estimate the average overall effect of the AI treatment on market participation, finding a large negative effect on average. Results of OLS estimates are reported in Table 3, with robust standard errors in parentheses. In model (1), we regress Market Participation on AI-Matching and Human-Matching treatment indicators, along with a constant. The constant term reflects the baseline market participation rate within the control group. The coefficient on the AI-Matching treatment dummy in model (1) is negative and significant at  $p < 0.05$ , with a coefficient of  $-0.024$  (s.e. = 0.010), indicating a 2.4 percentage-point reduction relative to a baseline 9.4% participation rate in the control group. This represents a 25.5% relative reduction in participation. By contrast, the coefficient on the Human-Matching treatment indicator is statistically indistinguishable from zero, indicating average Market Participation is the same in the Human-Matching treatment as in the control group.

Given the randomized design, these estimates can be interpreted as causal effects. Adding further controls should not materially alter the estimates. Consistent with this expectation, adding a full set of fixed effects for each job posting, as in model (2), yields coefficients that are substantively unchanged.

To provide an additional gauge of magnitude, model (3) includes the randomized wage (with a dummy variable for cases in which wage was not disclosed). Based on coefficient estimates, the AI treatment effect is equivalent to roughly  $7.5\times$  the impact of a \$10 wage reduction, reinforcing our interpretation of the AI treatment as substantively large. Finally, model (4) re-estimates this same

model, but adds back the job-posting fixed effects, again leaving the AI coefficient effectively unchanged.

**Table 3 Baseline OLS Estimates of AI- and Human-Matching Disclosure on Market Participation**

Dep. Var.:	<i>Market Participation</i>				
	Model:	(1)	(2)	(3)	(4)
<i>AI Matching</i>	-0.025**	-0.024**	-0.024**	-0.024**	
	(0.010)	(0.010)	(0.010)	(0.010)	
<i>Human Matching</i>	0.000	0.001	0.000	0.000	
	(0.011)	(0.011)	(0.011)	(0.011)	
<i>Wage</i>			0.032**	0.025*	
			(0.014)	(0.014)	
<i>No Wage Shown</i>			0.113**	0.099*	
			(0.055)	(0.055)	
Constant	0.095***		0.069***		
	(0.01)		(0.01)		
Job Posting FE's		Y		Y	
<i>Adjusted-R<sup>2</sup></i>	0.001	0.011	0.002	0.012	
No. Obs.	4562	4562	4562	4562	

Notes. OLS estimates with robust standard errors in parentheses.  
 \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

## 5.2 Match Prediction Quality (H1)

We now turn to testing our main hypotheses (Section 3), beginning with H1. H1 predicts that if jobseekers expect AI-predicted matches to differ in accuracy—either better or worse—then ex-ante well-matched candidates will respond more strongly than other candidates, raising or lowering their willingness to participate accordingly (Section 3.1).

To test this idea, we require a clearly identifiable subset of jobseekers who are widely accepted to be ex-ante well-matched. In our setting, platform managers and employers identified the most relevant degree programs for each position. Anecdotally in this market, candidates with master’s-level training in highly relevant programs are deemed to more likely to be good matches: they more often tend to combine advanced training, prior industry experience, and flexible course loads conducive to short-term assignments. Using LinkedIn records, we constructed an indicator, *Ex-Ante Well-Matched*, for master’s-level candidates from programs flagged as highly relevant to each. This definition is narrower than the broader keyword matching relevance criteria used to form the overall risk set (Section 4.1).

Results are reported in Table 4. Model (1) regresses *Market Participation* on *Ex-Ante Well-Matched* with a constant; model (2) adds job-posting fixed effects. In both models, the coefficient on *Ex-Ante Well-Matched* is positive and significant. That they are more likely to participate, in general, is consistent with these individuals being better ex-ante matches. The coefficient of 0.056 in model (1) indicates that well-matched candidates were about 5.6 percentage points more likely to participate than others, relative to a baseline participation rate of 9.5%. This represents nearly a 60% relative increase.

Models (3) and (4) re-estimate the model for the control group; models (5) and (6) for the AI-Matching treatment; and models (7) and (8) for the Human-Matching treatment. We find that *Ex-Ante Well-Matched* is strongly and positively related to *Market Participation* only in the Human-Matching treatment, consistent with ex-ante well-matched candidates expecting higher match-quality recognition from human judgment. For example, in the Human-Matching treatment (model 7), the coefficient of 0.103 implies that ex-ante well-matched candidates were more than 10 percentage points more likely to participate than their peers, effectively doubling the baseline participation rate. In the AI-Matching treatment, the coefficient is smaller and statistically indistinguishable from zero. The difference in the *Ex-Ante Well-Matched* coefficient between the AI and Human conditions is significant at  $p < 0.05$ . The control-group estimate lies between those for AI and Human.

