Employer and Employee Responses to Generative AI: Early Evidence*

Philip G. Berger
University of Chicago

Wei Cai
Columbia University

Lin Qiu
Purdue University

Cindy Xinyi Shen
Stanford University

First Version: February 2024
This Version: September 2024

ABSTRACT

This paper examines the responses of both employers and employees to the release of Generative AI. We leverage high-frequency data that directly tracks real-time employer demand and construct a forward-looking measure of exposure to Generative AI at the occupation level. Contrary to the conventional view of automation uniformly displacing workers, we find a *heterogeneous* effect: Generative AI complements high-level white-collar jobs (e.g., executive positions) but substitutes low-level white-collar jobs (e.g., entry-level office positions). We further find that firms with greater exposure to Generative AI significantly increase emphasis on Generative AI and machine learning skills in job listings. Our results suggest that heightened exposure to Generative AI does not inherently result in widespread job displacement but could instead drive higher demand for high-skilled labor and management roles. Moreover, we provide novel evidence on employee reactions and find a notable decline in the long-term outlook, although current employee ratings of firms remain stable. Furthermore, potential employees seek fewer interviews, indicating a shift in hiring dynamics. Our findings reveal a mismatch between how employers and employees perceive and react to the rapid advancement of Generative AI in the workplace.

JEL classification: G32, G24, G28, J48, J20, M12, M14.
Keywords: Generative AI, Employer and Employee Response, Labor Market.

*We thank Jung Ho Choi, Rebecca Hann, Valeri Nikolaev, and Yuqing Zhou for their helpful comments and suggestions. Berger, University of Chicago (email: Philip.Berger@chicagobooth.edu), Cai, Columbia University (e-mail: wc2419@gsb.columbia.edu), Qiu, Purdue University (email: qiu223@purdue.edu), Shen, Stanford University (email: xshen99@stanford.edu).

1 Introduction

On November 30, 2022, OpenAI unveiled ChatGPT, a groundbreaking achievement in artificial intelligence that has quickly become a focal point in discussions about the future of work. Powered by a large language model (LLM), ChatGPT performs a wide range of tasks through prompt engineering, from coding and debugging to composing music, stories, and essays.[1] This versatility, combined with its rapid adoption, has positioned ChatGPT as a key driver of the so-called Fourth Industrial Revolution, reshaping the labor market in profound ways. Even OpenAI and its competitors were surprised by the speed of ChatGPT's rise in terms of both popularity and capability.[2] What sets this generation of Generative AI apart from its predecessors is its proficiency in automating office jobs, which were previously considered immune to technological disruption due to their reliance on high levels of education and skill. However, the advent of Generative AI exposes various white-collar jobs to a labor market shock. Despite its potential to massively influence the workplace, it remains unclear how employers and employees are reacting to the rapid advancement of Generative AI technologies.

In this paper, we examine the responses of both employers and employees to Generative AI. We leverage high-frequency job posting data from LinkUp, which directly sources daily job listings from over 60,000 employer websites. This allows us to track real-time employer demand responses based on actual job openings rather than self-reported, aggregated data from different job boards (e.g., LinkedIn). Previous studies on technological advancements often face challenges in isolating demand shifts from supply-side changes, as well as tracking gradual technological advancement over time. Unlike previous generations of automation, the release of ChatGPT is a key milestone that was somewhat unexpected, which serves as a natural shock to the labor market. We use high-frequency job posting data to identify rapid shifts in employer demand triggered by the disruption of Generative AI. To capture the effects of Generative AI beyond the aggregated industry level,

[1] By January 2023, ChatGPT successfully passed MBA exam given by a Wharton professor (Rosenblatt, 2023). Additionally, Google's internal documents showed that ChatGPT passed the Google coding interview for Level 3 engineer with $183K salary (Elias, 2023).

[2] Even engineers at OpenAI did not expect how well accepted ChatGPT quickly became upon its release (Cowen, 2023).

we decompose the job listings into occupations and treat them as a collection of tasks. Building on recent research, we use detailed task data from the O*NET database and apply GPT API to categorize job requirements based on their exposure to Generative AI (Eloundou, Manning, Mishkin, and Rock, 2023; Felten, Raj, and Seamans, 2021). This allows us to examine the rapid changes in job market demand driven by Generative AI at an occupation-week level.

While past technologies primarily impacted blue-collar jobs, our findings reveal that Generative AI disproportionately affects white-collar jobs that require high levels of critical thinking and creativity. The most exposed occupations include computer system engineers, writers and authors, climate change policy analysts, and statistical assistants. These occupations generally require high technical expertise, analytical abilities, and effective writing skills. By contrast, the least exposed occupations mainly involve manual labor tasks, such as refuse and recyclable material collectors, wellhead pumpers, stockers and order fillers, machine feeders and offbearers, and recycling and reclamation workers. Industries that demand high educational attainment and specialized skills, such as Professional, Scientific, and Technical Services, as well as the Information sector, are among the most exposed to Generative AI. Conversely, industries like Accommodation and Food Service and Construction, which are more reliant on manual labor, face significantly less exposure. Geographically, states along the East and West coasts, and in the Great Lakes region, exhibit the highest levels of Generative AI exposure, while rural areas in the Midwest and industrial regions remain less affected. These findings stand in stark contrast to previous research on the impact of technology and robotics on the labor market (Acemoglu and Autor, 2011), where the most substantial shocks were observed among blue-collar workers.

Using a generalized difference-in-differences approach, we examine the responses of employers and employees to Generative AI. We find that firms with greater exposure to Generative AI significantly reduce their job postings following ChatGPT's release, reflecting a rapid adjustment to the advancement of the technology. We further examine impact of Generative AI at the granular occupational level. Contrary to the conventional view of automation uniformly displacing jobs highly exposed to the technology shock, our findings reveal a *heterogeneous* effect. The reduction in job opportunities is concentrated among occupations demanding lower levels of knowledge, skills, and education. In contrast, the demand for occupations that require higher levels of knowledge, skill, and education significantly increases after the release of ChatGPT. This heterogeneity in employer hiring responses suggests that Generative AI is substitutive for low-level office jobs (e.g., entry-level office positions) but complementary for high-level office jobs (e.g., executive positions). Our results suggest that heightened exposure to Generative AI does not inherently result in employee displacement but could instead drive higher demand for high-level labor and management roles. This further suggests that Generative AI may polarize the labor market by creating a disparity between jobs with higher and lower requirements.

We then examine whether companies adjust their job requirements in response to their exposure to Generative AI. Our analysis reveals that firms more exposed to Generative AI significantly increase the use of machine-learning and AI-related keywords following the release of ChatGPT. This shift in job descriptions suggests a strategic realignment, with companies placing greater emphasis on technology and innovation to stay competitive. The increased emphasis on AI-related skills also suggests that these firms are actively seeking candidates with expertise in machine learning and Generative AI, likely to address emerging skill gaps and support the integration of these technologies into their operations. Our findings suggest that Generative AI drives an increased demand for AI-related roles as companies prepare for the future.

Cross-sectional analysis reveals that the decline in hiring is more pronounced for firms with higher R&D intensity and greater cash holdings. Companies with high R&D intensity, often at the forefront of technological adoption, likely experience more significant shifts in workforce requirements as they integrate AI solutions.[3] Additionally, firms with substantial cash reserves can more quickly invest in transformative technologies like Generative AI. The increased demand for machine learning and Generative AI skills is more pronounced in larger firms, which benefit from

[3]Anecdotally, industries such as life sciences and chemicals have initiated the adoption of Generative AI foundation models within their R&D endeavors, a practice commonly referred to as "generative design." These foundation models have the capacity to generate prospective molecules, expediting the advancement of novel pharmaceuticals and materials. For instance, Entos, a biotechnology pharmaceutical firm, has integrated Generative AI with automated tools for synthetic development, enabling the design of small-molecule therapeutics. Importantly, these same principles extend beyond pharmaceuticals to encompass various other products, including larger-scale physical items and electrical circuits, among others (Mckinsey, 2023).

scale advantages, manage diverse projects, and handle complex datasets, thereby amplifying their need for skilled personnel proficient in machine learning and Generative AI.

How does the release of Generative AI affect employees? Just as Luddite workers during the Industrial Revolution feared being replaced by machines, today's workforce faces similar anxieties with the advent of Generative AI.[4] Our study provides novel evidence on employee reactions to this technological disruption. We find that prospective employees are less likely to engage in job interviews post-ChatGPT, particularly in non-management positions. This suggests a shift in hiring practices as companies increasingly prioritize specialized skill sets and technological expertise over general qualifications, with non-managerial roles being most affected.

We next examine the perceptions of current employees within organizations exposed to Generative AI. While initial exposure may have offered a sense of job security due to reduced external competition, concerns about potential downsizing and organizational restructuring persist. Our analysis shows that employee ratings on aspects such as Work/Life Balance, Career Opportunities, Compensation and Benefits, Senior Management, Culture and Values, Diversity and Inclusion, and Overall Recommendation remain largely unchanged in the short term. However, we observe a significant decline in ratings for CEO Approval and Business Outlook, particularly in firms highly exposed to Generative AI after the release of ChatGPT. This suggests that while current employees may not see immediate disruptions, they are increasingly worried about the long-term prospects of their organizations in the face of Generative AI.

Our findings reveal a mismatch between the reactions of employers and employees toward Generative AI. Employers are rapidly adapting by reducing job postings for standard white-collar positions and increasing the emphasis on AI-related skills, viewing Generative AI as a strategic tool to enhance efficiency and competitiveness. In contrast, employees are increasingly concerned about their long-term prospects, as reflected in significant declines in CEO Approval and Business Outlook ratings, especially in firms highly exposed to AI. While employers focus on the immediate benefits

[4] The Luddites were a group of early 19th-century English workers who destroyed machinery, particularly in the textile industry, as a form of protest. They feared that the introduction of machines would threaten their jobs and livelihoods. The term "Luddite" has since become associated with resistance to technological change.

of AI, such as cost savings and technological advancement, employees are more concerned about the future impacts on their roles, career growth, and organizational stability. This disconnect is further manifested in the reduced engagement of prospective employees, particularly in non-management roles, who may fear being replaced or marginalized by AI. The mismatch in perspectives between employers and employees suggests potential challenges in aligning organizational strategies with employee concerns, potentially affecting workforce morale and trust in leadership.

We conduct a battery of robustness tests to corroborate our findings. Our main analysis uses high-frequency weekly hiring data, which offers the advantage of capturing labor market responses in a timely manner but may be subject to random fluctuations. To address this potential concern, we examine alternative data frequency with monthly hiring data. In addition, we mitigate concerns about endogeneity between hiring strategies and firm characteristics using entropy balancing. Furthermore, we control for the possible influence of the COVID-19 pandemic, as the post-ChatGPT period overlaps with the post-COVID period. We employ various alternative measures of hiring quantity and Generative AI exposure to address potential measurement biases. To ensure our results are specifically attributable to Generative AI, we control for exposure to broader AI and other contemporary technologies. Our results remain consistent across all specifications.

Our contributions are threefold. First, we contribute to research on automation and the labor market by examining how employers react to the rapid rise of Generative AI. Prior literature has predominantly focused on the effects of automation and traditional AI on labor markets, particularly the impact on blue-collar workers performing routine tasks (e.g. Acemoglu, Autor, Hazell, and Restrepo, 2022; Acemoglu and Restrepo, 2022). Recent studies on Generative AI mostly focus on event studies on firm value and suggest a substitution effect where firms highly exposed to Generative AI are uniformly displaced (e.g. Eisfeldt, Schubert, and Zhang, 2023). Unlike previous research that focuses on industry- or firm-level analyses, we contribute by examining the effects of Generative AI at the occupational level. We uncover a *heterogeneous* effect: Generative AI substitutes for low-level office jobs (e.g., entry-level office positions) but complements high-level office jobs (e.g., executive positions). Our results suggest that heightened exposure to Generative AI does not inherently result in employee displacement but could instead drive higher demand for high-level labor and management roles. This also indicates that Generative AI could potentially polarize the labor market by widening the disparity between jobs requiring different levels of skills and expertise.

Second, we contribute by providing novel evidence on employee reactions. Prior studies predominantly focus on the macro labor outcomes such as employment rate and minimum wage, leaving internal firm dynamics largely unexplored (e.g., Acemoglu and Restrepo, 2020; Eloundou et al., 2023). To the best of our knowledge, our study is the first to provide large-scale evidence of the responses of both employers and employees to the introduction of Generative AI. We find a decline in long-term job outlook and a reduced interest in job interviews, indicating concerns over future prospects. Moreover, we identify a mismatch between how employers and employees perceive and react to the disruption brought by Generative AI. While employers are quickly adapting by increasing job postings for management positions and placing more emphasis on AI-related skills, employees are increasingly concerned about their future prospects. This disconnect is further manifested in the reduced engagement of prospective employees. By examining employer and employee responses in tandem, we shed light on how Generative AI disruption affects the workplace dynamic.

Third, we contribute to the emerging literature on Generative AI technology that has recently attracted attention from accounting researchers. Existing studies focus primarily on Generative AI's capabilities in performing specialized tasks that traditionally require human expertise, particularly in financial analysis and textual interpretation. For example, Kim, Muhn, and Nikolaev (2024) find that GPT4 outperforms financial analysts in predicting earnings changes. ChatGPT can facilitate processing information in complex corporate disclosure (Kim, Muhn, and Nikolaev, 2023a), and uncovering hard-to-quantify corporate risks (Kim, Muhn, and Nikolaev, 2023b). Bertomeu, Lin, Liu, and Ni (2023) find the ban of ChatGPT in Italy reduced analysts' information processing capacity. While these studies demonstrate the capability of Generative AI in specific tasks such as financial analysis, our study adds to this literature by systematically examining its impact across a broader universe of jobs, focusing on employer and employee responses to Generative AI.

More broadly, this paper examines real-time employer and employee reactions to the disruption of Generative AI. Existing research on the effects of technological advancement has predominantly focused on impacts on employment share changes or cross-sectional analysis at the industry level (e.g., Acemoglu and Autor, 2011; Acemoglu et al., 2022). We leverage high-frequency, forward-looking data that captures rapid shifts in job market demand at the occupation-week level. Our approach enables the tracking of immediate changes in hiring strategies and employee sentiment. Future research can build on this approach to explore real-time changes in internal firm dynamics in response to Generative AI adoption. For instance, future studies could examine how firms restructure roles, transitioning employees from lower-level tasks like bookkeeping to higher-level strategic or innovation roles, and how employer demand for AI-related skills shifts rapidly after new technologies are introduced, potentially exacerbating income inequality. Additionally, researchers could explore organizational restructuring, assessing whether Generative AI leads to flatter hierarchies and investigating the impacts on team dynamics and skill acquisition as firms adapt to the challenges of a rapidly evolving technological landscape.

The rest of the paper proceeds as follows. Section 2 describes the literature review. Section 3 presents our data, sample, and measure. Section 4 presents the descriptive statistics. Section 5 presents the research design. Section 6 presents the main results. Section 7 discusses robustness tests. Section 8 concludes.

2 Literature

2.1 Automation and Labor Market

Recent research has examined the impact of automation and technology adoption on the labor market, focusing primarily on general technology and robots. These studies show that blue-collar workers engaged in routine, manual tasks are the most affected, often experiencing reduced earnings and fewer employment opportunities in roles that can easily be replaced by machines(Acemoglu and Restrepo, 2020; Korinek and Juelfs, 2024). There is also evidence that traditional AI contributes to reductions in hiring for non-AI roles, as it automates repetitive processes (Acemoglu et al., 2022). Unlike traditional AI, which performs specific tasks using predefined rules, Generative

AI—such as OpenAI's ChatGPT—introduces new dynamics by impacting white-collar jobs that require creativity, critical thinking, and decision-making. Unlike prior technologies, Generative AI's versatility allows it to generate creative works, assist in content creation, and perform tasks previously considered immune to automation.

Our study contributes to this literature by examining real-time responses of both employers and employees to Generative AI. Our study departs from prior research, which predominantly focused on long-term impacts on blue-collar workers. We examine the immediate impacts of Generative AI, particularly on white-collar jobs that require critical thinking and creativity. A common challenge in previous studies has been isolating demand shifts from supply-side changes and capturing the exposure to technological disruption beyond the aggregated industry level. We leverage high-frequency, forward-looking data that captures rapid shifts in job market demand at the occupation-week and firm-week levels. While prior studies have generally found that automation leads to job displacement, we provide a more nuanced view, showing that Generative AI complements high-skilled white-collar roles, such as executive and managerial positions, while substituting for lower-skilled white-collar roles. Our findings departure from the traditional narrative surrounding automation's impact, expanding our understanding of how technology influences labor markets beyond routine manual jobs.

2.2 Generative AI Technology Shock

Generative AI's capabilities differ fundamentally from previous technologies, particularly its ability to handle tasks involving creativity and problem-solving rather than routine manual labor. ChatGPT, a key milestone in this evolution, has attracted considerable attention for its potential to transform the workplace by automating complex, cognitive tasks.

An emerging recent literature examines the capability in specific application of ChatGPT in accounting. Existing studies focus primarily on Generative AI's capabilities in performing specialized tasks that traditionally require human expertise, particularly in financial analysis and textual interpretation. For example, Kim et al. (2024) find that GPT4 outperforms financial analysts in predicting earnings changes. ChatGPT can facilitate processing information in complex corporate disclosure (Kim et al., 2023a) and uncovering hard-to-quantify corporate risks (Kim et al., 2023b). Bertomeu et al. (2023) find the ban of ChatGPT in Italy reduced analysts' information processing capacity. Other studies explore the cross-sectional variations in exposure to Generative AI at the industry level (Eloundou et al., 2023). Eisfeldt et al. (2023) conduct event studies on the effects of Generative AI firm value and find a substitution but not complementary effect for firms with high exposure. While previous research has largely focused on specific tasks, our study systematically examines Generative AI's effects on the labor market by examining employer and employee responses across a wide spectrum of jobs.

Existing research on Generative AI and, more broadly, on technological change and automation in the labor market have predominantly focused on macro-level outcomes like employment rate and minimum wage, leaving the internal firm dynamics largely unexplored. To the best of our knowledge, our study is the first to provide large-scale evidence of the responses of both employers and employees to the introduction of Generative AI. We find a heterogeneous effect and provide novel evidence on employee reactions. Our results suggest a mismatch between the reactions of employer and employee towards Generative AI.

3 Data, Sample, and Measure

3.1 LinkUp Data

We obtain data on job postings, vacancies, and descriptions from LinkUp in order to capture real-time responses of employer demand to Generative AI. LinkUP data offers several advantages. Unlike studies that rely on industry-level measures or measures using self-reported employment status or aggregated listings from job boards (Eisfeldt et al., 2023; Eloundou et al., 2023), LinkUp data allows us to directly assess a firm's exposure to Generative AI by tracking changes in actual job openings disclosed by the employers, capturing firms' response to Generative AI in real-time. LinkUp sources job listing data daily directly from over 60,000 employer websites, starting from

August 2007. The comprehensive dataset captures information on each job posting's creation date, deletion date, job description, O*NET occupation code, job title, location, and so on. By 2023, LinkUp covered about 89.3% of US public firms, which together accounted for 97.9% of total assets of 2022.[5]

3.2 Glassdoor Data

We obtain potential and current employees' reaction data from Glassdoor. Glassdoor is an online platform providing insights into companies and workplaces, offering job search functionalities, company reviews, and salary information. Launched in 2008, it allows employees to anonymously share their experiences, opinions, and salary details about their workplaces. The information helps job seekers make more informed decisions by providing a transparent view of company culture, work environment, and compensation practices.

Glassdoor uses multiple methods to ensure data reliability. First, it mandates the use of an active email address or a legitimate social networking account, to ensure authenticity and prevent companies from generating fake reviews or self-promotion. Second, using an algorithm, Glassdoor identifies and flags potentially fraudulent reviews. These flagged reviews are then assessed by an employee to eliminate any invalid submissions (Green, Huang, Wen, and Zhou, 2019). Moreover, starting in 2015, Glassdoor implemented a "give to get" model to counter concerns regarding potential bias in submissions. This model limits a job seeker's access to online information until they provide their review of an employer. Marinescu, Chamberlain, Smart, and Klein (2021) find that Glassdoor's "give to get" system effectively minimizes biases in the selection of individuals who post reviews.

3.3 O*NET data

We obtain the occupation data from the O*NET Resource Center, which is sponsored by the U.S. Department of Labor, Employment & Training Administration. It serves as a valuable

[5] We compare the number of unique firms covered in LinkUp to the number of US firms covered in the Compustat merged data.

hub for occupational information and workforce analysis and provides a comprehensive database known as O*NET. It provides rich information on various aspects of occupations, including skills, abilities, knowledge, tasks, work activities, and work context. This comprehensive dataset allows for a thorough analysis of occupation characteristics. To ensure the reliability and consistency of records, O*NET uses a standardized framework for classifying and describing occupations and gathers occupation information through rigorous job analyses and surveys.

O*NET is regularly updated to reflect changes in the labor market and evolving job roles, so that users have access to the most recent characteristics of different occupations. We use O*NET 27.1 for our study because it was released in November 2022, making it the most recent version of the occupational characteristics database available before the launch of ChatGPT.

3.4 Generative AI Exposure

We measure a company's Generative AI exposure in three steps. First, we obtain the task listing associated with each occupation from O*NET 27.1 database. Subsequently, we employ GPT-3.5 to categorize each occupation-task pair based on exposure to Generative AI technologies.[6] Second, we determine the exposure of each occupation to Generative AI by aggregating occupation-task measurements. Third, we obtain a company's Generative AI exposure score by aggregating the occupation-level Generative AI exposure measure using firm-occupational job posting data from LinkUp. We illustrate how we construct the exposure measure at the task, occupation, and firm level below.

3.4.1 Task Exposure

O*NET provides a detailed mapping of O*NET-SOC codes (occupations) to tasks associated with each occupation. O*NET 27.1 contains 17,953 distinct occupation-task pairs, with an average

[6]Eloundou et al. (2023) employs GPT-4, which includes vision capabilities, for assessing exposure to Generative AI. We posit that GPT-3.5 performs similarly for our task-based assessment, as both models share core language capabilities crucial for evaluating text-based tasks. For the purposes of this study, the difference in vision capabilities does not impact the results, as evaluating occupation exposure does not rely on visual processing.

of 21 distinct tasks per occupation.[7] Following Eloundou et al. (2023) and Eisfeldt et al. (2023), we use GPT API 3.5 Turbo to assess each task to determine if it can be executed more efficiently using ChatGPT and similar large language models (LLMs) or future applications built upon their capabilities. The format of the prompt submitted to the OpenAI GPT API is presented in Appendix C.[8]

We provide GPT with a detailed explanation of the rubric for scoring LLM exposure. Additionally, we present four example interactions between a user and an assistant to help GPT understand the desired responses. The task statement, along with the occupation's full title and task descriptions, is then submitted. GPT responds with a label, explanation, and confidence level. For each occupation-task pair, GPT selects one label from four categories: (1) E_1 - Direct exposure: direct access to the LLM through an interface like ChatGPT or the OpenAI playground alone can reduce the time it takes to complete the task with equivalent quality by at least half; (2) E_2 - Exposure by LLM-powered applications: Having access to the LLM alone may not reduce the time it takes to complete the task by at least half, but it is easy to imagine additional software that could be developed on top of the LLM that would reduce the time it takes to complete the task by half; (3) E_3 - Exposure given image capabilities: There is a significant reduction in the time it takes to complete the task given access to an LLM and image capabilities; (4) E_0 - No exposure: Exposure by LLM does not decrease the time it takes for an experienced worker to complete the task with high quality by at least half. Table 2 presents randomly selected examples of GPT scores assigned to task statements and GPT-provided explanations.

[7] O*NET provides highly comprehensive task listings with detailed descriptions for each task. There is no overlap of tasks between different occupations.

[8] Several studies have confirmed the consistency and reliability of using GPT API to measure task exposure to Generative AI. For example, Eloundou et al. (2023) compare the GPT tag and human tag, and find approximately 80% agreement between the two. Eisfeldt et al. (2023) conduct a comparison of scores assigned in three different GPT runs, revealing a very high level of agreement. In nearly 88% of all cases, they consistently arrive at the same score.

3.4.2 Occupation Exposure

We then aggregate task-level exposure scores to the occupation level. We view each occupation as a bundle of tasks and measure occupation-level Generative AI exposure as the share of directly exposed tasks(E_1) and exposed by LLM-powered applications (E_2) of all tasks, following Eloundou et al. (2023) and Eisfeldt et al. (2023). We half the weight of E_2 to reflect the fact that the task is supplemented only by the potential development of LLM. The occupation exposure is weighted by the proportion of tasks exposed to Generative AI.

$$E_o = \sum_{tasks\ in\ o} \frac{E_1 + 0.5 * E_2}{E_0 + E_1 + E_2 + E_3}.$$

3.4.3 Firm Exposure

We calculate a company's generative AI exposure score by aggregating the occupation-level Generative AI exposure measure using firm-occupational job posting data from LinkUp. For each job posting, LinkUp collects a comprehensive set of variables, including job creation date, job deletion date, source company, and O*NET code. We denote $emp_{f,o}$ as the total number of jobs created by firm i on occupation o from January 1st, 2022 to November 30th, 2022 (the launch date of ChatGPT), and emp_f the total number of job created by firm i in that time period. The firm exposure to Generative AI (E_f) is the weighted average of occupational exposure.

$$E_f = \sum_{occupations\ in\ f} \frac{emp_{f,o}}{emp_f} * E_o.$$

3.5 Other Occupation Characteristics and Firm Characteristics

We obtain occupation characteristic data from the O*NET 27.1 database. O*NET provides detailed descriptions of various occupations, including tasks performed, skills required, education, and job interests. O*NET's occupation characteristic data is regularly updated to reflect changes in the workforce, technology, and industry trends. We obtain other firm characteristic data from Compustat, including firm size, tangibility, labor intensity, market-to-book ratio, ROA, cash holdings, sales growth, and stock returns.[9]

4 Descriptive Statistics

4.1 Summary Statistics

In Table 1, we present descriptive statistics for variables used in our analyses, with detailed definitions provided in Appendix A, Table A. 1. Panel A summarizes firm-week level variables obtained from LinkUp. We obtain 540,575 unique firm-week observations. $Log(JobCreated + 1)$ has a mean of 2.535 and a median of 2.485. These values indicate that, on average, each company posts 11.5 new vacancies every week. $Log(MachineLearning + 1)$ and $Log(GenerativeAI + 1)$ have means of 0.069 and 0.014, suggesting the mention of keywords related to machine learning and Generative AI is generally rare in job descriptions. *Post-ChatGPT* has a mean of 0.072, indicating that most of our sample period is before the release of ChatGPT. The average of *Gen. AI Exposure* is 0.37, showing that a substantial portion of the firms in our sample are meaningfully exposed to Generative AI.

In Panel B, we report descriptive statistics for the firm characteristics from Compustat. Panel C presents descriptive statistics for occupation-week level variables constructed using LinkUp's job postings and O*NET occupation data. Specifically, we first decompose the LinkUp database into sub-samples based on occupation characteristics obtained from O*NET 27.1. Then we aggregate LinkUp's job postings database into occupation-week level.

Panels D and E contain descriptive statistics for Glassdoor interview and review data, which we collected from the Glassdoor website, covering the period from January 1, 2021, to August 23, 2023. The average of *Interview Num* is 2.862, indicating that approximately three new interview comments are posted weekly on average. For manager and non-manager positions, the averages of interview posts are 0.306 and 2.547, respectively, reflecting a higher number of interview posts for non-manager positions. *Experience Score* is a discrete variable measuring the interviewee's impres-

[9] Due to the unavailability of firm characteristics for 2023, we backfill the data from 2022.

sion of the interview that takes on values of -1 (negative experience), 0 (neutral experience), and 1 (positive experience). Similarly, *Difficulty Score* is a discrete variable measuring the interviewee's view of the interview's difficulty that takes on values of -1 (easy), 0 (medium), and 1 (difficult). *Experience Score* has a mean of 0.443, suggesting that overall, employees hold slightly positive attitudes towards interview experiences. The average of *Difficulty Score* is -0.182, indicating that interviews are perceived as not particularly challenging overall. For employee review variables in Panel E, *Approves of CEO*, *Recommend*, *Business Outlook* are discrete variables with values of -1, 0, and 1. *Overall Rating*, *Work/Life Balance*, *Career Opportunities*, *Career Opportunities*, *Compensation and Benefits*, *Senior Management*, *Culture* & *Values*, *Diversity* & *Inclusion* are discrete variables taking integer values from 1 to 5.

4.2 Generative AI Exposure by Occupation, Firm, Industry, and State

Next, we examine how Generative AI exposure varies by occupation, firm, industry, and state. We present examples of the most and least exposed occupations in Table 3. The most exposed occupations include computer system engineers, writers and authors, climate change policy analysts, and statistical assistants. These occupations generally require high technical expertise, analytical abilities, and effective writing skills. By contrast, the least exposed occupations mainly involve manual labor tasks, such as refuse and recyclable material collectors, wellhead pumpers, stockers and order fillers, machine feeders and offbearers, and recycling and reclamation workers. This is in stark contrast to findings from previous research on the impact of technology and robotics on the labor market. For example, Acemoglu and Autor (2011) find that the most substantial shocks are observed among blue-collar workers, whereas jobs requiring cognitive abilities such as problem-solving, judgment, and creativity are much less susceptible to automation.

In Table 4, we present examples of the most exposed and least exposed firms by Generative AI. Companies with high exposure are predominantly operating in technical and information-related industries, while those with the least exposure are in manufacturing industries. Table 5 displays the Generative AI exposure across different industries. The exposure score is the average of all companies in that industry, and we also provide the number of companies in our sample that belong to that industry. Industries that require high education and skills, such as Professional, Scientific, and Technical Services, as well as Information are the most exposed, while the Accommodation and Food Service and Construction industries are the least exposed.

Substantial variations of Generative AI exposure exist across different states, as shown in Figure 1. The state-level exposure score is calculated by aggregating the occupation-level exposure score weighted by state occupational employment in 2022.[10] The states most exposed are those located on the East and West coasts and in the Great Lakes region.[11] The least exposed states include Nevada, Wyoming, Alabama, North Dakota, Mississippi, Arkansas, Louisiana, and Indiana. The notable variations in Generative AI exposure across regions and industries suggest that the development of Generative AI may exacerbate existing labor market inequalities. Coastal and urban regions with higher exposure, which are already home to tech hubs and industries requiring high-skilled labor, may see a greater demand for AI-related skills and experience faster economic growth driven by AI. In contrast, low-exposure regions such as the Midwest and areas reliant on agriculture and manufacturing may experience less immediate disruption in their traditional industries, but also fewer opportunities for AI-driven innovation. As a result, the gap between high-exposure and low-exposure regions may widen, leading to a more pronounced economic and skill divide between tech-forward urban areas and more industrial, rural regions.

4.3 Generative AI Exposure and Occupation Characteristics

We explore the correlation between Generative AI exposure and occupational characteristics to discern the types of occupations most profoundly impacted by this technological advancement. Figure 2 presents the Generative AI exposure score across occupation requirements for various

[10]We obtain state occupational employment data from the U.S. Bureau of Labor Statistics.

[11]Colorado is among the most exposed states because of its employment structure. According to the statistics of Colorado's occupational employment in 2022, the top 4 occupations that have the highest employment are Office and Administrative Support Occupations, Sales and Related Occupations, Food Preparation and Serving Related Occupations, and Business and Financial Operations Occupations.

knowledge elements, skill elements, and ability elements. The O*NET dataset characterizes specific elements that describe the type of knowledge, skill, or ability required for each occupation, assigning a score to each occupation at the element level. In Figure 2, the exposure score at the element level is calculated as the average occupation-level exposure score, weighted by the value of each occupation across elements from O*NET 27.1. Among knowledge elements, the most affected element is *Communication and Media*, while the least affected element is *Mechanical*. Among skill elements, the most affected element is *Programming*, while the least affected element is *Operation and Control*. Among ability elements, the most affected element is *Written Expression*, while the least affected element is *Multilimb Coordination*.[12]

We then estimate the univariate regression to examine the relation between various occupation characteristics and the occupation exposure score. Binscatter plots and regression coefficients, shown in Figure B.1 in Section Appendx B, reveal several notable patterns. We find a positive correlation between the requirements for knowledge, skills, and education, and the exposure score, indicating that occupations demanding higher levels of these attributes are more affected by Generative AI. Conversely, we observe a negative correlation between the need for ability and training, and the exposure score, suggesting that jobs requiring less technical training are less impacted. These findings highlight that Generative AI is more likely to influence roles that require advanced expertise and formal education.

We further explore the variation in Generative AI exposure across different categories of jobs using O*NET's classification of occupations into five job zones. Each job zone reflects varying levels of education, experience, and on-the-job training required for different occupations. Job Zone 1 includes roles like fast-food cooks, fishing and hunting workers, and amusement attendants, which require minimal preparation. At the other end of the spectrum, Job Zone 5 includes highly skilled professions such as chief executives, administrative law judges, economists, and acupuncturists, which demand extensive expertise and education.

In Panel f of Figure B.1, we present a box plot of exposure scores across the five job zones.

[12] The definition of *Multilimb Coordination* given by O*NET is the ability to coordinate two or more limbs (for example, two arms, two legs, or one leg and one arm) while sitting, standing, or lying down.

Our analysis shows that Generative AI exposure tends to increase with job zone levels. However, we observe that occupations in Zone 5, which require the highest levels of skill, knowledge, and education, are less exposed to Generative AI compared to those in Zone 4. This non-linear pattern suggests that while standardized white-collar jobs (found in Zone 4) are more vulnerable to Generative AI, both blue-collar jobs (Zone 1) and high-level executive roles (Zone 5) face lower exposure. This indicates that the impact of Generative AI on the labor market is concentrated in mid-level white-collar jobs, where routine and standardized tasks are more likely to be automated.

5 Research Design

We employ a generalized difference-in-differences design to examine employer and employee reactions to Generative AI exposure. Unlike previous generations of automation, which advanced more gradually, the release of ChatGPT was relatively unexpected. ChatGPT's launch marked a significant breakthrough in natural language processing and generative capabilities, drawing widespread attention to the potential applications of Generative AI systems. Notably, even OpenAI was surprised by the immediate and substantial impact ChatGPT had upon its release. This makes the introduction of ChatGPT on November 30, 2022, a natural shock for labor market participants, providing a unique opportunity to examine real-time responses. Specifically, we estimate the following equation:

$$Y_{i,t} = \beta_0 + \beta_1 \textit{Post-ChatGPT}_t \times \textit{Gen. AI Exposure}_i + \beta_2 \textit{Controls}_{i,t} + \gamma_i + \theta_t + \epsilon_{i,t}, \tag{1}$$

where we regress a series of employers' and employees' reactions in firm i and week t ($Y_{i,t}$) on the measure of firm exposure to the technology ($\textit{Gen. AI Exposure}_i$) and its interaction with an indicator which equals one when the week t is after the release of ChatGPT ($\textit{Post-ChatGPT}_t$). The continuous Generative AI exposure score at occupation or firm level allows for a more nuanced representation of the exposure to the technology compared with discrete exposure measures. This granular measurement captures variation in exposure across companies, providing a more detailed understanding of how different levels of exposure may impact labor market outcomes. Moreover, by leveraging high-frequency (weekly) labor market data, we are able to capture subtle shifts in employment dynamics over shorter intervals and track how exposure to Generative AI impacts employment responses on a week-to-week basis.