**Table 4 Market Participation by Candidates who are Ex-Ante Well-Matched, by Treatment**

Dep. Var.:	<i>Market Participation</i>							
	<i>All</i>		<i>Control Group</i>		<i>AI Matching</i>		<i>Human Matching</i>	
	Model: (1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Ex-Ante Well-Matched</i>	0.056***	0.053***	0.039	0.046	0.030	0.015	0.103***	0.104***
<i>[Master's in Rel. Fields]</i>	(0.018)	(0.019)	(0.032)	(0.032)	(0.025)	(0.028)	(0.034)	(0.037)
Constant	0.053***		0.046***		0.048***		0.030**	
	(0.008)		(0.014)		(0.011)		(0.015)	
Job Posting FE's		Y		Y		Y		Y
<i>Adjusted-R<sup>2</sup></i>	0.003	0.012	0.00	0.02	0.001	0.013	0.009	0.011
No. Obs.	4562	4562	1500	1500	1525	1525	1537	1537

Notes. OLS estimates with robust standard errors in parentheses.  
 \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

As an additional investigation of H1 and particularly to assess evidence with respect to H4, we test whether effects are most present among those with technical backgrounds who at least today might better understand the workings of machine-based prediction. Results stratifying those with STEM and non-STEM backgrounds are reported in Table 5.

Models (1)-(4) are consistent with the earlier results and also with H4, as we see similar patterns as those above, which are highly significant in the case of the Human-Matching treatment. We predict that patterns in relation to models (5)-(8) for those without technical backgrounds might not conform as well to the patterns predicted by H1. The patterns presented are consistent with this idea, inasmuch as the coefficient on *Ex-Ante Well-Matched* is not significant in the Human-Matching treatment group. However, we caution this result is difficult to interpret. This lack of significance could in part reflect lesser conformance to the predictions of H1, but it could also reflect simply that these coefficients are estimated with less statistical power, leading to larger standard errors and less precisely estimated point estimates.<sup>6</sup> Therefore, the significant results among those with STEM backgrounds affirm H1 and H4, but we caution in interpreting the results for non-STEM candidates.

**Table 5 Market Participation by Candidates who are Ex-Ante Well-Matched, by Treatment, Stratified by those with and without Technical Backgrounds**

Dep. Var.:	<i>Market Participation</i>							
	<i>STEM = 1</i>				<i>STEM = 0</i>			
Subsample:	<i>All</i>	<i>Control Group</i>	<i>AI Matching</i>	<i>Human Matching</i>	<i>All</i>	<i>Control Group</i>	<i>AI Matching</i>	<i>Human Matching</i>
Treatment:	<i>All</i>	<i>Control Group</i>	<i>AI Matching</i>	<i>Human Matching</i>	<i>All</i>	<i>Control Group</i>	<i>AI Matching</i>	<i>Human Matching</i>
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Ex-Ante Well-Matched</i>	0.048**	0.043	0.008	0.101***	0.106	0.046	0.182	0.112
<i>[Master's in Rel. Fields]</i>	(0.019)	(0.033)	(0.029)	(0.038)	(0.104)	(0.156)	(0.217)	(0.185)
Job Posting FE's	Y	Y	Y	Y	Y	Y	Y	Y
<i>Adjusted-R<sup>2</sup></i>	0.02	0.027	0.026	0.02	0.005	0.001	0.003	-0.007
No. Obs.	2769	918	910	941	1643	535	559	549

Notes. OLS estimates with robust standard errors in parentheses.  
\* p < 0.10, \*\* p < 0.05, \*\*\* p < 0.01

### 5.3 Sources and Inputs of Match Predictions (H2)

We now test H2, which concerns expectations about **what inputs** are used to generate match predictions. H2 predicts that jobseekers who have invested heavily in human-oriented matching assets—especially professional networks—will be differentially less likely to participate when recommendations are disclosed as AI-generated, anticipating that machine-based evaluation does not fully leverage human-centered assets.

<sup>6</sup> Among the 2,769 STEM candidates, 1,339 were from directly relevant fields and 398 had master's-level training in those fields. Among the 1,643 non-STEM candidates, 447 were from directly relevant fields and only 15 had master's-level training. Dropping job-posting fixed effects to gain degrees of freedom leaves results unchanged. Redefining Ex-Ante Well-Matched to include all individuals with relevant training (not limited to master's-level) lowers standard errors but still yields statistically null estimates.

Results are reported in Table 6. Model (1) regresses *Market Participation* on the size of professional network. This is measured as the number of LinkedIn connections, expressed in 100s, to make the coefficients easily legible. Model (2) adds job-posting fixed effects. In both, the coefficient on network size is positive and significant, consistent with those with larger networks being more active jobseekers or higher-quality applicants.

Our main test of H2 is to examine whether those with large professional networks are more or less likely to participate under the AI-Matching treatment. Models (3)-(9) re-estimate models (1) and (2), but stratified by control group, AI-Matching treatment, and Human-Matching treatment. Broadly, we see that the coefficient estimates in the AI treatment are weakest and not significantly distinguishable from zero, consistent with those with large (human) professional networks anticipating lower returns to these assets in the AI-Matching treatment. However, while the coefficients on the professional network variable across the control group, models (3) and (4), and the Human-Matching treatment, models (5) and (6), are larger (more than 2×) in each case and more statistically significant, only that estimated in model (3) is significant at  $p < 0.05$ . Further, the differences in estimated coefficients on professional network size across models (3) and (5) is not statistical at conventional levels, nor is the difference across models (4) and (6).