$Controls_{i,t}$ represent a vector of firm-year level controls. Firm size accounts for operating scale, while tangibility and labor intensity capture investments in both tangible and intangible assets. Sales growth reflects current growth, and the market-to-book ratio captures growth opportunities. We include return on assets (ROA) to control for current profitability and stock return to account for market signals about future profitability. Cash holdings are included to capture liquidity and financial constraints. We incorporate firm-fixed effects to control for time-invariant firm characteristics, such as company location, industry, and community culture. Year-week fixed effects are used to account for time trends and events affecting the overall labor market, such as GDP growth, population changes, and contemporary policy shifts.[13] Standard errors are clustered at the firm level. Appendix A, Table A. 1 provides detailed variable definitions.

6 Main Results

In this section, we analyze how various stakeholders in the labor market respond to the introduction of Generative AI. We begin by examining employer behavior, including changes in hiring quantity and job descriptions. We also invetigate whether the effects of Generative AI differ across occupations with varying skill requirements and firm characteristics. Next, we turn to potential employees and assess their job-seeking behavior using Glassdoor interview data. Lastly, we analyze current employees' perceptions of their companies by examining changes in Glassdoor reviews.

[13]For instance, the interest rate hikes from August 2022 to March 2023 are absorbed in the year-week fixed effects.

6.1 Employer Reactions

6.1.1 Hiring Quantity Overall Effect

We begin by examining the effects of Generative AI exposure on companies' hiring quantity. We predict a negative relation between hiring quantity and the Generative AI exposure score. Table 6 presents the results from estimating equation (1). The adjusted R-squared is 0.707 when including control variables, firm fixed effects, and year-week fixed effects, indicating that the model has satisfactory explanatory power. Column (3) shows that the coefficient of the difference-in-differences interaction term *Post-ChatGPT*$\times$*Gen.AI Exposure* is significantly negative at the 1% level, supporting our prediction that more exposed companies reduce their hiring quantity compared with less exposed companies after the launch of ChatGPT. This reduction of hiring quantity in response to Generative AI is economically significant. A one-standard-deviation increase in Generative AI exposure (0.139) corresponds to a decrease of 18.2% (0.139 $\times$ 1.310) of weekly hiring quantity after the launch of ChatGPT. This effect is comparable to major labor market disruptions observed in past events. For example, Chetty, Friedman, and Stepner (2024) estimate a 14% hiring decline for lower-wage employees during the COVID-19 pandemic.

We use $Log(Job\ Created + 1)$ as our dependent variable because the raw hiring quantity $Job\ Created$ is skewed and equal to zero for many observations. However, recent studies find that estimating linear regressions of the natural logarithm of one plus the outcome leads to estimates that lack a straightforward interpretation and may show an unexpected sign on average (Chen and Roth, 2023; Cohn, Liu, and Wardlaw, 2022). Additionally, adding an arbitrary positive constant can distort the original data structure, with even small adjustments significantly impacting empirical results (Duan, Manning, Morris, and Newhouse, 1983; N'guessan, Featherstone, Odeh, and Upendram, 2017). To address these potential concerns, we conduct a series of robustness tests presented in Table 6, Panel B. In Column (1), we conduct a Poisson regression for the raw value of hiring quantity ($Job\ Created$). In subsequent columns, we further test different transformations to mitigate skewness: using the raw hiring quantity *Job Created* in Column (2), the natural logarithm of hiring quantity *Log(Job Created)* in Column (3), the IHS transformation of hiring quantity

IHS Job Created in Column (4), and an indicator for non-zero hiring in Column (5).[14] We find consistent results across all specifications.

6.1.2 Hiring Quantity Heterogeneous Effect

We next investigate the heterogeneous effects of Generative AI on hiring quantity across various occupation characteristics. We segment the LinkUp job postings into sub-groups based on the median values of occupation characteristics, such as knowledge, skills, education, and technical expertise classified by O*NET. Then, we aggregate job postings at the occupation-week level to generate the occupation-level exposure score *Gen. AI Exposure*. As shown in Table 7, occupations with higher requirements for knowledge, skills, education, and technical expertise exhibit a positive relationship between Generative AI exposure and hiring quantity. Conversely, occupations with lower requirements in these areas show a negative relationship. This heterogeneous effect is further manifested in the significantly positive coefficient of the hiring ratio between high and low-requirement jobs (*High/Low Ratio*) across all job requirement categories.

Our results reveal a *heterogeneous* effect of exposure to Generative AI on employers' hiring intentions. The reduction in hiring is concentrated in occupations that require lower levels of knowledge, skill, education, training, and technical skill. In terms of economic manganitude, a one-standard-deviation increase in occupation-level Generative AI exposure (0.232) corresponds to a decrease of 7.3% in weekly hiring for low-knowledge-requirement occupations, a 7.0% decrease for low-skill-requirement occupations, a 11.6% decrease for low-education-requirement occupations, a 2.2% decrease for low-training-requirement occupations, and a 11.8% decrease for low-technical-skill-requirement occupations after the launch of ChatGPT.

In contrast, the demand for occupations that require higher levels of knowledge, skill, education, training, technical skill increases after the release of ChatGPT. Specifically, a one-standard-deviation increase in occupation-level Generative AI exposure (0.232) corresponds to an increase

[14] The inverse hyperbolic sine (IHS) transformation is widely used as alternative way to transform right-skewed variables that include multiple zeros instead of the log transformation. This is because it can avoid the arbitrary manipulation problem of log transformation and can be easily applied to zeros (Aihounton and Henningsen, 2021; Bellemare and Wichman, 2020).

of 3.3% in weekly hiring for high-knowledge-requirement occupations, a 3.2% increase for high-skill-requirement occupations, a 7.6% decrease for high-education-requirement occupations, a 1.8% decrease for high-training-requirement occupations, and a 8.4% increase for high-technical-skill-requirement occupations after the launch of ChatGPT.

Table 7, Columns 16-20 present the heterogeneous effects of Generative AI across five job zones characterized by O*NET.[15] For example, Zone 5 includes CEOs, investment fund managers and lawyers, Zone 4 includes supply chain managers, credit analysts, and computer programmers. Zone 3 includes tax preparers, technicians, and travel agents, Zone 2 includes food service managers and title examiners. Zone 1 includes dishwashers, truck and tractor operators, and taxi drivers. The results indicate that jobs created for positions more exposed to Generative AI decrease within subsamples of Zone 1, Zone 2, and Zone 3, where the jobs need less preparation. In contrast, a significant positive relation is observed within Zone 4 and Zone 5. In terms of economic magnitude, a one-standard-deviation increase in occupation-level Generative AI exposure (0.232) corresponds to a decrease of 0.5% in weekly hiring for Zone 1 occupations, a 6.2% decrease for Zone 2 occupations, and a 3.7% decrease for Zone 3 occupations. Conversely, Zone 4 occupations experience a 6.0% increase, and Zone 5 occupations see an 1.4% increase. These results further indicates the heterogeneous effect of Generative AI on hiring demand across different job zones, with lower-skill jobs being more negatively affected and higher-skill jobs seeing a rise in demand.

To summarize, our findings reveal a heterogeneous effect that Generative AI substitutes for lower-level jobs, automating tasks that typically require less specialized knowledge or training effort. On the other hand, it complements and increase demand for higher-level roles, possibly by enhancing productivity and efficiency in complex tasks. Contrary to the conventional view that automation uniformly displaces jobs that are highly exposed to the new technology, our results suggest that heightened exposure to Generative AI does not inherently result in employee displacement but could instead drive higher demand for high-level labor and management roles. This further suggests that

[15] We include occupation fixed effects to control for time-invariant occupation characteristics. We also include week fixed effects to control for time trends in employment. Standard errors are clustered at the occupation level. We didn't include control variables because of the serious missing data problem of occupation variables in public survey databases such as the Bureau of Labor Statistics.

Generative AI may polarize the labor market for office jobs by creating a disparity between jobs with higher and lower requirements.

6.1.3 Job Posting Description

In this section, we investigate whether companies adjust the content of their job postings in response to their exposure to Generative AI. Job descriptions provide valuable insights into companies' potential developing strategies. We posit that more exposed companies increase the use of words related to machine learning and Generative AI after the release of ChatGPT. We obtain detailed original job posting content posted by the company from LinkUp. As shown in Appendix A, Figure A.1, a typical job posting includes essential details such as company information, a brief job summary, specific responsibilities and duties, required qualifications, key skills, and information about the work environment. We present the keywords used to identify machine learning and Generative AI-related vacancies in Appendix A, Table A. 2.

In Table 8 we report the regression results. We examine various dependent variables, including (1) the ratio of the count of keywords in job descriptions to the total number of jobs posted for firm i in week t (*Machine Learning Ratio* and *Generative AI Ratio*) in Column 1 and Column 4; (2) an indicator variable that equals one when the company posts new job vacancies that contain keywords in week t (*Machine Learning Dummy* and *Generative AI Dummy*) in Column 2 and 5; and (3) the natural logarithm of one plus the count of keywords in job descriptions ($Log(Machine\ Learning{+}1)$ and $Generative\ AI + 1)$) in Column 3 and 6. We find a significantly positive effect on the use of machine-learning-related keywords and Generative-AI-related keywords across all columns. In terms of economic magnitude, a one-standard-deviation increase in Generative AI exposure is related to a 2.3% increase in Machine Learning keywords in job descriptions and a 1.3% increase in Gen AI keywords. These increases correspond to a 33.3% and 92.9% rise, respectively, compared to the unconditional sample means of 6.9% for $Log(Machine\ Learning + 1)$ and 1.4% for $Log(Generative\ AI + 1)$. The economic magnitude is substantial compared to prior research on labor demand shifts. For instance, Modestino, Shoag, and Ballance (2020) report that the surge in unemployment during the Great Recession led to an 18% to 25% increase in skill requirements between 2007 and 2010. The increased emphasis on machine learning and Generative AI skills in job postings is consistent with heterogeneous effects of occupation characteristics in Table 7 in the sense that companies shift job demands to positions that require higher skills and education, such as Generative AI and machine learning skills.

The shift in job descriptions in response to Generative AI suggests that companies may be placing a greater emphasis on technology and innovation to remain competitive in an evolving landscape. Additionally, the increased focus on machine learning and Generative AI skills suggests firms are actively seeking candidates to fill specific technical skill gaps within their workforce.

6.1.4 Cross-Sectional Analysis

To shed more light on the changes in employment in response to Generative AI exposure, we estimate several cross-sectional analyses based on company characteristics by using the following equation:

$$\begin{aligned} Y_{i,t} = \beta_0 + \ & \beta_1 Firm\ Characteristic_{i,t} \times Post\text{-}ChatGPT_t \times Gen.\ AI\ Exposure_i \\ & + \ \beta_2 Post\text{-}ChatGPT_t \times Gen.\ AI\ Exposure_i + \ \beta_2 FirmCharacteristic_{i,t} \\ & + \ \beta_3 Controls_{i,t} + \ \gamma_i + \ \theta_t + \ \epsilon_{i,t}, \end{aligned} \tag{2}$$

where the independent variable $Y_{i,t}$ is the same as the variables used in Table 6 and Table 8. Other variables are the same as the ones in equation (1). The coefficient of interest is β_1. We explore various firm characteristics, including the firm size (*LogSize*), R&D intensity (*R&D Intensity*), cash holdings (*CashHold*), and corporation age (*CorpAge*).

We report the regression results in Table 9. Panel A presents findings regarding hiring quantity, while Panels B and C present outcomes related to the use of machine-learning and Generative-AI keywords in job descriptions. Results show that the negative effects on hiring quantity are more pronounced for firms with higher R&D intensity and more cash holdings. Companies with high R&D intensity, often leading in technological adoption, are likely to undergo more pronounced shifts in workforce requirements as they integrate innovative AI solutions. This contributes to the observed decrease in overall hiring quantity. Moreover, the trend of more serious effects on companies with high cash holdings might be attributed to the financial flexibility afforded by substantial cash reserves, allowing companies to invest in transformative technologies like generative AI. The higher cash holdings empower these companies to navigate the transitional phase associated with AI adoption more efficiently, influencing the observed decrease in overall hiring quantity.

Panels B and C show that the effects of increased demand for technical skills are more pronounced for larger firms. This may be attributed to the scale advantages that larger companies enjoy, facilitating their more extensive utilization of Generative AI technologies. Larger organizations typically engage in a wider range of operations, manage diverse projects, and handle complex datasets, thus amplifying their demand for skilled personnel proficient in machine learning and Generative AI.

6.2 Potential Employee Reactions

We next investigate whether the introduction of Generative AI impacts job seekers' behavior by analyzing interview feedback shared on the Glassdoor Interview forum. In line with the observed decrease in hiring quantity, we find that the number of interviews conducted with highly exposed firms declines. Column (1) of Table 10 shows a significant reduction in the number of interviews for companies with higher Generative AI exposure during the post-ChatGPT period. This suggests that job seekers may be increasingly hesitant to pursue opportunities with firms highly exposed to AI.

Furthermore, we examine the subsamples of manager-position interviews and non-manager-position interviews. As shown in Columns (2) and (3), although the effects are both negative for manager positions and non-manager positions, the reduction in interview counts is primarily driven by non-managerial positions. This suggests a mismatch between employer and employee reactions to Generative AI. As discussed earlier, we find a heterogeneous effect of Generative AI exposure on employer demand with increased demand for higher-level positions such as management roles. However, employees appear not to differentiate between management and non-management job opportunities, maintaining a generally pessimistic outlook across both categories.

Column (5) shows that interview difficulty significantly increases for companies with higher Generative AI exposure, while interview experience scores remain unaffected, indicating that companies maintain consistent interview quality. This rise in difficulty likely reflects a shift toward hiring candidates with more advanced technical knowledge or those who can adapt to rapidly evolving technological environments. Our findings on potential employee reactions point to a growing polarization in the labor market, where demand for specialized expertise is rising, while more routine positions may be increasingly vulnerable or replaced by automation.

6.3 Current Employee Reactions

In this section, we examine the changes in current employee perceptions of the company in response to Generative AI exposure. Existing employees may initially experience a sense of job security due to reduced external competition for positions. However, concerns about potential downsizing or restructuring within the company may persist. To evaluate the current employees' reactions, we utilize the Glassdoor review database, one of the largest platforms for employees to share insights about their employers. Glassdoor offers various numerical ratings that reflect key aspects of employee feedback, including *Approves of CEO*, *Overall Rating*, *Work/Life Balance*, *Career Opportunities*, *Compensations and Benefits*, *Senior Management*, *Culture* & *Values*, *Diversity* & *Incloution*, *Recommend*, and *Business Outlook*.

We present the results in Table 11. We find ratings in *Work/Life Balance*, *Career Opportunities*, *Compensations and Benefits*, *Senior Management*, *Culture* & *Values*, *Diversity* & *Inclusions*, and *Recommend*, are not significantly affected in the short run, indicating that the immediate work environment has not worsened for current employees in highly exposed companies. However, there is a significant negative effect on *Approves of CEO* and *Business Outlook*, indicating a noteworthy shift in the future outlook of firms, particularly those highly exposed to Generative AI following the ChatGPT release. The decrease in recognition of CEO might reflect a concern for the CEO's ability to adapt to this technology transition. In addition, the decline in future outlook suggests that while current employees may not perceive immediate disruptions, they are increasingly worried about the long-term prospects of their organizations in light of Generative AI. Our findings suggest that while current employee experiences within their companies do not decline in the short term, their perceptions of long-term organizational development and future competitiveness significantly worsen after the release of ChatGPT, particularly in companies highly exposed to Generative AI.

7 Robustness

7.1 Alternative Data Frequency

In our main analysis, we exploit firm-week level hiring data to examine whether firms that are highly exposed to Generative AI reduce hiring quantity after the release of ChatGPT. The high-frequency data enables us to estimate the employer's reaction to the technology change in a timely manner. However, the use of high-frequency data may capture random disturbances. To alleviate this concern, we aggregate dependent variables into the firm-month level and reestimate equation (1). As shown in Panel A of Table 12, we obtain robust results consistent with our main findings using alternative data frequency. The economic magnitude is meaningful, with a one-standard-deviation increase in exposure to Generative AI associated with approximately a 19.1% decrease in monthly job postings following the launch of ChatGPT.

7.2 Alternative Measures of Firm Exposure to Generative AI

In this section, we use alternative measures of a firm's Generative AI exposure to mitigate potential measurement bias. We compute a company's generative AI exposure score by aggregating the occupation-level Generative AI exposure measure, weighting it based on the share of jobs created by firm i in occupation o from January 1st, 2022, to the launch of ChatGPT. We use alternative weights to test the robustness of our results and present results in Panel B, Table 12. We use the weight of the share of jobs created within one month prior to the launch of ChatGPT in Column (1), the share of jobs created within six months prior to the launch of ChatGPT in Column (2), within eighteen months in Column (3), and within two years in Column (4). The coefficients are significantly negative across all columns, which confirms that our results are not sensitive to our measure of firm Generative AI exposure.

7.3 Entropy Balancing Test

Another concern arises from the potential endogeneity of firms' hiring strategies with respect to firm characteristics, beyond their Generative AI exposure. Consequently, we use entropy balancing to address imbalances in firm characteristics between those with above-median Generative AI exposure scores (treatment group) and those with below-median Generative AI exposure scores (control group). This process is widely adopted to enhance comparability by re-weighting observations, mitigating the impact of confounding variables, and thereby improving the validity of causal inferences (Basri, Felix, Hanna, and Olken, 2021; Hainmueller, 2012). Panel C shows that our estimated coefficients are consistently negative and the economic magnitude remains similar to the main analysis.

7.4 Control for the Effects of COVID

Recent research indicates the effects of COVID-19 on the labor market, including exacerbated inequality, reduced hiring quantity, and shifting working styles (Baek, McCrory, Messer, and Mui, 2021; Barry, Campello, Graham, and Ma, 2022; Giupponi and Landais, 2023). Considering the overlap between the post-ChatGPT period and the post-COVID period, the effects of COVID might introduce bias into our results. For example, firms operating in the information industry are generally highly exposed to Generative AI. Simultaneously, the nature of information industry jobs, which often allows for remote work and digital collaboration, may have mitigated certain challenges faced by industries requiring physical presence. On the other hand, sectors heavily reliant on in-person operations, such as manufacturing or hospitality, which are generally less exposed, may have encountered more pronounced disruptions in the labor market during COVID-19. To alleviate this concern of confounding factors, we conduct several tests and present results in Panel D, Table 12.

First, we keep only observations after 2019 and rerun the regression. The results in Column (1) show that even when we limit the sample period to the post-COVID period, there is still a significant difference in hiring quantity among firms with varying exposure levels. Second, in Column (2), we control for the interaction between an indicator variable that equals one in the post-COVID period (*Post-COVID*) and the state-level COVID deaths (*COVID Deaths*). Finally, we show that there is no significant correlation between a firm's Generative AI exposure and state-level COVID severity in Column (3).

7.5 Control for Effects of Other Technologies

Another potential concern is that our Generative AI exposure might be confounded with general Artificial Intelligence or machine learning. This could bias our results, as we might capture only effects from AI in a broad sense, not specifically from Generative AI. To address this potential concern, we control for the interaction of firm-level AI exposure and $Post\text{-}ChatGPT_t$ in our main regression.

Following Acemoglu et al. (2022), we use three indicators of firm-level Artificial Intelligence or Machine Learning exposure. Each was originally assigned at the six-digit SOC occupation level and we aggregate it into firm level using the same strategy as for Generative AI. Each indicator is specifically crafted to encompass occupations that heavily involve tasks compatible with the current capabilities of AI technologies. We present the results in Panel E, Table 12. In Column (1), we use the measure from Felten, Raj, and Seamans (2019). They utilize data from the AI Progress Measurement project, which identifies nine AI application areas. They then assess their relevance to 52 O*NET ability scales through crowdsourced evaluations. The resulting AI occupational impact for each O*NET occupation is determined by a weighted sum of the 52 AI application-ability scores. In Column (2), we use the measure from Webb (2019). The core idea of Webb's measure is to evaluate the capabilities of AI by pinpointing similarities between claims about AI capabilities in patents and job descriptions within O*NET. Occupations exhibiting a greater proportion of shared tasks are classified as more exposed to AI. The measure we use in Column (3) is SML from Brynjolfsson, Li, and Raymond (2023). They first construct a measure of task suitability for machine learning by creating a 23-item rubric. They then calculate Suitability for Machine Learning (SML) scores through its application to the textual descriptions of all O*NET occupations using the crowdsourcing platform. The results are presented in Panel E and show our main results are robust after controlling for alternative measures of AI exposure or machine learning exposure.

7.6 Other Robustness Tests

In addition, we conduct several other robustness checks for our findings. Importantly, our results are robust after excluding all samples in one year, all firms located in one state, and all firms belonging to one industry at a time, respectively. We plot the coefficients in Appendix A, Figure A.2. This test alleviates the concern that our results are mainly driven by specific years, states, or industries.

8 Conclusion

This paper examines the immediate responses of both employers and employees to the release of Generative AI. Using a generalized difference-in-differences design, we find that, in contrast of previous generations of automation, which primarily affect blue-collar workers (Acemoglu and Autor, 2011), Generative AI primarily affects white-collar jobs that require critical thinking and creativity, while blue-collar jobs remain largely unaffected. Firms with greater exposure to Generative AI significantly curtail their job postings after ChatGPT's release. Contrary to the conventional view of automation uniformly displacing workers, we find a *heterogeneous* effect on jobs that are highly exposed to the technology: Generative AI complements high-skilled white-collar roles (e.g., executive positions) but substitutes low-skilled white-collar roles (e.g., entry-level office positions). Moreover, we find that firms with greater exposure to Generative AI significantly increase emphasis on Generative AI and machine learning skills in job listings. Our findings suggest that heightened exposure to Generative AI does not inherently result in employee displacement. Instead, the upper echelons of the labor force distribution experience increased demand due to the evolving employment landscape.

Our study provides novel evidence of employee reactions to this technological disruption. We find that prospective employees are less likely to engage in job interviews post-ChatGPT, especially in non-management positions. For current employees we find that the overall Glassdoor ratings do not exhibit significant changes in response to Generative AI adoption. However, there is a noteworthy shift in current employees' future outlook about their firms, particularly at firms highly exposed to Generative AI following the ChatGPT release. The decline in future outlook suggests that while current employees may not perceive immediate disruptions, they are increasingly concerned about the future prospects of their organizations in light of Generative AI. Our results are robust to a variety of alternative specifications and potential alternative explanations.

Our findings highlight a mismatch between the reactions of employers and employees toward Generative AI. Employers are rapidly adapting by reducing job postings for standard white-collar positions and increasing the emphasis on AI-related skills, viewing Generative AI as a strategic tool to enhance efficiency and competitiveness. This is particularly pronounced in firms with higher R&D intensity and substantial cash reserves. On the other hand, employees express growing concerns about their long-term prospects, as evidenced by declines in CEO approval and business outlook ratings. While employers focus on immediate benefits such as cost savings and innovation, employees worry about job security, career growth, and organizational stability. This disconnect is further manifested in the reduced engagement of prospective employees, particularly in non-management roles, who may fear being replaced or marginalized by AI. The divergence in perspectives could pose challenges in aligning organizational strategies with employee expectations, potentially affecting workforce morale and trust in leadership.

We contribute to the emerging literature on the effects of automation by systematically examining the responses of both employers and employees to the release of Generative AI. Our study expands the understanding of how labor market participants react in response to significant technological shocks. To the best of our knowledge, our study is the first to provide large-scale evidence of the responses of both employers and employees to the introduction of Generative AI. We leverage high-frequency, forward-looking data that captures rapid shifts in job market demand at the occupation-week and firm-week levels. This approach can facilitate future research into the long-term consequences of Generative AI on labor market dynamics. Our findings provide timely insights for policymakers and regulators, especially in light of the rapid advancement of Generative AI and the increasing debate around AI-induced job displacement and regulatory scrutiny.

References

Acemoglu, D., Autor, D., 2011. Skills, tasks and technologies: Implications for employment and earnings. In: *Handbook of Labor Economics*, Elsevier, vol. 4, pp. 1043–1171.

Acemoglu, D., Autor, D., Hazell, J., Restrepo, P., 2022. Artificial intelligence and jobs: evidence from online vacancies. Journal of Labor Economics 40, S293–S340.

Acemoglu, D., Restrepo, P., 2020. Robots and jobs: Evidence from us labor markets. Journal of Political Economy 128, 2188–2244.

Acemoglu, D., Restrepo, P., 2022. Tasks, automation, and the rise in us wage inequality. Econometrica 90, 1973–2016.

Aihounton, G. B., Henningsen, A., 2021. Units of measurement and the inverse hyperbolic sine transformation. The Econometrics Journal 24, 334–351.

Baek, C., McCrory, P. B., Messer, T., Mui, P., 2021. Unemployment effects of stay-at-home orders: Evidence from high-frequency claims data. Review of Economics and Statistics 103, 979–993.

Barry, J. W., Campello, M., Graham, J. R., Ma, Y., 2022. Corporate flexibility in a time of crisis. Journal of Financial Economics 144, 780–806.

Basri, M. C., Felix, M., Hanna, R., Olken, B. A., 2021. Tax administration versus tax rates: evidence from corporate taxation in indonesia. American Economic Review 111, 3827–3871.

Bellemare, M. F., Wichman, C. J., 2020. Elasticities and the inverse hyperbolic sine transformation. Oxford Bulletin of Economics and Statistics 82, 50–61.

Bertomeu, J., Lin, Y., Liu, Y., Ni, Z., 2023. Capital market consequences of generative ai: Early evidence from the ban of chatgpt in italy. Available at SSRN 4452670.

Brynjolfsson, E., Li, D., Raymond, L. R., 2023. Generative AI at work. Working Paper 31161, National Bureau of Economic Research.

Chen, J., Roth, J., 2023. Logs with zeros? some problems and solutions. The Quarterly Journal of Economics p. qjad054.

Chetty, R., Friedman, J. N., Stepner, M., 2024. The economic impacts of covid-19: Evidence from a new public database built using private sector data. The Quarterly Journal of Economics 139, 829–889.

Cohn, J. B., Liu, Z., Wardlaw, M. I., 2022. Count (and count-like) data in finance. Journal of Financial Economics 146, 529–551.

Cowen, T., 2023. Chatgpt is also an impressive feat of marketing. Bloomberg. Available at: https://www.bloomberg.com/opinion/articles/2023-05-23/chatgpt-is-also-an-impressive-feat-of-marketing.

Duan, N., Manning, W. G., Morris, C. N., Newhouse, J. P., 1983. A comparison of alternative models for the demand for medical care. Journal of business & economic statistics 1, 115–126.

Eisfeldt, A. L., Schubert, G., Zhang, M. B., 2023. Generative AI and firm values. Working Paper 31222, National Bureau of Economic Research.

Elias, J., 2023. Google is asking employees to test potential chatgpt competitors, including a chatbot called ‘apprentice bard’. CNBC News. Available at: https://www.cnbc.com/2023/01/31/google-testing-chatgpt-like-chatbot-apprentice-bard-with-employees.html.

Eloundou, T., Manning, S., Mishkin, P., Rock, D., 2023. Gpts are gpts: An early look at the labor market impact potential of large language models. arXiv preprint arXiv:2303.10130 .

Felten, E., Raj, M., Seamans, R., 2021. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. Strategic Management Journal 42, 2195–2217.

Felten, E., Raj, M., Seamans, R. C., 2019. The effect of artificial intelligence on human labor: An ability-based approach. In: *Academy of Management Proceedings*, Academy of Management Briarcliff Manor, NY 10510, vol. 2019, p. 15784.

Giupponi, G., Landais, C., 2023. Subsidizing labour hoarding in recessions: the employment and welfare effects of short-time work. The Review of Economic Studies 90, 1963–2005.

Green, T. C., Huang, R., Wen, Q., Zhou, D., 2019. Crowdsourced employer reviews and stock returns. Journal of Financial Economics 134, 236–251.

Hainmueller, J., 2012. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. Political Analysis 20, 25–46.

Kim, A., Muhn, M., Nikolaev, V., 2024. Financial statement analysis with large language models. arXiv preprint arXiv:2407.17866 .

Kim, A. G., Muhn, M., Nikolaev, V. V., 2023a. Bloated disclosures: can chatgpt help investors process information? Chicago Booth Research Paper .

Kim, A. G., Muhn, M., Nikolaev, V. V., 2023b. From transcripts to insights: Uncovering corporate risks using generative ai. Chicago Booth Research Paper .

Korinek, A., Juelfs, M., 2024. Preparing for the (Non-Existent?) Future of Work. In: *The Oxford Handbook of AI Governance*, Oxford University Press.

Marinescu, I., Chamberlain, A., Smart, M., Klein, N., 2021. Incentives can reduce bias in online employer reviews. Journal of Experimental Psychology: Applied 27, 393–407.

Mckinsey, 2023. The economic potential of generative AI: The next productivity frontier. Mckinsey Digital. Available at: https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontierintroduction.

Modestino, A. S., Shoag, D., Ballance, J., 2020. Upskilling: Do employers demand greater skill when workers are plentiful? Review of Economics and Statistics 102, 793–805.

N'guessan, Y. G., Featherstone, A., Odeh, O., Upendram, S., 2017. Choice of the empirical definition of zero in the translog multiproduct cost functional form. Applied Economics Letters 24, 1112–1120.

Rosenblatt, K., 2023. ChatGPT passes MBA exam given by a Wharton professor. NBC News. Available at: https://www.nbcnews.com/tech/tech-news/chatgpt-passes-mba-exam-wharton-professor-rcna67036.

Webb, M., 2019. The impact of artificial intelligence on the labor market. Available at SSRN 3482150 .

Figure 1.

This figure presents the Generative AI exposure across states. The state-level exposure score is calculated by aggregating the occupation-level exposure score weighted by state occupational employment in 2022.

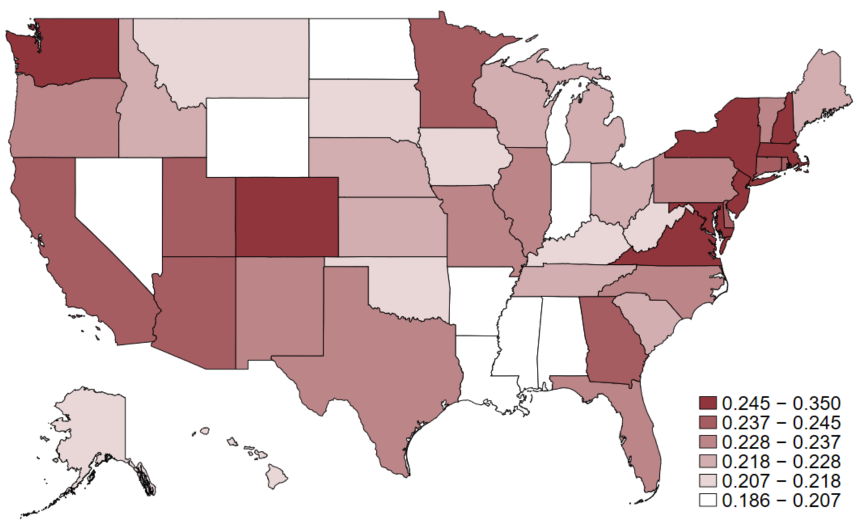

Figure 2.

This figure presents the Generative AI exposure score across occupation requirements for various knowledge elements, skill elements, and ability elements. The exposure score is the average occupation-level exposure score, weighted by each occupation's value across elements from O*NET 27.1.

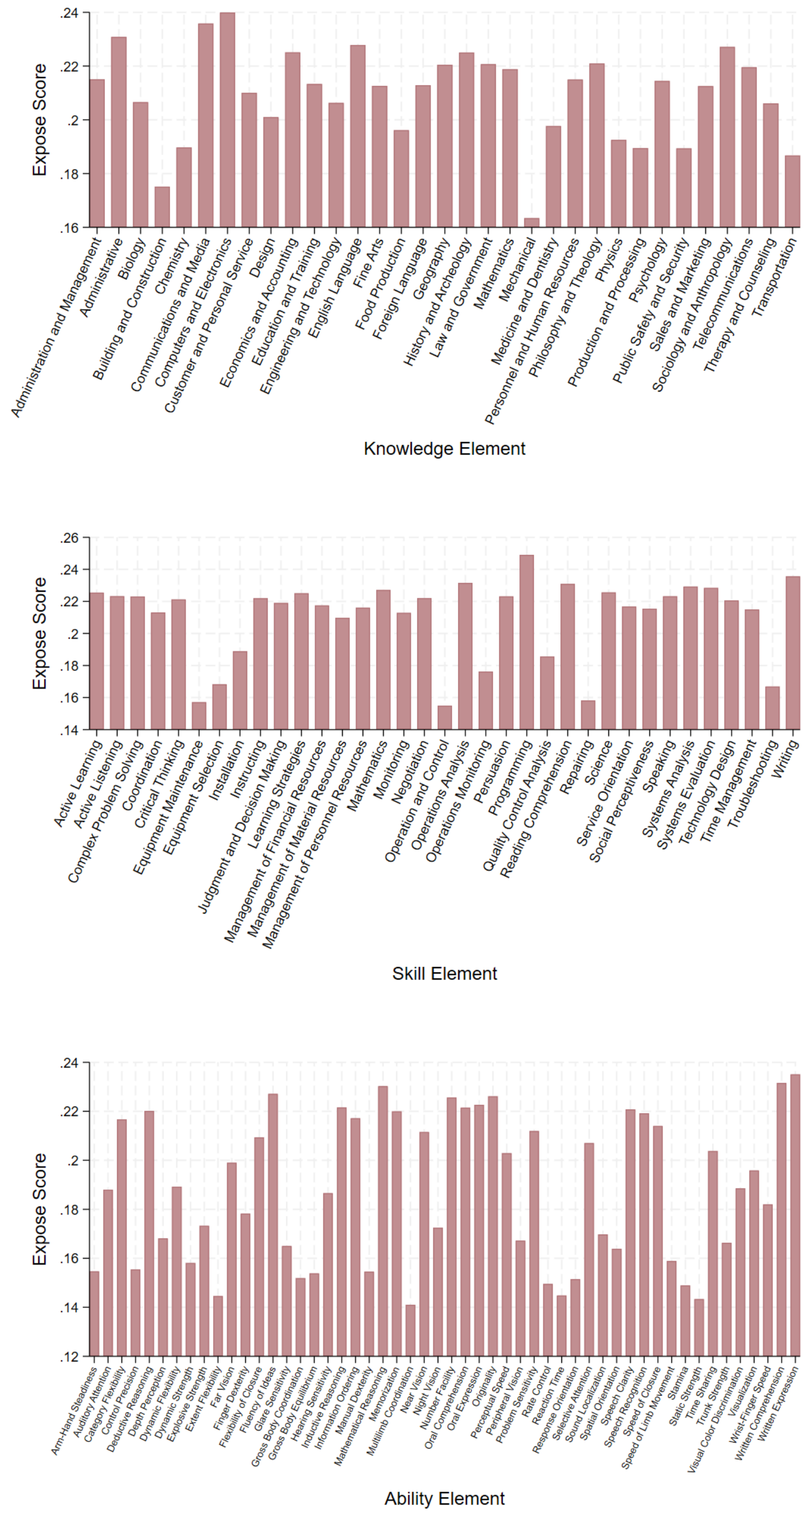

Table 1.