These results are consistent with those with large prior investments in (human) professional networks being differentially less likely to participate, anticipating the lower returns to these assets for postings mediated by AI predictions, supporting H2. For example, in the Human-Matching treatment, a coefficient of about 0.004 translates into a 0.4 percentage-point increase per 100 connections—roughly 4% of baseline—while the corresponding estimate in AI-Matching is less than half and statistically zero.

**Table 6 Market Participation by Candidates with Large (Human) Professional Networks, by Treatment**

Dep. Var.:	<i>Market Participation</i>								
	Treatment:	<i>All</i>		<i>Control Group</i>		<i>AI Matching</i>		<i>Human Matching</i>	
	Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Prof'l Network Size [No. LinkedIn Connections]</i>	0.004*** (0.001)	0.003** (0.001)	0.005** (0.002)	0.004* (0.002)	0.002 (0.002)	0.001 (0.002)	0.004* (0.002)	0.004 (0.002)	
Constant	0.076*** (0.005)		0.079*** (0.010)		0.065*** (0.008)		0.083*** (0.009)		
Job Posting FE's		Y		Y		Y		Y	
<i>Adjusted-R<sup>2</sup></i>	0.003	0.012	0.006	0.018	0.003	0.013	0.003	0.005	
No. Obs.	4562	4562	1500	1500	1525	1525	1537	1537	

Notes. OLS estimates with robust standard errors in parentheses. No. of LinkedIn Connections is in 100s. \* p < 0.10, \*\* p < 0.05, \*\*\* p < 0.01

To more closely test H2 and to assess whether evidence is consistent with H4, we proceed by stratifying results by those with technical versus non-technical backgrounds in Table 7. Model (1) confirms that for candidates with STEM backgrounds that the relationship between *Market Participation* and size of professional network is positive. Models (2), (3), and (4) then compare results for the control group, the AI-Matching treatment, and the Human-Matching treatment. Again, consistent with H2, we see positive and statistically significant coefficients on the professional network variable in the case of the control group and Human-Matching, but not for AI-Matching. The magnitude of effect for the control group is between those of the Human and AI-Matching treatments, leading that estimate to be less significant than the estimate for the Human-Matching treatment.

Consistent with H4, the magnitudes and significance of estimates for the Human treatment group and control group are higher for the subset of our study population with STEM backgrounds. Further, the difference of coefficient estimates between Human-Matching and AI-Matching becomes marginally statistically significant among those with STEM backgrounds, where this difference is significant at p < 0.10.

Model (5) estimates the model for non-STEM individuals and finds no significant relationship. The analogous stratification by treatment in the case of non-STEM individuals finds insignificant coefficients, each much closer to zero, in each case. The results are considerably stronger among those with technical backgrounds, consistent with H4.

**Table 7 Effect of AI Disclosure on Market Participation for Jobseekers with Large Professional Networks**

Dep. Var.:	<i>Market Participation</i>							
Subsample:								
Treatment:	<i>All</i>	<i>Control Group</i>	<i>AI Matching</i>	<i>Human Matching</i>	<i>All</i>	<i>Control Group</i>	<i>AI Matching</i>	<i>Human Matching</i>
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Prof'l Network Size [No. LinkedIn Connections]</i>	0.005** (0.002)	0.006* (0.003)	0.001 (0.002)	0.007** (0.003)	0.001 (0.002)	0.002 (0.003)	0.000 (0.004)	0.000 (0.003)
Job Posting FE's	Y	Y	Y	Y	Y	Y	Y	Y
<i>Adjusted-R<sup>2</sup></i>	0.02	0.02	0.026	0.027	0.004	0.001	0.001	0.008
No. Obs.	2769	918	910	941	1643	535	559	549

Notes. OLS estimates with robust standard errors in parentheses.  
 \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

### 5.4 Scale of Competition (H3)

We turn now to H3, which predicts that under the AI-Matching treatment jobseekers' willingness to participate will be less sensitive to the scale of local competition. Because this analysis relies on variation in the number of competitors across postings (rather than experimental assignment), we proceed in steps and interpret results with caution.

#### 5.4.1 Negative Relationship between Participation and Competition

Model (1) of Table 8 estimates the baseline association between *Market Participation* and the log of the number of local competitors in the full sample. The coefficient is negative, consistent with standard crowding effects in rank-order contests. However, to attempt to detect whether this pattern is the result of spurious effects, models (2)-(6) progressively add controls for education, graduation year, demographics, and wages; across specifications, the competition coefficient remains negative and remarkably stable. Stratifying by STEM and non-STEM, in models (7) and (8), also yields this same negative pattern.