Descriptive Statistics

This table presents descriptive statistics for our sample. Panel A presents summary statistics of variables measured at the firm-week level from LinkUp. Panel B presents summary statistics of variables measured at the firm-week level from Compustats and CRSP. Panel C presents summary statistics of variables measured at the occupation-week level from LinkUp. Panel D and E present summary statistics of variables measured at the firm-week level from the Glassdoor interview and review, respectively. All variables are defined in Appendix A, Table A. 1. All continuous variables are winsorized at the 1% and 99% levels.

Panel A: LinkUp Summary Statistics (Firm-Week Level)

Variable	N	Mean	SD	P25	P50	P75
Log(Job Created + 1)	540575	2.535	1.875	1.099	2.485	3.829
Log(Machine Learning + 1)	354527	0.069	0.327	0	0	0
Log(Generative AI + 1)	354527	0.014	0.134	0	0	0
Post-ChatGPT	540575	0.072	0.258	0	0	0
Gen. AI Exposure	540575	0.370	0.139	0.269	0.377	0.469

Panel B: Compustat Summary Statistics (Firm-Year Level)

Variable	N	Mean	SD	P25	P50	P75
LgSize	540575	8.289	1.644	7.135	8.141	9.391
Tangibility	540575	0.465	0.390	0.160	0.336	0.702
Labor_Intensity	540575	-1.211	1.086	-1.831	-1.091	-0.353
MtB	540575	4.216	5.246	1.554	2.614	4.538
ROA	540575	0.052	0.082	0.019	0.051	0.091
CashHold	540575	0.147	0.150	0.038	0.094	0.205
SaleGrowth	540575	0.100	0.216	-0.002	0.069	0.158
StockReturn	540575	0.012	0.029	-0.005	0.012	0.028
R&D Intensity	330654	0.246	7.754	0.004	0.029	0.103
CorpAge	540200	31.450	19.030	17	26	46

Table 1. -Continued.

Panel C: LinkUp Summary Statistics (Occupation-Week Level)

Variable	N	Mean	SD	P25	P50	P75
High Knowledge Job Created	604920	1.014	1.738	0	0	1.609
Low Knowledge Job Created	604920	1.116	1.860	0	0	1.792
High/Low Ratio Knowledge Job Created	604920	35.66	146.9	0.167	1	5
High Skill Job Created	604920	1.155	1.885	0	0	1.792
Low Skill Job Created	604920	0.974	1.707	0	0	1.386
High/Low Ratio Skill Job Created	604920	50.70	203.9	0.25	1	6
High Education Job Created	604920	1.147	1.867	0	0	1.792
Low Education Job Created	604920	0.983	1.730	0	0	1.386
High/Low Ratio Education Job Created	604920	44.39	170.6	0.25	1	6
High Training Job Created	604920	1.017	1.740	0	0	1.609
Low Training Job Created	604920	1.112	1.859	0	0	1.792
High/Low Ratio Training Job Created	604920	36.09	150.0	0.167	1	5
High TechSkill Job Created	604920	1.294	2.042	0	0	2.197
Low TechSkill Job Created	604920	0.834	1.484	0	0	1.099
High/Low Ratio TechSkill Job Created	604920	69.22	261.3	0.333	1	9
JobZone 1 Job Created	604920	0.056	0.391	0	0	0
JobZone 2 Job Created	604920	0.605	1.415	0	0	0
JobZone 3 Job Created	604920	0.513	1.303	0	0	0
JobZone 4 Job Created	604920	0.687	1.640	0	0	0
JobZone 5 Job Created	604920	0.232	0.812	0	0	0
Gen. AI Exposure	604920	0.208	0.232	0	0.120	0.314

Panel D:Glassdoor Interview Summary Statistics (Firm-Week Level)

Variable	N	Mean	SD	P25	P50	P75
Interview Num	41618	2.862	3.671	1	1	3
Interview Num - Manager	41618	0.306	0.620	0	0	0
Interview Num - Non-Manager	41618	2.547	3.374	1	1	3
Experience Score	41618	0.443	0.669	0	0.667	1
Difficulty Score	41618	-0.182	0.519	-0.500	0	0

Table 1. -Continued.

Panel E: Glassdoor Review Summary Statistics (Firm-Week Level)

Variable	N	Mean	SD	P25	P50	P75
Approves of CEO	75965	0.385	0.613	0	0.500	1
Overall Rating	89180	3.591	0.941	3	3.714	4.143
Work/Life Balance	82773	2.723	1.654	1.500	3	4
Career Opportunities	83018	2.707	1.627	1.500	3	4
Compensation and Benefits	82900	2.798	1.625	2	3.250	4
Senior Management	82589	2.540	1.598	1	3	3.759
Culture & Values	82729	2.779	1.690	1.333	3.200	4
Diversity & Inclusion	82507	2.951	1.729	2	3.500	4.200
Recommend	79251	0.290	0.728	-0.091	0.429	1
Business Outlook	77565	0.375	0.631	0	0.500	1

Table 2.

Examples of GPT Scores Assigned to Task Statements and GPT-Provided Explanations

This Table presents examples of GPT scores assigned to task statements and GPT-provided explanations. Prompts used are illustrated in Appendix C. Occupation-task description comes from O*NET 27.1.

Occupation	Task	GPT Score	GPT Explanation	GPT Confidence
Advertising and Promotions Managers	Coordinate with the media to disseminate advertising.	E0	The task of coordinating with the media to disseminate advertising involves direct human interaction and decision-making that cannot be fully automated by the LLM. It requires negotiating and collaborating with media outlets, understanding target audiences, and making strategic decisions based on market trends and advertising goals.	High
Chief Executives	Interpret and explain policies, rules, regulations, or laws to organizations, government or corporate officials, or individuals.	E1	The task involves interpreting and explaining policies, rules, regulations, or laws, which can be done by providing written or spoken explanations. The LLM can assist in generating these explanations based on its knowledge and understanding of the subject matter.	High
Marketing Managers	Negotiate contracts with vendors or distributors to manage product distribution, establishing distribution networks or developing distribution strategies.	E0	Negotiating contracts with vendors or distributors involves direct human interaction and decision-making that cannot be fully automated by the LLM. It requires understanding the specific needs and requirements of both parties, as well as considering various factors such as pricing, terms, and conditions. The LLM can provide information and recommendations, but the final negotiation and agreement process would still require human involvement.	High
Financial Managers	Evaluate financial reporting systems, accounting or collection procedures, or investment activities and make recommendations for changes to procedures, operating systems, budgets, or other financial control functions.	E2	The task involves analyzing financial data and making recommendations, which can be done by leveraging the LLM's capabilities to analyze written information and provide counsel on financial matters. However, additional software could be developed on top of the LLM to retrieve up-to-date financial information and assist in analyzing complex financial data.	High
Online Merchants	Compose images of products, using video or still cameras, lighting equipment, props, or photo or video editing software.	E3	The task involves creating or editing digital images according to instructions, which can be significantly aided by the LLM's image capabilities. The model can provide guidance on composition, lighting, and editing techniques, and can even generate images based on textual descriptions.	High
Computer and Information Research Scientists	Apply theoretical expertise and innovation to create or apply new technology, such as adapting principles for applying computers to new uses.	E1	The model can learn theoretical expertise during training as part of its general knowledge base, and the principles to adapt can be captured in the text input to the model.	Medium
Barbers	Cut and trim hair according to clients' instructions or current hairstyles, using clippers, combs, hand-held blow driers, and scissors.	E0	This task requires physical interaction and expertise in using tools such as clippers, combs, blow dryers, and scissors. The LLM does not have the capability to perform physical tasks or manipulate objects.	High
Editors	Prepare, rewrite and edit copy to improve readability, or supervise others who do this work.	E1	The LLM can assist with tasks such as rewriting and editing copy, providing suggestions to improve readability, and even supervising others in this work. It can generate alternative phrasings, offer grammar and style corrections, and provide feedback on the overall quality of the text.	High

Table 3.

Highest and Lowest Generative AI Exposure Score Occupations

This Table presents the Generative AI exposure across different industries. The exposure score is the average of all companies in that industry.

SOC Code	Occupation Title	Exposure Score
15-1299.08	Computer Systems Engineers/Architects	0.964
27-3043.00	Writers and Authors	0.938
19-2041.01	Climate Change Policy Analysts	0.929
43-9111.00	Statistical Assistants	0.929
13-2082.00	Tax Preparers	0.917
15-1211.00	Computer Systems Analysts	0.909
43-9081.00	Proofreaders and Copy Markers	0.909
15-2031.00	Operations Research Analysts	0.906
31-9094.00	Medical Transcriptionists	0.900
43-5061.00	Production, Planning, and Expediting Clerks	0.882
15-1251.00	Computer Programmers	0.882
43-6013.00	Medical Secretaries and Administrative Assistants	0.875
15-1253.00	Software Quality Assurance Analysts and Testers	0.870
43-9022.00	Word Processors and Typists	0.850
19-3011.00	Economists	0.846
27-3091.00	Interpreters and Translators	0.824
43-6011.00	Executive Secretaries and Executive Administrative Assistants	0.818
43-4161.00	Human Resources Assistants, Except Payroll and Timekeeping	0.816
15-1254.00	Web Developers	0.815
15-1243.01	Data Warehousing Specialists	0.806
43-6014.00	Secretaries and Administrative Assistants, Except Legal, Medical, and Executive	0.797
15-2099.01	Bioinformatics Technicians	0.789
15-2041.00	Statisticians	0.789
15-1242.00	Database Administrators	0.786
43-9061.00	Office Clerks, General	0.786
19-3022.00	Survey Researchers	0.781
⋮	⋮	⋮
51-9192.00	Cleaning, Washing, and Metal Pickling Equipment Operators and Tenders	0.000
51-9195.00	Molders, Shapers, and Casters, Except Metal and Plastic	0.000
51-9195.03	Stone Cutters and Carvers, Manufacturing	0.000

Table 3 -Continued.

SOC Code	Occupation Title	Exposure Score
51-9195.04	Glass Blowers, Molders, Benders, and Finishers	0.000
51-9195.05	Potters, Manufacturing	0.000
51-9196.00	Paper Goods Machine Setters, Operators, and Tenders	0.000
51-9197.00	Tire Builders	0.000
53-2011.00	Airline Pilots, Copilots, and Flight Engineers	0.000
53-2012.00	Commercial Pilots	0.000
53-2031.00	Flight Attendants	0.000
53-3052.00	Bus Drivers, Transit and Intercity	0.000
53-4041.00	Subway and Streetcar Operators	0.000
53-5011.00	Sailors and Marine Oilers	0.000
53-5022.00	Motorboat Operators	0.000
53-5031.00	Ship Engineers	0.000
53-6011.00	Bridge and Lock Tenders	0.000
53-6021.00	Parking Attendants	0.000
53-7021.00	Crane and Tower Operators	0.000
53-7031.00	Dredge Operators	0.000
53-7041.00	Hoist and Winch Operators	0.000
53-7051.00	Industrial Truck and Tractor Operators	0.000
53-7061.00	Cleaners of Vehicles and Equipment	0.000
53-7062.04	Recycling and Reclamation Workers	0.000
53-7063.00	Machine Feeders and Offbearers	0.000
53-7065.00	Stockers and Order Fillers	0.000
53-7073.00	Wellhead Pumpers	0.000
53-7081.00	Refuse and Recyclable Material Collectors	0.000

Table 4.

Highest and Lowest Generative AI Exposure Score Firms

This Table presents the most exposed and least exposed firms by Generative AI. The exposure score is calculated by the weighted average of occupation-level exposure score using the number of employees hired from January 1, 2022.

Company Name	NAICS Industry	State	Exposure Score
Perficient Inc	Professional, Scientific, and Technical Services	Missouri	0.756
Black Knight Inc	Information	Florida	0.735
Dxc Technology Co	Professional, Scientific, and Technical Services	Virginia	0.726
Kyndryl Holdings Inc	Professional, Scientific, and Technical Services	New York	0.715
Synchronoss Technologies	Information	New Jersey	0.714
Eastman Kodak Co	Manufacturing	New York	0.712
Accenture Plc	Professional, Scientific, and Technical Services		0.688
Epam Systems Inc	Professional, Scientific, and Technical Services	Pennsylvania	0.676
Iridium Communications Inc	Information	Virginia	0.673
Federal National Mortga Assn	Finance and Insurance	District of Columbia	0.672
Verisign Inc	Information	Virginia	0.668
Gartner Inc	Professional, Scientific, and Technical Services	Connecticut	0.667
Ion Geophysical Corp	Mining, Quarrying, and Oil and Gas Extraction	Texas	0.667
Akamai Technologies Inc	Information	Massachusetts	0.657
Ptc Inc	Information	Massachusetts	0.652
⋮	⋮	⋮	⋮
Olin Corp	Manufacturing	Missouri	0.025
Ryder System Inc	Real Estate and Rental and Leasing	Florida	0.023
Warrior Met Coal Inc	Mining, Quarrying, and Oil and Gas Extraction	Alabama	0.019
Roadrunner Trans Systems Inc	Transportation and Warehousing	California	0.016
Group 1 Automotive Inc	Retail Trade	Texas	0.002
Advanced Energy Inds Inc	Manufacturing	Colorado	0.000
Amgen Inc	Manufacturing	California	0.000
Cerner Corp	Professional, Scientific, and Technical Services	Missouri	0.000
First Midwest Bancorp Inc	Finance and Insurance	Illinois	0.000
Sanmina Corp	Manufacturing	California	0.000
Summit Hotel Properties Inc	Real Estate and Rental and Leasing	Texas	0.000
Tidewater Inc	Transportation and Warehousing	Texas	0.000
Trustco Bank Corp/Ny	Finance and Insurance	New York	0.000
Two Harbors Investment Corp	Finance and Insurance	Minnesota	0.000

Table 5.

Generative AI Exposure Score across Industries

This Table presents the Generative AI exposure across different industries. The exposure score is the average of all companies in that industry. The most exposed two industries are Professional, Scientific, and Technical Services, and Information. The least exposed two industries are Accommodation and Food Services, and Retail Trade.

NAICS Code	NAICS Industry	Number of Companies	Exposure Score
11	Agriculture, Forestry, Fishing and Hunting	2	0.38
21	Mining, Quarrying, and Oil and Gas Extraction	56	0.318
22	Utilities	45	0.353
23	Construction	25	0.252
31-33	Manufacturing	536	0.375
42	Wholesale Trade	38	0.347
44-45	Retail Trade	67	0.231
48-49	Transportation and Warehousing	40	0.312
51	Information	118	0.501
52	Finance and Insurance	203	0.42
53	Real Estate and Rental and Leasing	88	0.316
54	Professional, Scientific, and Technical Services	41	0.512
56	Administrative and Support and Waste Management and Remediation Services	29	0.359
61	Educational Services	9	0.388
62	Health Care and Social Assistance	25	0.315
71	Arts, Entertainment, and Recreation	6	0.283
72	Accommodation and Food Services	35	0.155
81	Other Services (except Public Administration)	5	0.279

Table 6.

Employers' Hiring Quantity in Response to Exposure to Generative AI

This table presents an analysis of employer's hiring response to Generative AI exposure. Panel A presents results using the full sample and various fixed effects structures. Panel B presents results using an alternative model and dependent variable specifications. Panel A uses the natural logarithm of one plus weekly new job posted (*log(1+Job Created)*) as the dependent variable. Panel B, Column 1 uses a Poisson regression. The dependent variables are, respectively, the weekly new job posted(*Job Created*) in columns 1 and 2, the natural logarithm of weekly new job posted (*log(Job Created)*) in Column 3, the Inverse hyperbolic sine transformation of weekly new job posted (*IHS Job Created*) in Column 4, and an indicator variable for whether new job posted is non-zero (*Job Created>0*) in Column 5. All continuous variables are winsorized at the 1% and 99% levels. t-statistics based on standard errors clustered by firm are in parentheses. Levels of significance are presented as follows: *p<0.1; **p<0.05; ***p<0.01.