**Table 8 Negative Relationship between Market Participation and Number of Local Competitors**

Dep. Var.:	<i>Market Participation</i>						
	<i>All</i>					<i>STEM = 1</i>	<i>STEM = 0</i>
	Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>log No. Local Competitors</i>	-0.020*** (0.007)	-0.023*** (0.007)	-0.023*** (0.007)	-0.023*** (0.007)	-0.023*** (0.007)	-0.023*** (0.007)	-0.022*** (0.007)
<i>Masters</i>	0.076*** (0.015)	0.082*** (0.015)	0.086*** (0.015)	0.086*** (0.015)	0.087*** (0.015)	0.086*** (0.015)	0.085*** (0.017)
<i>Doctoral</i>	-0.048*** (0.015)	-0.051*** (0.015)	-0.044*** (0.015)	-0.044*** (0.015)	-0.044*** (0.015)	-0.044*** (0.015)	-0.036 (0.023)
<i>Graduation Year</i>		0.012* (0.006)					
Constant	0.223*** (0.045)	0.232*** (0.045)					
Grad Year Dummies			Y	Y	Y	Y	Y
Gender Dummy				Y	Y	Y	Y
Wage Controls					Y	Y	Y
<i>Adjusted-R<sup>2</sup></i>	0.00	0.02	0.02	0.03	0.03	0.00	0.00
No. Obs.	4562	4562	4562	4562	4562	2769	1643

Notes. OLS estimates with robust standard errors in parentheses.  
 \* p < 0.10, \*\* p < 0.05, \*\*\* p < 0.01

### 5.4.2 Treatment Differences

Table 9 reports our tests of H3, where we are most interested in whether there is a weaker relationship in the AI-Matching treatment between *Market Participation* and number of competitors in the local market. For comparison's sake, we first report the relationship between willingness to participate and the scale of the market for the entire study population for both the uncontrolled regression in model (1) and the fully saturated model (i.e., with the full set of controls from the preceding analysis) in model (2). Models (3) and (4) then report results for just the control group, where the negative effect of competition is even stronger.

We report the relationship in the AI-Matching treatment in models (5) and (6). We find indeed that the coefficient is weaker, indistinguishable from zero, and more than an order of magnitude smaller than the relationship estimated in the Control group. Thus, consistent with H3, the number of competitors in the local market is no longer a predictor for willingness to participate in the AI-

Matching treatment. Comparing either the coefficient estimate of model (6) with that of model (3) or that of model (5) with (4), we find the difference in these estimates to be significant at  $p < 0.10$ .<sup>7</sup>

In models (7) and (8), we find that results for the Human-Matching treatment are also weaker than in the control group. We speculate that this pattern is plausibly consistent with subjects in the Human-Matching treatment forming expectations that Human-Matching might lead to matching that does not correlate with the full scale of the local market of direct competitors, perhaps limited by scale of information processing.<sup>8</sup>

**Table 9 Market Participation in Relation to Scale of Local Competition, by Treatment**

Dep. Var.:	<i>Market Participation</i>							
	<i>All</i>		<i>Control Group</i>		<i>AI Matching</i>		<i>Human Matching</i>	
Treatment:								
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>log No. Local Competitors</i>	-0.020*** (0.007)	-0.024*** (0.007)	-0.037*** (0.012)	-0.041*** (0.012)	-0.009 (0.011)	-0.012 (0.011)	-0.013 (0.012)	-0.017 (0.012)
Constant	0.223*** (0.045)		0.341*** (0.080)		0.131* (0.074)		0.182** (0.079)	
Full Set of Controls	Y		Y		Y		Y	
<i>Adjusted-R<sup>2</sup></i>	0.003	0.018	0.008	0.026	0.000	0.008	0.000	0.015
No. Obs.	4562	4562	1500	1500	1525	1525	1537	1537

Notes. OLS estimates with robust standard errors in parentheses.  
\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

To evaluate H4, we repeat the analysis by technical background, as reported in Table 10. For both STEM and non-STEM subgroups, patterns broadly mirror the full sample: a strong negative association with competition in Control, attenuated toward zero under AI-Matching, and intermediate under Human-Matching. We do not find significant differences between STEM and non-STEM in the competition-response relationship.

<sup>7</sup> We find evidence consistent with these results in a post-experiment survey for which a small number of participants responded ( $N = 79$ , evenly divided across the two treatments and the control group). Consistent with this result and H3, those in the AI-Matching treatment reported they faced over 700 more competitors on average than those in the Human-Matching treatment. The survey also found those in the AI treatment reported a lower expected likelihood of winning the job and a lower expectation of the quality of the match.

<sup>8</sup> As footnoted earlier, a small post-experimental survey found that, on average, respondents from the Human-Matching treatment expected to face fewer competitors than those in the AI treatment.