Panel A: Main Results

	(1)	(2)	(3)	(4)	(5)
Variables	Log (Job Created + 1)				
Post-ChatGPT × Gen. AI Exposure	-1.087***	-1.229***	-1.310***	-1.309***	-1.325***
	(-4.37)	(-6.68)	(-7.28)	(-7.27)	(-6.31)
Post-ChatGPT	0.933***			-0.039	
	(9.18)			(-0.56)	
Gen. AI Exposure	-1.473***				-1.603***
	(-4.22)				(-5.08)
LgSize			0.534***	0.533***	0.681***
			(10.03)	(10.02)	(35.37)
Tangibility			-0.191	-0.191	-1.300***
			(-0.80)	(-0.80)	(-7.73)
Labor_Intensity			0.305***	0.304***	0.474***
			(3.99)	(3.99)	(8.23)
Market-to-Book			0.010***	0.010***	0.029***
			(2.61)	(2.60)	(7.70)
ROA			0.185	0.184	-0.084
			(1.18)	(1.18)	(-0.35)
CashHold			0.013	0.012	0.308*
			(0.09)	(0.08)	(1.71)
SaleGrowth			0.062	0.062	-0.061
			(1.48)	(1.50)	(-0.96)
StockReturn			0.113	0.117	-0.303
			(0.40)	(0.41)	(-0.79)
Constant	3.042***	2.567***	-1.459***	-1.450***	-1.463***
	(22.14)	(533.56)	(-3.02)	(-3.01)	(-6.35)
Observations	540,575	540,575	540,575	540,575	540,575
Adjusted R^2	0.017	0.702	0.708	0.690	0.501
Mean of Dep. Var	2.535	2.535	2.535	2.535	2.535
Firm FE		✓	✓	✓	
Year#Week FE		✓	✓		✓
Industry FE					✓
Year FE				✓	

Table 6. -Continued.

Panel B: Alternative specifications

	(1)	(2)	(3)	(4)	(5)
	Poisson	OLS			
Variables	Job Created	Job Created	Log (Job Created)	IHS Job Created	Job Created>0
Post-ChatGPT × Gen. AI Exposure	-0.549***	-146.584***	-1.034***	-1.430***	-0.115***
	(-3.39)	(-4.78)	(-6.44)	(-7.13)	(-3.10)
LgSize	0.586***	29.251***	0.524***	0.597***	0.061***
	(7.31)	(3.69)	(10.69)	(9.90)	(4.18)
Tangibility	-0.208	-5.972	-0.398*	-0.209	0.011
	(-0.59)	(-0.23)	(-1.66)	(-0.78)	(0.18)
Labor_Intensity	0.233**	23.906***	0.303***	0.341***	0.040*
	(1.96)	(2.63)	(4.27)	(3.94)	(1.89)
Market-to-Book	-0.000	1.048*	0.012***	0.010**	-0.002**
	(-0.11)	(1.94)	(3.60)	(2.30)	(-2.25)
ROA	0.390	-19.593	0.229	0.222	0.043
	(1.26)	(-1.04)	(1.55)	(1.26)	(0.98)
CashHold	0.023	2.517	-0.181	0.011	0.035
	(0.08)	(0.13)	(-1.23)	(0.07)	(0.74)
SaleGrowth	-0.053	-7.046	0.020	0.081*	0.035***
	(-0.83)	(-1.61)	(0.47)	(1.70)	(2.84)
StockReturn	-0.108	24.273	0.164	0.120	-0.027
	(-0.22)	(0.69)	(0.59)	(0.37)	(-0.32)
Constant	0.123	-138.389**	-0.927**	-1.458***	0.346***
	(0.16)	(-2.06)	(-2.04)	(-2.68)	(2.71)
Observations	538,166	540,575	438,117	540,575	540,575
Pseudo/Adjusted R^2	0.777	0.572	0.748	0.693	0.350
Mean of Dep. Var	72.15	71.85	2.998	3.006	0.810
Firm FE	✓	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓	✓

Table 7.

Heterogeneous Effects on Hiring Quantity across Occupation Characteristics

This table presents the analysis of heterogeneous employers' hiring response to Generative AI exposure across occupation characteristics. All continuous variables are winsorized at the 1% and 99% levels. t-statistics based on standard errors clustered by occupation are in parentheses. Levels of significance are presented as follows: *p<0.1; **p<0.05; ***p<0.01.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Knowledge			Skill			Education		
Variables	High	Low	High/Low Ratio	High	Low	High/Low Ratio	High	Low	High/Low Ratio
Post-ChatGPT × Gen. AI Exposure	0.144**	-0.315***	43.203***	0.141**	-0.301***	69.121***	0.329***	-0.499***	81.743***
	(2.50)	(-5.07)	(3.25)	(2.24)	(-5.31)	(3.57)	(5.36)	(-8.79)	(4.99)
Constant	1.012***	1.119***	35.184***	1.153***	0.978***	49.934***	1.143***	0.989***	43.478***
	(1,581.60)	(1,616.06)	(237.63)	(1,648.92)	(1,549.75)	(231.52)	(1,671.18)	(1,564.41)	(238.53)
Observations	604,920	604,920	604,920	604,920	604,920	604,920	604,920	604,920	604,920
Adjusted R^2	0.917	0.915	0.783	0.926	0.904	0.765	0.929	0.902	0.782
Mean of Dep. Var	1.014	1.116	35.67	1.155	0.974	50.70	1.147	0.983	44.39
Occupation FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 7. -Continued.

Variables	(10) High	(11) Low	(12) High/Low Ratio	(13) High	(14) Low	(15) High/Low Ratio	(16) Zone 1	(17) Zone 2	(18) Zone 3	(19) Zone 4	(20) Zone 5
	Training			Technical Skill			Job Zone				
Post-ChatGPT × Gen. AI Exposure	-0.077	-0.093	33.925***	0.362***	-0.508***	98.677***	-0.022**	-0.268***	-0.161***	0.257***	0.059*
	(-1.56)	(-1.36)	(2.96)	(5.86)	(-9.71)	(4.68)	(-2.23)	(-6.05)	(-4.45)	(5.34)	(1.87)
Constant	1.018***	1.113***	35.716***	1.290***	0.840***	68.119***	0.056***	0.608***	0.515***	0.684***	0.231***
	(1,855.21)	(1,470.98)	(279.36)	(1,876.81)	(1,441.91)	(290.03)	(512.87)	(1,232.02)	(1,277.36)	(1,279.63)	(655.46)
Observations	604,920	604,920	604,920	604,920	604,920	604,920	604,920	604,920	604,920	604,920	604,920
Adjusted R^2	0.919	0.914	0.794	0.941	0.867	0.765	0.909	0.905	0.924	0.954	0.873
Mean of Dep. Var	1.017	1.112	36.09	1.294	0.834	69.22	0.0561	0.605	0.513	0.687	0.232
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 8.

Job Description in Response to Exposure to Generative AI

This table presents an analysis of changes in the job description in response to Generative AI exposure. Columns 1 to 3 present the changes in machine-learning-related keywords used in job descriptions. Columns 4 to 6 present the changes in Generative-AI-related keywords used in job descriptions. The keywords used are listed in Appendix A, Table A. 2. The dependent variables are, respectively, the ratio of the count of machine-learning-related keywords in job descriptions to the total number of jobs posted (*Machine Learning Ratio*) in Column 1, an indicator variable that equals one when the company posted new job vacancies that contain machine-learning-related keywords (*Machine Learning Dummy*) in Column 2, the natural logarithm of one plus the count of machine-learning-related keywords in job descriptions (*Log(Machine Learning + 1)*) in Column 3, the ratio of the count of generative-AI-related keywords in job descriptions to the total number of job posted (*Generative AI Ratio*) in Column 4, an indicator variable that equals one when the company posted new job vacancies that contain generative-AI-related keywords (*Generative AI Dummy*) in Column 5, and the natural logarithm of one plus the count of generative-AI-related keywords in job descriptions (*Log(Generative AI + 1)*) in Column 6. All continuous variables are winsorized at the 1% and 99% levels. t-statistics based on standard errors clustered by firm are in parentheses. Levels of significance are presented as follows: *p<0.1; **p<0.05; ***p<0.01.

Variables	(1) Machine Learning Ratio	(2) Machine Learning Dummy	(3) Log(Machine Learning + 1)	(4) Generative AI Ratio	(5) Generative AI Dummy	(6) Log(Generative AI + 1)
Post-ChatGPT × Gen. AI Exposure	0.010***	0.115***	0.163***	0.006***	0.055***	0.091***
	(4.81)	(4.73)	(4.37)	(2.93)	(3.40)	(3.21)
LgSize	0.002	0.018	0.023	0.000**	0.002	0.003
	(1.63)	(1.49)	(1.09)	(2.26)	(0.65)	(0.63)
Tangibility	0.001	-0.064**	-0.092*	-0.001	-0.014	-0.012
	(0.53)	(-2.05)	(-1.72)	(-0.81)	(-1.47)	(-1.01)
Labor_Intensity	-0.000	0.023	0.037	0.000	0.005	0.007
	(-0.25)	(1.46)	(1.32)	(1.38)	(0.93)	(0.96)
Market-to-Book	0.000	0.003***	0.005***	0.000**	0.001***	0.002***
	(1.59)	(3.68)	(3.44)	(2.56)	(2.75)	(2.72)
ROA	0.000	0.038	0.056	-0.002*	0.012	0.021
	(0.03)	(1.45)	(1.28)	(-1.96)	(0.98)	(1.27)
CashHold	-0.002	-0.109***	-0.177***	-0.000	-0.036**	-0.058**
	(-0.74)	(-2.64)	(-2.62)	(-0.27)	(-2.32)	(-2.52)
SaleGrowth	-0.000	-0.008	-0.012	-0.000	-0.004*	-0.006*
	(-0.79)	(-1.23)	(-1.09)	(-0.82)	(-1.65)	(-1.92)
StockReturn	-0.006	-0.073*	-0.154**	-0.001	-0.048**	-0.076***
	(-1.20)	(-1.89)	(-2.31)	(-0.58)	(-2.53)	(-2.65)
Constant	-0.014	-0.047	-0.038	-0.003	0.003	0.001
	(-1.28)	(-0.42)	(-0.20)	(-1.62)	(0.10)	(0.03)
Observations	354,493	354,493	354,493	354,493	354,493	354,493
Adjusted R^2	0.054	0.314	0.349	0.028	0.123	0.129
Mean of Dep. Var	0.00220	0.0501	0.0687	0.000420	0.0131	0.0136
Firm FE	✓	✓	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓	✓	✓

Table 9.

Cross-Sectional Analysis

This table presents the cross-sectional analysis of employer's response to Generative AI exposure. Dependent variables are, respectively, the natural logarithm of one plus weekly new job posted (*log(1+Job Created)*) in Panel A, the natural logarithm of one plus the count of machine-learning-related keywords in job descriptions (*Log(Machine Learning + 1)*) in Panel B, and the natural logarithm of one plus the count of generative-AI-related keywords in job descriptions (*Log(Generative AI + 1)*) in Panel C. *LgSize* is the natural log of total asset (in millions). *R&D Intensity* is research and development expense divided by total sales. *CashHold* is cash and short-term investments divided by total assets. *CorpAge* is years since corporation foundation. All continuous variables are winsorized at the 1% and 99% levels. t-statistics based on standard errors clustered by firm are in parentheses. Levels of significance are presented as follows: *p<0.1; **p<0.05; ***p<0.01.

Panel A: Hiring Quantity

	(1)	(2)	(3)	(4)
Variables	Log(Job Created + 1)			
LogSize × Post-ChatGPT × Gen. AI Exposure	0.082**			
	(2.33)			
R&D Intensity × Post-ChatGPT × Gen. AI Exposure		-1.254***		
		(-2.62)		
CashHold × Post-ChatGPT × Gen. AI Exposure			-1.396***	
			(-4.31)	
CorpAge × Post-ChatGPT × Gen. AI Exposure				0.001
				(0.23)
Observations	540,575	330,654	540,575	540,200
Adjusted R^2	0.708	0.719	0.708	0.707
Mean of Dep. Var	2.535	2.536	2.535	2.534
Controls	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓

Table 9. -Continued.

Panel B: Machine Learning Mentioned in Job Descriptions

	(1)	(2)	(3)	(4)
Variables	Log(Machine Learning + 1)			
LogSize × Post-ChatGPT × Gen. AI Exposure	0.071***			
	(6.24)			
R&D Intensity × Post-ChatGPT × Gen. AI Exposure		0.110		
		(0.88)		
CashHold × Post-ChatGPT × Gen. AI Exposure			-0.066	
			(-0.84)	
CorpAge × Post-ChatGPT × Gen. AI Exposure				-0.000
				(-0.23)
Observations	354,493	215,149	354,493	354,187
Adjusted R^2	0.350	0.381	0.349	0.349
Mean of Dep. Var	0.0687	0.0790	0.0687	0.0686
Controls	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓

Panel C: Generative AI Mentioned in Job Descriptions

	(1)	(2)	(3)	(4)
LogSize × Post-ChatGPT × Gen. AI Exposure	0.042***			
	(4.02)			
R&D Intensity × Post-ChatGPT × Gen. AI Exposure		0.204		
		(1.60)		
CashHold × Post-ChatGPT × Gen. AI Exposure			0.124	
			(1.54)	
CorpAge × Post-ChatGPT × Gen. AI Exposure				0.000
				(0.16)
Observations	354,493	215,149	354,493	354,187
Adjusted R^2	0.132	0.153	0.129	0.129
Mean of Dep. Var	0.0136	0.0172	0.0136	0.0136
Controls	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓

Table 10.

Potential Employee's Response to Exposure to Generative AI

This table presents an analysis of potential employees' response to Generative AI exposure using the Glassdoor Interview dataset. The dependent variables are, respectively, the total number of interview experiences posted (*Interview Num*) in Column 1, the total number of interview experiences on manager positions posted (*Interview Num - Manager*) in Column 2, the total number of interview experiences on non-manager positions posted (*Interview Num – Non-Manager*) in Column 3, experience score (*Experience Score*) in Column 4, and difficulty score (*Difficulty Score*) in Column 5. All continuous variables are winsorized at the 1% and 99% levels. t-statistics based on standard errors clustered by firm are in parentheses. Levels of significance are presented as follows: *p<0.1; **p<0.05; ***p<0.01.

Variables	(1) Interview Num	(2) Interview Num - Manager	(3) Interview Num - Non-Manager	(4) Experience Score	(5) Difficulty Score
Post-ChatGPT × Gen. AI Exposure	-0.692**	-0.051	-0.619**	-0.021	0.116***
	(-2.20)	(-1.10)	(-2.17)	(-0.37)	(3.20)
LgSize	0.414	-0.034	0.431	-0.008	0.039
	(1.34)	(-0.59)	(1.52)	(-0.10)	(0.79)
Tangibility	0.966	0.430***	0.482	-0.313	-0.121
	(0.71)	(2.67)	(0.37)	(-1.53)	(-0.74)
Labor_Intensity	-0.014	-0.189**	0.169	0.188**	0.048
	(-0.04)	(-2.55)	(0.49)	(1.97)	(0.86)
Market-to-Book	0.001	0.001	0.000	-0.000	-0.001*
	(0.17)	(1.05)	(0.02)	(-0.06)	(-1.75)
ROA	0.906	0.074	0.781	-0.118	0.136
	(1.60)	(0.65)	(1.45)	(-0.85)	(1.04)
CashHold	0.620	0.062	0.617	0.108	-0.207**
	(1.02)	(0.61)	(1.05)	(0.62)	(-2.07)
SaleGrowth	0.031	0.018	0.011	0.025	-0.030
	(0.19)	(0.68)	(0.07)	(0.77)	(-1.21)
StockReturn	0.693	-0.014	0.721	-0.148	0.169
	(0.65)	(-0.07)	(0.72)	(-0.67)	(0.98)
Constant	-1.525	0.179	-1.550	0.887	-0.420
	(-0.51)	(0.37)	(-0.56)	(1.33)	(-0.92)
Observations	41,575	41,575	41,575	41,575	41,575
Adjusted R^2	0.808	0.302	0.794	0.070	0.140
Mean of Dep. Var	2.864	0.306	2.549	0.443	-0.182
Firm FE	✓	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓	✓

Table 11.

Current Employees' Response to Exposure to Generative AI

This Table presents the changes in employees' Glassdoor review in response to firm-level exposure to Generative AI. Panel A illustrates the treatment effects on the whole group, Panel B illustrates the effects on the subgroups of senior employees (work experience of more than five years), and Panel C illustrates the effects on the subgroups of junior employees (work experience of equal or less than five years). All continuous variables are winsorized at the 1% and 99% levels. t-statistics based on standard errors clustered by firm are in parentheses. Levels of significance are presented as follows: *p<0.1; **p<0.05; ***p<0.01.

Panel A: All Employees

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Variables	Approves of CEO	Overall Rating	Work/Life Balance	Career Opportunities	Compensation and Benefits	Senior Management	Culture & Values	Diversity & Inclusion	Recommend	Business Outlook
Post-ChatGPT × Gen. AI Exposure	-0.113***	-0.052	-0.037	-0.033	-0.004	-0.074	-0.058	-0.070	-0.009	-0.168***
	(-2.62)	(-0.88)	(-0.67)	(-0.57)	(-0.09)	(-1.16)	(-0.99)	(-1.29)	(-0.20)	(-3.78)
Observations	75,957	89,172	82,764	83,009	82,892	82,580	82,720	82,498	79,243	77,557
Adjusted R^2	0.147	0.125	0.719	0.717	0.777	0.666	0.717	0.748	0.110	0.119
Mean of Dep. Var	0.385	3.591	2.723	2.706	2.798	2.540	2.779	2.951	0.290	0.375
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 11. -Continued.