**Table 10 Market Participation in Relation to Scale of Local Competition, by Treatment, Stratified by those with or without Technical Backgrounds**

Dep. Var.:	<i>Market Participation</i>							
	<i>STEM = 1</i>				<i>STEM = 0</i>			
Subsample:	<i>All</i>	<i>Control Group</i>	<i>AI Matching</i>	<i>Human Matching</i>	<i>All</i>	<i>Control Group</i>	<i>AI Matching</i>	<i>Human Matching</i>
Treatment:								
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>log No. Local Competitors</i>	-0.027*** (0.010)	-0.044*** (0.017)	-0.023 (0.017)	-0.012 (0.017)	-0.022** (0.010)	-0.035* (0.019)	-0.002 (0.014)	-0.026 (0.020)
Job Posting FE's	Y	Y	Y	Y	Y	Y	Y	Y
<i>Adjusted-R<sup>2</sup></i>	0.016	0.025	0.015	0.017	0.010	0.007	0.006	0.002
No. Obs.	2,769	918	910	941	1,643	535	559	549

Notes. OLS estimates with robust standard errors in parentheses.  
 \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

### 5.5 Overall Effects of the AI Treatment on STEM versus non-STEM Candidates

The preceding analyses and our earlier theorizing in Section 3.4, suggested possible differences in responses between those with technical backgrounds versus those without such backgrounds. Given these differences, we return to the overall AI-Matching treatment effect, and examine each group separately, dividing those with versus without STEM backgrounds. Given the earlier estimation of a significant and rather large negative effect on the AI-Matching treatment (Section 5.1), we focus attention on this coefficient.

Results are reported in Table 11. For those with STEM backgrounds, model (1) regresses *Market Participation* on the treatment dummies and finds no statistical difference with the Control group. Model (2) adds job posting fixed effects, in case it bears on the size of standard errors. We continue to find an insignificant coefficient on the AI-Matching treatment dummy. In each case, the coefficients on the AI-Matching treatment dummy are negative, but small.

Repeating these same estimates for those without STEM backgrounds, as in models (3) and (4), by contrast, finds large negative effects of the AI-Matching treatment. Therefore, it appears that the earlier negative effect documented in Section 5.1 is largely driven by those without technical backgrounds.

Repeating these same estimates for those without STEM backgrounds, as in models (3) and (4), by contrast, finds large negative effects of the AI-Matching treatment. Therefore, it appears that the earlier negative effect documented in Section 5.1 is largely driven by those without technical backgrounds.

Therefore, much of the overall negative AI-Matching treatment effect (Section 4.1) is driven by those without technical training. For those without technical training, the AI-Matching treatment dropped participation by about 4 percentage points, or 40% relative to the ~9.9% baseline participation rate in the Control group. The point estimate for those with technical backgrounds indicates a 17% drop relative to the ~9.3% baseline participation rate in the Control group, by comparison. To emphasize, although we find large statistical effects on the composition of market participation among those with technical backgrounds (e.g., Section 5.2 and 5.3), this 17% drop in levels of participation is not statistically significant at conventional levels.

**Table 11 Estimates of AI-Matching Treatment Effect, Re-Estimated by those with and without Technical Backgrounds**

Dep. Var.:	<i>Market Participation</i>			
	<i>STEM = 1</i>		<i>STEM = 0</i>	
Subsample:	(1)	(2)	(3)	(4)
Model:	(1)	(2)	(3)	(4)
<i>AI Matching</i>	-0.016 (0.013)	-0.014 (0.013)	-0.040** (0.016)	-0.040** (0.016)
<i>Human Matching</i>	0.009 (0.014)	0.010 (0.014)	-0.013 (0.018)	-0.014 (0.018)
Constant	0.093*** (0.010)		0.099*** (0.013)	
Job Posting FE's		Y		Y
<i>Adjusted-R<sup>2</sup></i>	0.001	0.019	0.003	0.007
No. Obs.	2769	2769	1643	1643

Notes. OLS estimates with robust standard errors in parentheses.  
\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

## 6 SUMMARY AND CONCLUSION

Much prior research has found examples of people altering their behavior and choices when interacting with machine-based intelligence—including both AI “aversion” and “appreciation”—with many plausible theories and hypotheses for why (Section 2.2).

The central contribution of this study is to document such evidence in job markets, where we assessed whether AI use and disclosure altered market participation by jobseekers. In making new progress in this context, our primary aim and contribution was to seek discriminating evidence for what is arguably the simplest explanation: that the use and disclosure of AI alters jobseekers’ expected payoffs from pursuing and competing for a job, as AI use might alter the nature of competition. To probe this broader idea, we theorized three distinct channels through which AI use and disclosure could plausibly influence behavior, related to (i) the quality of match predictions, (ii) the kinds of

inputs and information used in the process, and (iii) the perceived scale of competition. We formulated falsifiable empirical hypotheses to test these ideas, allowing for the possibility that the direction of effects for (i) and (iii) could go in either direction. For example, the quality of prediction and scale of competition could plausibly either increase or decrease with AI use. Further, given the many theories and hypothesized explanations for changing behavior in this context, we predicted that (iv) our predictions related to the three channels would be most apparent among those with the greatest understanding and experience with machine-based predictions, which we roughly proxied with the distinction between those with technical backgrounds in STEM versus those without.

Our analysis revealed systematic shifts in market participation choices in relation to the three predicted channels, particularly among those with technical STEM backgrounds. For example, ex-ante well-matched candidates were more than 10 percentage points more likely to participate under human recommendations but showed no such premium under AI—a swing larger than the baseline participation rate itself. Jobseekers with larger professional networks saw about a 0.4 percentage-point increase in participation per 100 LinkedIn connections in human or control conditions, but no such gains under AI. And whereas participation normally fell as local competitor counts increased, this relationship all but disappeared under AI disclosure. These systematic magnitudes—each expressed relative to a  $\sim 9\frac{1}{2}$  % baseline—were most pronounced among technically trained participants, who are best positioned to anticipate how machine-based prediction reshapes their expected payoffs. Therefore, the general patterns, and particularly those related to STEM individuals responding to AI use and disclosure, conformed with our series of predictions, consistent with changing behavior in relation to how AI is expected to shift payoffs from pursuing and competing for a job.

At the same time, we predicted that non-STEM individuals would be less likely to conform to our predictions related to shifts in expected payoffs, since they could possess less understanding of machine-based prediction than those with technical backgrounds. Indeed, we found several instances in which non-STEM subjects appear to respond differently to the treatments (Sections 5.2–5.5). We caution that these differences in response among those with less technical knowledge of machine-based prediction could have multiple alternative interpretations (especially see Sections 5.2 and 5.4). Nonetheless, the results are consistent with prior research indicating that many plausible mechanisms related to subjective attitudes underlie AI “aversion,” especially for those with less technical familiarity. However, relative to the aims and focus of this study, the clearest systematic evidence aligns with payoff-based mechanisms

Most striking of all in the differences in responses to the AI-Matching treatment between STEM and non-STEM jobseekers was the overall treatment effect on market participation. Whereas the overall average effect in our study population was a  $-2.4$  percentage point reduction in participation (a 25% decline relative to the 9.5% baseline participation rate)—roughly 7.5 times the impact of a \$10 wage reduction—much of this effect was driven by those without technical training. Among non-STEM jobseekers the overall AI-Matching treatment effect was to reduce market participation by about **4 percentage points**, or about a 40% drop relative to the  $\sim 9.9\%$  baseline participation rate in the Control group. The point estimate of the AI-Matching treatment effect among STEM jobseekers was only roughly one-third the magnitude of that among non-STEM jobseekers—about a 17% drop relative to the  $\sim 9.3\%$  baseline participation rate in the Control group—and statistically indistinguishable from zero.

Broadly, we interpret the sum of these patterns as supporting our general idea that (informed) expectations of payoffs from AI can shift behaviors. Further, the adherence of those with technical backgrounds to predictions of responding to reasoned payoffs leads us also to speculate that greater knowledge and understanding across individuals (and perhaps over time) could lead the expected payoffs rationale to create less ambit for subjective attitudes towards AI to play a role.

Therefore, these findings deepen our understanding of how AI-mediated matching systems reshape labor markets. Beyond reducing average participation, AI disclosure can alter *who* chooses to compete, potentially affecting efficiency, equity, and inclusion in the allocation of human capital (Mortensen 1986; Pissarides 2000). Because some responses appear consistent with rational expectations about structural shifts in competition, these effects are unlikely to dissipate over time, unlike attitudinal aversion documented in prior human–AI interaction studies (Dietvorst et al., 2015; Longoni et al., 2019). They may even intensify as jobseekers become more familiar with AI systems. At the same time, because expectations are shaped by firms’ strategies and technologies, different designs could elicit very different responses. For instance, while in our context informed jobseekers anticipated lower match quality, diminished returns to networks, and broader competition, alternative designs—such as credibly narrower targeting or demonstrably superior prediction quality—might generate the opposite expectations.

Our study makes three main contributions. First, in the literature on labor market digitization, we extend prior work on market thickening, search costs, and platform reach (Horton 2017; Goldfarb & Tucker 2019) by showing that AI disclosure can strategically influence not only the level but also the composition of participation. Second, in research on platform governance and algorithmic

management, we identify mechanisms—perceived match quality, reliance on codified versus tacit inputs, and scale of competition—through which disclosure shapes user behavior, advancing understanding of transparency, trust, and signaling in algorithmic platforms (Lee 2018; Castelo, Bos, & Lehmann 2019; Burton et al. 2020). Third, in the broader human–AI interaction domain, we contribute a rare large-scale field experiment in a high-stakes labor market, moving beyond documenting broad aversion to provide discriminating evidence of underlying payoff-based mechanisms (Luo et al. 2019; Tong et al. 2021; Wiles et al. 2023).

Taken together, these contributions highlight how AI disclosure affects both the participation margin and its composition in job markets. They also underscore that the behavioral consequences of AI depend not only on attitudes toward the technology but also on the structural payoffs it creates—features that firms and policymakers can influence through platform design and governance (Kim et al. 2024).