Panel B: Senior Employees

Variables	(1) Approves of CEO	(2) Overall Rating	(3) Work/Life Balance	(4) Career Opportunities	(5) Compensation and Benefits	(6) Senior Management	(7) Culture & Values	(8) Diversity & Inclusion	(9) Recommend	(10) Business Outlook
Post-ChatGPT × Gen. AI Exposure	-0.028	0.005	-0.025	-0.026	0.077	-0.021	-0.035	-0.049	0.086	-0.161***
	(-0.51)	(0.08)	(-0.36)	(-0.37)	(1.28)	(-0.26)	(-0.49)	(-0.78)	(1.37)	(-2.78)
Observations	49,252	68,655	56,343	56,764	56,557	56,142	56,333	56,103	51,463	50,185
Adjusted R^2	0.135	0.106	0.685	0.693	0.747	0.633	0.693	0.725	0.101	0.113
Mean of Dep. Var	0.396	3.689	2.726	2.756	2.825	2.563	2.820	2.987	0.345	0.409
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Panel C: Junior Employees

Variables	(1) Approves of CEO	(2) Overall Rating	(3) Work/Life Balance	(4) Career Opportunities	(5) Compensation and Benefits	(6) Senior Management	(7) Culture & Values	(8) Diversity & Inclusion	(9) Recommend	(10) Business Outlook
Post-ChatGPT × Gen. AI Exposure	-0.169***	-0.028	-0.058	-0.063	-0.067	-0.109	-0.094	-0.091	-0.025	-0.156***
	(-3.55)	(-0.39)	(-0.91)	(-0.95)	(-1.09)	(-1.48)	(-1.37)	(-1.46)	(-0.49)	(-3.19)
Observations	62,294	71,734	68,815	68,953	68,922	68,644	68,780	68,587	65,855	63,893
Adjusted R^2	0.136	0.116	0.689	0.687	0.747	0.633	0.687	0.722	0.104	0.110
Mean of Dep. Var	0.386	3.522	2.727	2.694	2.778	2.538	2.772	2.958	0.270	0.361
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 12.
Robustness Tests
This table presents robustness tests for the baseline estimates presented in Table 6. Panel A shows results using firm-month level data. Panel B presents results using alternative independent variables. Panel C presents results using entropy balancing. Panel D presents results excluding effects from COVID. Panel E presents results controlling effects from Artificial Intellience in a broad sense. All continuous variables are winsorized at the 1% and 99% levels. t-statistics based on standard errors clustered by firm are in parentheses. Levels of significance are presented as follows: *p<0.1; **p<0.05; ***p<0.01.

Panel A: Month-Level Results

	(1)	(2)	(3)	(4)	(5)
Variables	Log (Job Created + 1)				
Post-ChatGPT × Gen. AI Exposure	-1.109***	-1.297***	-1.376***	-1.375***	-1.377***
	(-3.78)	(-5.94)	(-6.42)	(-6.41)	(-5.55)
Post-ChatGPT	1.044***			0.080	
	(8.68)			(0.96)	
Gen. AI Exposure	-1.642***				-1.627***
	(-4.10)				(-4.41)
Observations	124,565	124,563	124,563	124,563	124,565
Adjusted R^2	0.018	0.699	0.705	0.699	0.488
Mean of Dep. Var	3.867	3.867	3.867	3.867	3.867
Controls			✓	✓	✓
Firm FE		✓	✓	✓	
Year#Month FE		✓	✓		✓
Industry FE					✓
Year FE				✓	

Table 12. -Continued.

Panel B: Use Alternative Independent Variables

	(1)	(2)	(3)	(4)
Variables	Log (Job Created + 1)			
Post-ChatGPT × Gen. AI Exposure (1-month-Avg)	-1.052***			
	(-6.81)			
Post-ChatGPT × Gen. AI Exposure (6-month-Avg)		-1.196***		
		(-6.47)		
Post-ChatGPT × Gen. AI Exposure (18-month-Avg)			-1.267***	
			(-6.87)	
Post-ChatGPT × Gen. AI Exposure(2-year-Avg)				-1.260***
				(-6.81)
LgSize	0.525***	0.524***	0.535***	0.531***
	(9.76)	(9.88)	(10.21)	(10.24)
Tangibility	-0.168	-0.163	-0.168	-0.162
	(-0.68)	(-0.68)	(-0.71)	(-0.69)
Labor_Intensity	0.278***	0.279***	0.298***	0.298***
	(3.52)	(3.68)	(3.97)	(4.01)
Market-to-Book	0.009**	0.009**	0.010***	0.010**
	(2.29)	(2.34)	(2.60)	(2.58)
ROA	0.154	0.184	0.190	0.195
	(0.94)	(1.16)	(1.23)	(1.28)
CashHold	-0.041	-0.031	-0.009	0.000
	(-0.26)	(-0.21)	(-0.06)	(0.00)
SaleGrowth	0.061	0.061	0.058	0.060
	(1.40)	(1.44)	(1.39)	(1.45)
StockReturn	0.242	0.184	0.121	0.088
	(0.83)	(0.65)	(0.43)	(0.31)
Constant	-1.411***	-1.405***	-1.492***	-1.466***
	(-2.86)	(-2.90)	(-3.14)	(-3.12)
Observations	519,623	533,768	550,484	559,201
Adjusted R^2	0.706	0.706	0.708	0.708
Mean of Dep. Var	2.575	2.553	2.518	2.507
Controls	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
Year#Week FE	✓	✓	✓	✓

Table 12. -Continued.

Panel C: Entropy Balancing

Panel C.1: Descriptive Statistics of Firm Characteristics Before Entropy Balancing

	Treat		Control		Diff
	Mean	Variance	Mean	Variance	Diff (High-Low)
LgSize	8.277	3.163	8.300	2.244	-0.022426***
Tangibility	0.335	0.113	0.596	0.157	-0.261085***
Labor_Intensity	-1.601	1.214	-0.823	0.842	-0.777243***
MtB	4.875	35.253	3.56	18.958	1.314751***
ROA	0.052	0.008	0.052	0.005	-0.000112
CashHold	0.191	0.029	0.103	0.012	0.087829***
SaleGrowth	0.114	0.051	0.085	0.042	0.0285459***
StockReturn	0.012	0.001	0.012	0.001	-0.000176**

Panel C.2: Descriptive Statistics of Firm Characteristics After Entropy Balancing

	Treat		Control		Diff
	Mean	Variance	Mean	Variance	Diff (High-Low)
LgSize	8.277	3.163	8.278	3.162	-0.00000874
Tangibility	0.335	0.113	0.335	0.113	-0.0000806
Labor_Intensity	-1.601	1.214	-1.600	1.214	-0.0001953
MtB	4.875	35.25	4.875	35.247	.0001692
ROA	0.052	0.008	0.052	0.008	0.000000327
CashHold	0.191	0.029	0.191	0.029	.0000122
SaleGrowth	0.114	0.051	0.114	0.051	0.00000494
StockReturn	.012	0.001	0.012	0.001	-0.000000123

Panel C.3: Hiring Quantity

	(1)	(2)	(3)	(4)	(5)
Variables	Log (Job Created + 1)				
Post-ChatGPT × Gen. AI Exposure	-1.052***	-1.099***	-1.178***	-1.179***	-1.120***
	(-2.82)	(-4.85)	(-5.20)	(-5.21)	(-3.64)
Post-ChatGPT	0.940***			-0.104	
	(5.28)			(-1.14)	
Gen. AI Exposure	-1.828***				-1.759***
	(-3.49)				(-3.91)
Observations	540,575	540,575	540,575	540,575	540,575
Adjusted R^2	0.022	0.731	0.737	0.721	0.548
Mean of Dep. Var	2.570	2.570	2.570	2.570	2.570
Controls			✓	✓	✓
Firm FE		✓	✓	✓	
Year#Week FE		✓	✓		✓
Industry FE					✓
Year FE				✓	

Table 12. -Continued.

Panel D: Control for COVID Effects

	(1)	(2)	(3)
Model	After 2019	Control COVID Exposure	Correlation
Variables	Log(Job Created + 1)	Log(Job Created + 1)	Log(COVID Deaths)
Post-ChatGPT × Gen. AI Exposure	-0.801***	-1.330***	
	(-5.09)	(-7.21)	
Post-COVID × Log(COVID Deaths)		-0.038*	
		(-1.91)	
Gen. AI Exposure			0.050
			(0.16)
LgSize	0.211*	0.539***	0.091***
	(1.86)	(10.01)	(3.51)
Tangibility	-0.343	-0.174	0.094
	(-0.93)	(-0.71)	(0.46)
Labor_Intensity	0.167	0.306***	-0.062
	(1.57)	(3.88)	(-0.89)
Market-to-Book	0.003	0.009**	-0.000
	(0.72)	(2.48)	(-0.05)
ROA	0.694***	0.174	-0.912**
	(3.56)	(1.10)	(-2.00)
CashHold	-0.121	-0.004	1.003***
	(-0.59)	(-0.03)	(3.38)
SaleGrowth	0.071	0.063	0.165
	(1.26)	(1.47)	(0.83)
StockReturn	1.055***	0.159	2.169
	(2.81)	(0.55)	(1.59)
Constant	1.556	-1.314***	13.417***
	(1.52)	(-2.62)	(44.88)
Observations	190,571	520,721	32,504
Adjusted R-squared	0.817	0.706	0.029
Mean of Dep. Var	2.965	2.512	14.40
Firm FE	✓	✓	
Year#Week FE	✓	✓	

Table 12. -Continued.

Panel E: Control for Other AI Effects

	(1)	(2)	(3)
Variables	Log(Job Created + 1)		
Post-ChatGPT × Gen. AI Exposure	-1.087***	-1.168***	-1.352***
	(-3.58)	(-5.99)	(-7.35)
Post-ChatGPT × AI Exposure (Felten)	-2.038		
	(-0.86)		
Post-ChatGPT × AI Exposure (Webb)		-0.438	
		(-1.59)	
Post-ChatGPT × AI Exposure (SML)			0.381
			(0.46)
Observations	539,604	539,604	539,604
Adjusted R-squared	0.707	0.707	0.707
Mean of Dep. Var	2.534	2.534	2.534
Controls	✓	✓	✓
Firm FE	✓	✓	✓
Year#Week FE	✓	✓	✓

INTERNET APPENDIX

Appendix A

Figure A.1.
This figure gives a sample of job description.

<job>

<hash>f000417caf84521a55112de410ff8ae6</hash>

<description><![CDATA[POSITION SUMMARY:

The Archery Technician performs various selling / customer service activities, to include greeting and acknowledging all customers in a prompt and friendly manner, handling merchandise with care, providing information, assistance, assembly of products and direction / instruction to customers.

ESSENTIAL FUNCTIONS:
* Supports a strong commitment to world class customer service and ensures a pleasant and productive shopping experience for all customers.
* Troubleshoots, diagnoses, and tunes all archery equipment to make certain it is safe and functioning properly.
* Installs accessory items on Bows and demonstrates product to customers.
* Assists customers by answering specific technical questions to help customer understand the operation of the unit and the work to be done or work already completed.
* Continues technical training to improve proficiency, quality of work and to achieve higher levels of product knowledge.
* Assists with maintenance of pricing and UPC integrity; determining proper assortments; accuracy of inventory; proper display, signing and pricing of all products including advertised items.
* Provides a legendary experience for every customer, every time by assisting customers in making buying decisions by:
* identifying and evaluating customers' needs,
* making product recommendations based on this analysis,
* promoting programs including, but not limited to CLUB Membership, VOC and In-Store Pick-up.
* Replenishes product on shelves as required per Merchandising guidelines.
* Keeps work area organized, clean and well stocked with supplies.
* Follows all Company Policies and Procedures.
* Performs other duties, assignments and responsibilities as needed.
* ALL OTHER DUTIES AS ASSIGNED

EXPERIENCE/QUALIFICATIONS:
* Minimum Degree Required: High School Diploma or Equivalent
* Years of experience: 2+ years' experience with bow set up

Figure A.1. -Continued.

KNOWLEDGE, SKILLS, AND ABILITY:
* Must be able to perform bow set up for customers including but not limited to; draw length, identifying eye dominance, and determining draw weight.
* Must be able to perform bow assembly and accessory installation such as kisser button, peep sights, D loop, knock location, and scope installation.
* Knowledge of arrow set up to include but not limited to; arrow length, spine, and insert installation.
* Must be able to operate various equipment and tools including a bow press and chop saw.
* Instructs customers on safety, accessories, bow features and brand awareness.

TRAVEL REQUIREMENTS:
* N/A

PHYSICAL REQUIREMENTS:
* Constantly stand and/or walk during shift
* Constantly communicate with others to exchange information
* Constantly repeat motions that may include the wrists, hands and/or fingers
* Constantly lift or move objects weighing up to 20 pounds
* Occasionally operate machinery and/or power tools
* Occasionally ascend or descend ladders, stairs, ramps, etc.
* Occasionally work in noisy environments
* Light work that includes constantly moving or lifting objects up to 20 pounds, occasionally move and lift objects up to 100 pounds or more (utilizing a team lift as needed)

INDEPENDENT JUDGEMENT:
* Performs tasks and duties under general supervision, using established procedures and innovation. Chooses from limited alternatives to resolve problems. Occasional independent judgment is required to complete work assignments. Often makes recommendations to work procedures, policies, and practices.

Part Time Benefits Summary:
Enjoy discounts on retail merchandise, our restaurants, world-class resorts and conservation attractions!
* Dental
* Vision
* Voluntary benefits
* 401k Retirement Savings
* Paid holidays
* Paid vacation
* Bass Pro Cares Fund
* And more!

Bass Pro Shops is an equal opportunity employer. Hiring decisions are administered without regard to race, color, creed, religion, sex, pregnancy, sexual orientation, gender identity, age, national origin, ancestry, citizenship status, disability, veteran status, genetic information, or any other basis protected by applicable federal, state or local law.

Reasonable Accommodations
Qualified individuals with known disabilities may be entitled to reasonable accommodation under the Americans with Disabilities Act and certain state or local laws.
If you need a reasonable accommodation for any part of the application process, please visit your nearest location or contact us at hrcompliancebasspro.com.

Figure A.2.

This figure provides robustness checks for the baseline estimates presented in Panel A of Table 6. Panels (a), (b), and (c) address concerns regarding the potential influence of specific time periods or companies on the results. These panels show the estimates after excluding one year, one state, and one industry at a time, respectively. The standard errors used to construct the 90% confidence intervals, denoted by the spikes, are clustered at the firm level.

(a) Sensitivity to Year

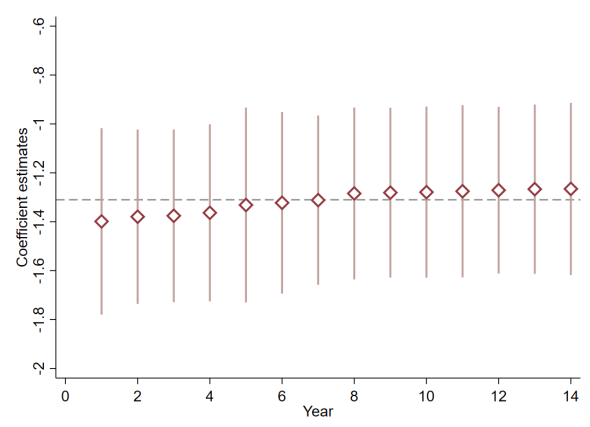

(b) Sensitivity to State

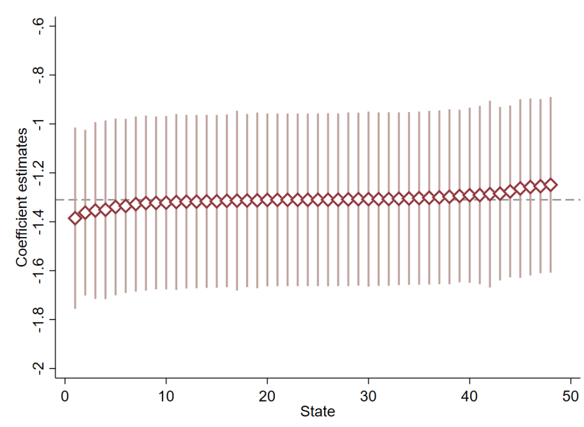

(c) Sensitivity to Industry

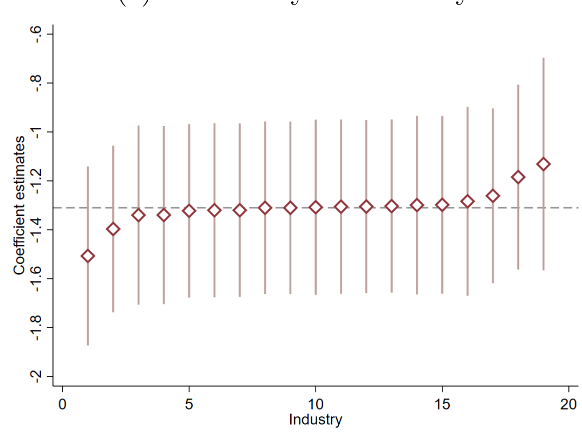

Table A. 1.

Variable Definitions

This table contains definitions of the primary variables used throughout the paper. All continuous variables are winsorized at 1% and 99%.

Panel A: LinkUp Variable Definition (Firm-Week Level)

Variable	Definition
Log(Job Created+1)	Natural logarithm of one plus job vacancy newly posted
Log(Machine Learning+1)	Natural logarithm of one plus the count of machine-learning-related keywords in job descriptions
Log(Generative AI+1)	Natural logarithm of one plus the count of generative-AI-related keywords in job descriptions
Post-ChatGPT	An indicator variable that equals one after the initial release date of ChatGPT (November 30, 2022)
Gen. AI Exposure	Generative AI exposure score

Panel B: Compustat Variable Definition (Firm-Week Level)

Variable	Definition
LgSize	Natural log of total asset (in millions)
Tangibility	Physical capital divided by total assets
Labor_Intensity	Natural logarithm of tangibility
Market to Book	The market value of a firm's equity divided by the book value of equity
ROA	Net income divided by total assets
CashHold	Cash and short-term investments divided by total assets
SaleGrowth	Change in cash sales
StockReturn	Average of monthly cumulative stock return
RD Intensity	Research and development expense divided by total sales
CorpAge	Years since corporation foundation

Panel C: LinkUp Variable Definition (Occupation-Week Level)

Variable	Definition
High Knowledge Job Created	Natural logarithm of one plus job vacancy newly posted (above median of knowledge score by O*NET 27.1)
Low Knowledge Job Created	Natural logarithm of one plus job vacancy newly posted (below median of knowledge score by O*NET 27.1)
High/Low Ratio Knowledge Job Created	Ratio of High Knowledge Job Created to Low Knowledge Job Created
High Skill Job Created	Natural logarithm of one plus job vacancy newly posted (above median of skill score by O*NET 27.1)

Table A.1. -Continued.