This study also has limitations that point to directions for future research. While we emphasize payoff-based mechanisms, the broader human–AI interaction literature makes clear that responses to AI disclosure often involve subjective attitudes related to many factors, such as perceived fairness, legitimacy, or trust. Several of our findings—especially among non-STEM jobseekers—may reflect these and deserve further careful study. Further, our design captures only a single decision point, a single experiment in time. It remains to be demonstrated how responses might change with repeated exposure or over longer time horizons; longitudinal or repeated-exposure studies could help clarify whether learning leads to greater reliance on payoff-based reasoning. We suspect there will be greater adherence to the mechanisms studied here, over time. Moreover, our analysis focuses on the margin of participation—whether jobseekers enter at all—while other margins may also be important, there may be other decisions and other margins worth investigating. For example, this might include how much effort or investment candidates devote to a given opportunity, or how they adapt to perceived changes in competition over time with AI. Future work might also study attempts at remedies, such as alternative disclosure strategies or system designs, and examine how expectations and behaviors shift in equilibrium as both jobseekers and employers adjust. Finally, our context—a high-skill, university-based job platform—provides a useful but bounded setting; replications in other types of labor markets would help test the scope and generality of these findings.

In conclusion, our study reveals that the integration of AI into job matching processes is not merely a technological shift and advance in information-processing, but one that fundamentally alters the behavior and decision-making of jobseekers. As AI continues to transform labor markets,

understanding and addressing these behavioral responses will be crucial for realizing the potential benefits of AI-driven matching while mitigating unintended negative consequences.

## REFERENCES

- Albright, A. (2025). *The hidden effects of algorithmic recommendations* (Institute Working Paper No. 104). Federal Reserve Bank of Minneapolis.
- Arnold, D., Dobbie, W., & Yang, C. (2018). Racial bias in bail decisions. *Quarterly Journal of Economics*, 133(4), 1885–1932.
- Autor, D. H. (2001). Wiring the labor market. *Journal of Economic Perspectives*, 15(1), 25–40.
- Avery, M., Leibbrandt, A., & Vecci, J. (2024). Does artificial intelligence help or hurt gender diversity? Evidence from two field experiments on recruitment in tech. *Working paper*.
- Baker, G. P. (1992). Incentive contracts and performance measurement. *Journal of Political Economy*, 100(3), 598–614.
- Bansak, K., Hainmueller, J., & Hangartner, D. (2016). How economic, humanitarian, and religious concerns shape European attitudes toward asylum seekers. *Science*, 354(6309), 217–222.
- Bojd, B., Garimella, A., & Yin, H. (2024). Overcoming the stigma barrier: Conversational information-seeking from AI chatbots vs. humans. *Working paper*.
- Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton.
- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108, 43–47.
- Burks, S. V., Cowgill, B., Hoffman, M., & Housman, M. (2015). The value of hiring through employee referrals. *Quarterly Journal of Economics*, 130(2), 805–839.
- Burton, J. W., Stein, M. K., & Jensen, T. B. (2020). A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making*, 33(2), 220–239.
- Cabral, L., & Hortaçsu, A. (2010). The dynamics of seller reputation: Evidence from eBay. *Journal of Industrial Economics*, 58(1), 54–78.
- Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research*, 56(5), 809–825.
- Cowgill, B., Dell’Acqua, F., & Deng, S. (2021). Causal machine learning and algorithmic management: Experiments and theory. *Working paper*.
- Dellarocas, C. (2003). The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49(10), 1407–1424.
- Diederich, S., Brendel, A. B., & Lichtenberg, S. (2022). Algorithm aversion in information systems research: A comprehensive literature review. *Business & Information Systems Engineering*, 64(1), 47–64.
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126.
- Fradkin, A., Bhole, M., & Horton, J. J. (2025). Competition avoidance vs. herding in job search: Evidence from large-scale field experiments on an online job board. *Management Science, Articles in Advance*.
- Fügener, A., Grahl, J., Gupta, A. and Ketter, W., 2021. Will humans-in-the-loop become borgs? Merits and pitfalls of working with AI. *MIS Quarterly*, 45(3), pp.1527-1556.
- Fuller, J. B., & Raman, M. (2021). *Hidden workers: Untapped talent*. Harvard Business School / Accenture.
- Gee, L. K. (2019). The more you know: Information effects on job application rates in a large field experiment. *Management Science*, 65(5), 2077–2094.
- Gibbons, R., & Waldman, M. (2004). Task-specific human capital. *American Economic Review*, 94(2), 203–207.
- Goldfarb, A., & Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1), 3–43.

- Granovetter, M. (1995). *Getting a job: A study of contacts and careers* (2nd ed.). University of Chicago Press.
- Green, B., & Chen, Y. (2019). Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 90–99). ACM.
- Haesevoets, T., De Cremer, D., Dierckx, K., & Van Hiel, A. (2021). Human–machine collaboration in managerial decision making. *Computers in Human Behavior*, *119*, 106730.
- Hoffman, M., Kahn, L. B., & Li, D. (2018). Discretion in hiring. *Quarterly Journal of Economics*, *133*(2), 765–800.
- Horton, J. J. (2017). The effects of algorithmic labor market recommendations: Evidence from a field experiment. *Journal of Labor Economics*, *35*(2), 345–385.
- Jussupow, E., Spohrer, K., Heinzl, A., & Gawlitza, J. (2023). A review of algorithm aversion in human–AI interaction. *Information Systems Frontiers*, *25*(3), 755–776.
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, *14*(1), 366–410.
- Kim, H., Glaeser, E. L., Hillis, A., Kominers, S. D., & Luca, M. (2024). Decision authority and the returns to algorithms. *Strategic Management Journal*, *45*(4), 619–648.
- Kuhn, P. and Mansour, H. (2014) 'Is internet job search still ineffective?', *The Economic Journal*, *124*(581), pp. 1213-1233.
- Lazear, E. P., & Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, *89*(5), 841–864.
- Lee, M. K. (2018). Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society*, *5*(1), 1–16.
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, *151*, 90–103.
- Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research*, *46*(4), 629–650.
- Luo, T., Zhang, J., & Mu, X. (2025). Robo-interviews and jobseeker responses: Evidence from a survey-based experiment. *Working paper*.
- Luo, X., Tong, S., Fang, Z., & Qu, Z. (2019). Machines vs. humans: The impact of AI chatbot disclosure on customer purchases. *Marketing Science*, *38*(6), 937–947.
- Mahmud, A., Shafiq, A., & Haider, M. (2020). Algorithm aversion: A systematic literature review. *Journal of Behavioral Sciences*, *30*(1), 1–21.
- Mayer-Schönberger, V., & Cukier, K. (2013). *Big data: A revolution that will transform how we live, work, and think*. Houghton Mifflin Harcourt.
- Mortensen, D. T. (1986). Job search and labor market analysis. In O. Ashenfelter & R. Layard (Eds.), *Handbook of Labor Economics* (Vol. 2). North-Holland.
- Narayanan, A., & Kapoor, S. (2024). *AI snake oil: What artificial intelligence can do, what it can't, and how to tell the difference*. Princeton University Press.
- Oyer, P., & Schaefer, S. (2011). Personnel economics: Hiring and incentives. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (Vol. 4B, pp. 1769–1823). Elsevier.
- Pissarides, C. A. (2000). *Equilibrium unemployment theory* (2nd ed.). MIT Press.
- Prahl, A., & Van Swol, L. (2017). Understanding algorithm aversion: When is advice from automation discounted? *Journal of Forecasting*, *36*(6), 691–702.
- Schellmann, H. (2023, November 4). AI may not steal your job, but it could stop you getting hired. *Wired*.
- Sinclair-Desgagné, B. (2024). On the (non-)reliance on algorithms—A decision-theoretic account. *Journal of Mathematical Economics*, *109*, 102001.

- Starke, C., Baleis, J., Keller, B., & Marcinkowski, F. (2022). Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature. *Big Data & Society*, 9(2), 1–16.
- Stelmaszak, M., Möhlmann, M. and Sørensen, C., 2025. When algorithms delegate to humans: Exploring human-algorithm interaction at Uber. *MIS Quarterly*, 49(1), pp.305-330.
- Tambe, P., Cappelli, P., & Yakubovich, V. (2019). Artificial intelligence in human resource management: Challenges and a path forward. *California Management Review*, 61(4), 15–42.
- The Economist. (2023, May 7). Your job is (probably) safe from artificial intelligence. *The Economist*.
- Tong, S., Jia, N., Luo, X., & Fang, Z. (2021). The Janus face of artificial intelligence feedback: Deployment versus disclosure effects on employee performance. *Strategic Management Journal*, 42(9), 1600–1631.
- Turel, O., and Kalhan, S. 2023. “Prejudiced Against the Machine? Implicit Associations and the Transience of Algorithm Aversion,” *MIS Quarterly* (47:4), pp. 1369–1394.
- U.S. Citizenship and Immigration Services. (2022). *STEM designated degree program list*. U.S. Department of Homeland Security.
- Waldman, M. (1984). Job assignments, signalling, and efficiency. *RAND Journal of Economics*, 15(2), 255–267.
- Wiles, E., & Horton, J. J. (2025). Generative AI and labor market matching efficiency. *Working paper*.
- Wiles, E., Munyikwa, Z., & Horton, J. J. (2023). Algorithmic writing assistance on jobseekers’ résumés increases hires (NBER Working Paper No. 30886). National Bureau of Economic Research.
- Yeomans, M., Shah, A., Mullainathan, S., & Kleinberg, J. (2019). Making sense of recommendations. *Journal of Behavioral Decision Making*, 32(4), 403–416.
- Zhang, S. & P. Kuhn, "Measuring Bias in Job Recommender Systems: Auditing the Algorithms," NBER Working Paper 32889 (2024), <https://doi.org/10.3386/w32889>.