Panel C: LinkUp Variable Definition (Occupation-Week Level) -Continued

Variable	Definition
Low Skill Job Created	Natural logarithm of one plus job vacancy newly posted (below median of skill score by O*NET 27.1)
High/Low Ratio Skill Job Created	Ratio of High Skill Job Created to Low Skill Job Created
High Education Job Created	Natural logarithm of one plus job vacancy newly posted (above median of education requirement score by O*NET 27.1)
Low Education Job Created	Natural logarithm of one plus job vacancy newly posted (below median of education requirement score by O*NET 27.1)
High/Low Ratio Education Job Created	Ratio of High Education Job Created to Low Education Job Created
High Training Job Created	Natural logarithm of one plus job vacancy newly posted (above median of training requirement score by O*NET 27.1)
Low Training Job Created	Natural logarithm of one plus job vacancy newly posted (below median of training requirement score by O*NET 27.1)
High/Low Ratio Training Job Created	Ratio of High Training Job Created to Low Training Job Created
High TechSkill Job Created	Natural logarithm of one plus job vacancy newly posted (above median of technical skill requirement score by O*NET 27.1)
Low TechSkill Job Created	Natural logarithm of one plus job vacancy newly posted (below median of technical skill requirement score by O*NET 27.1)
High/Low Ratio TechSkill Job Created	Ratio of High TechSkill Job Created to Low TechSkill Job Created
JobZone 1 Job Created	Natural logarithm of one plus job vacancy newly posted (Job Zone 1 classification by O*NET 27.1)
JobZone 2 Job Created	Natural logarithm of one plus job vacancy newly posted (Job Zone 2 classification by O*NET 27.1)
JobZone 3 Job Created	Natural logarithm of one plus job vacancy newly posted (Job Zone 3 classification by O*NET 27.1)
JobZone 4 Job Created	Natural logarithm of one plus job vacancy newly posted (Job Zone 4 classification by O*NET 27.1)
JobZone 5 Job Created	Natural logarithm of one plus job vacancy newly posted (Job Zone 5 classification by O*NET 27.1)

Table A.1. -Continued.

Panel D: Glassdoor Interview Variable Definition (Firm-Week Level)

Variable	Definition
Interview Num	Natural logarithm of Glassdoor interview posted
Interview Num - Manager	Natural logarithm of Glassdoor interview posted on manager position
Interview Num - Non-Manager	Natural logarithm of Glassdoor interview posted on non-manager position
Experience Score	Average interview experience score
Difficulty Score	Average interview experience score

Panel E: Glassdoor Review Variable Definition (Firm-Week Level)

Variable	Definition
Approves of CEO	Approves of CEO rating
Overall Rating	Overall rating
Work/Life Balance	Work/Life balance rating
Career Opportunities	Career opportunities rating
Compensation and Benefits	Compensation and benefits rating
Senior Management	Senior management rating
Culture & Values	Culture & values rating
Diversity & Inclusion	Diversity & inclusion rating
Recommend	Recommend rating
Business Outlook	Business outlook rating

Table A. 2.

List of Keywords Used to Identify Machine Learning and Generative AI related Vacancies

This table listed the keywords used to identify Machine Learning and Generative AI related vacancies used in Table 8. The keywords are sourced from the list provided by Eisfeldt et al. (2023)

Machine Learning		Generative AI	
1	Deep Learning	1	LLM
2	ML	2	ChatGPT
3	Machine Learning	3	GPT
4	Natural Language	4	GPT3
5	Neural Net	5	GPT3.5
6	Neural Network	6	GPT4
7	NLP	7	GPT-3
		8	GPT-3.5
		9	GPT-4
		10	Generative
		11	Natural Language Generation

Appendx B

Figure B.1.

This figure plotted the binscatter of occupation-level exposure score and occupation characteristics. Coefficients and standard error of single-variable regression are also included. Occupation characteristics are calculated from O*NET 27.1. Occupational exposure score is the average of task exposure score (E1+0.5*E2). Levels of significance are presented as follows: *p<0.1; **p<0.05; ***p<0.01.

(a) Knowledge Requirement

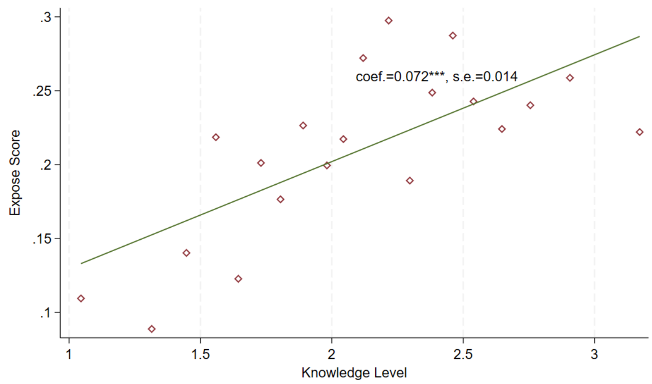

(b) Skill Requirement

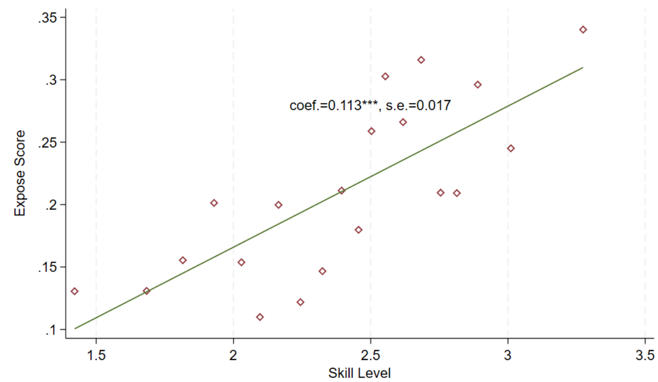

(c) Ability Requirement

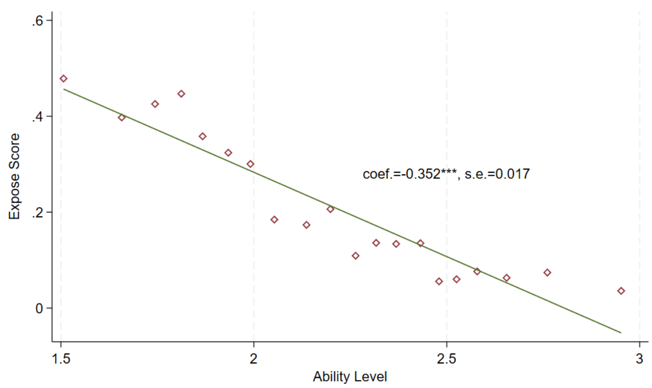

(d) Education Requirement

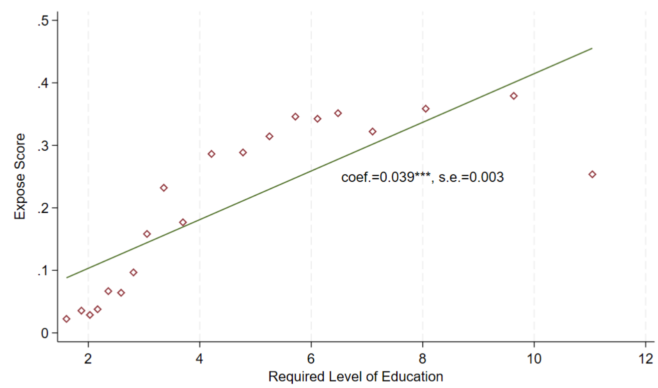

(e) Training Requirement

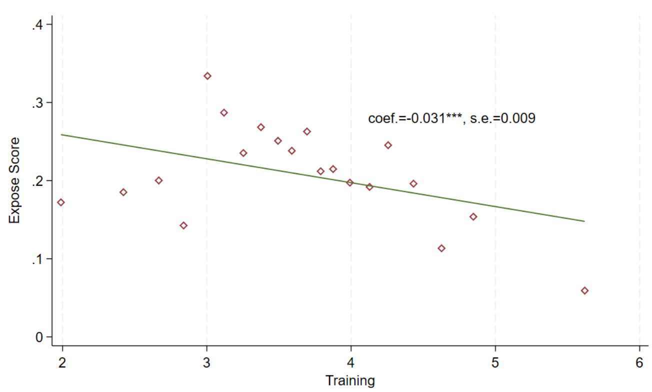

(f) Job Zone

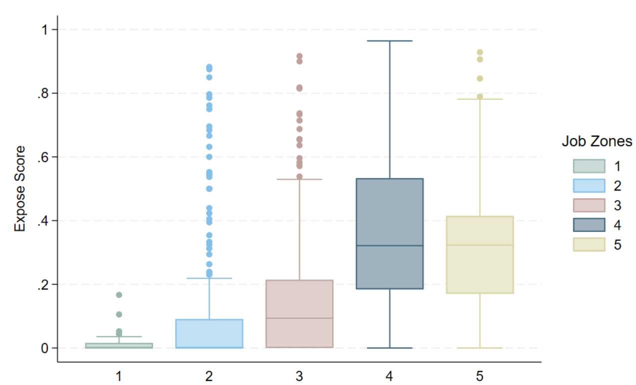

(g.1) Job Interests-Artistic

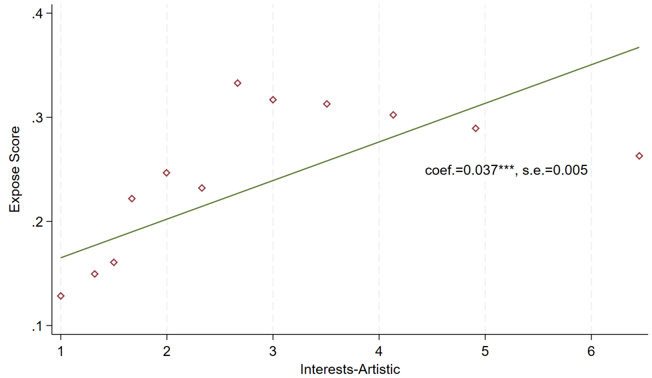

(g.2) Job Interests-Conventional

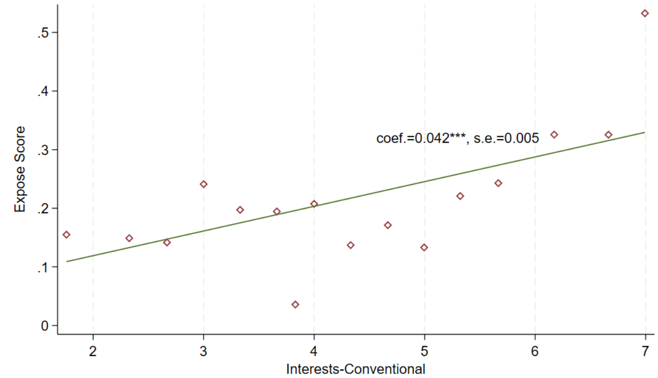

(g.3) Job Interests-Enterprising

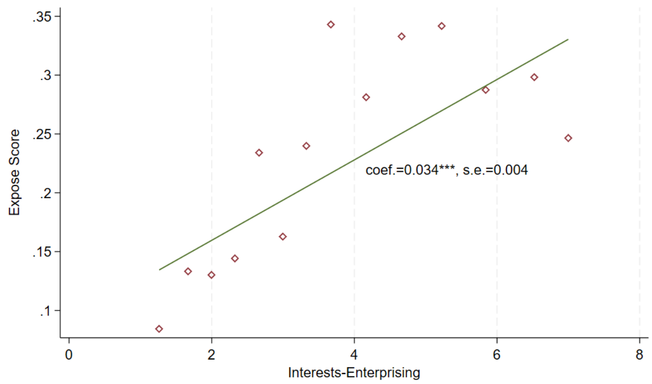

(g.4) Job Interests-Investigative

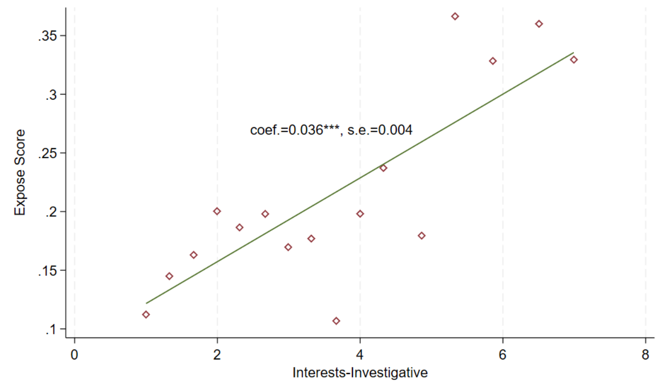

(g.5) Job Interests-Realistic

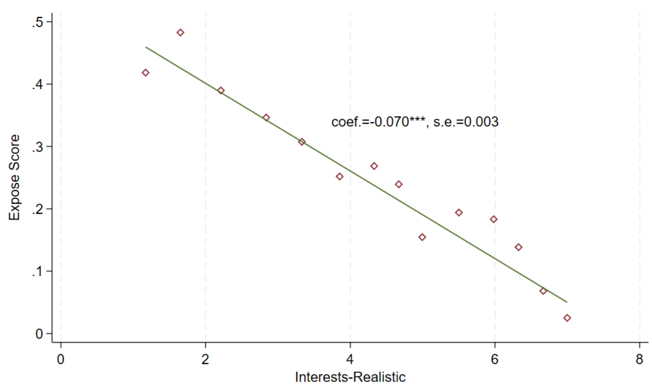

(g.6) Job Interests-Social

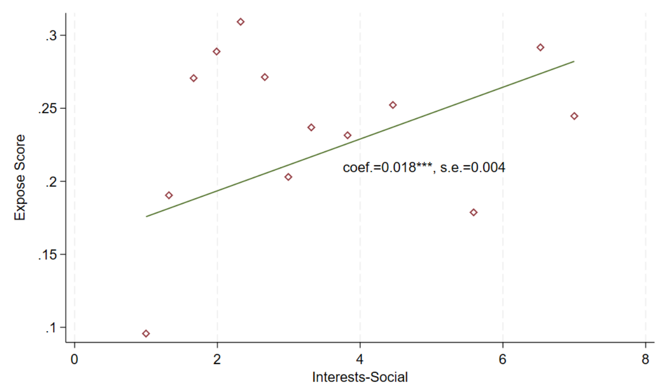

(h.1) Work Styles - Achivement/Effort

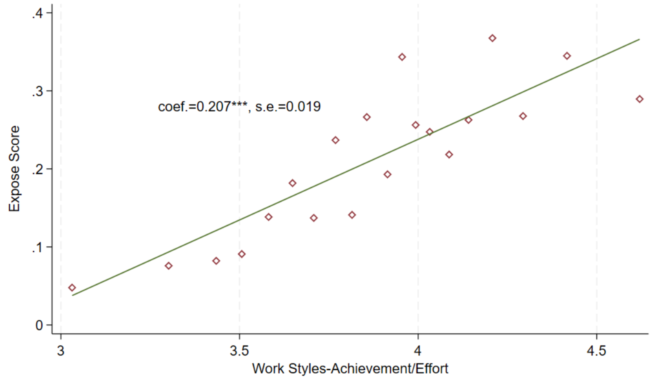

(h.2) Work Styles - Adaptability/Flexibility

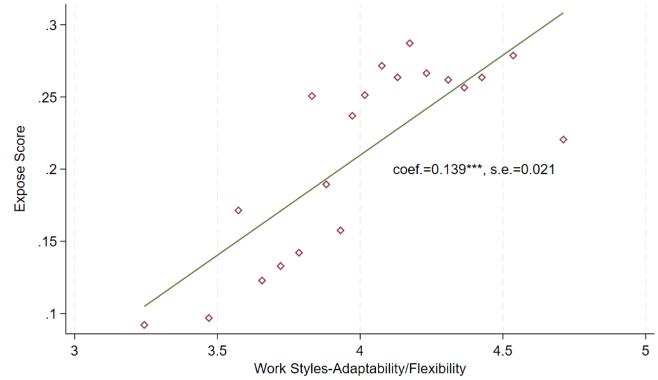

(h.3) Work Styles - Analytical Thinking

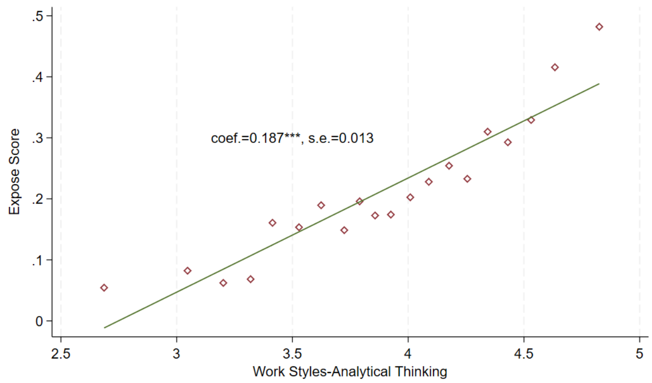

(h.4) Work Styles - Attention to Detail

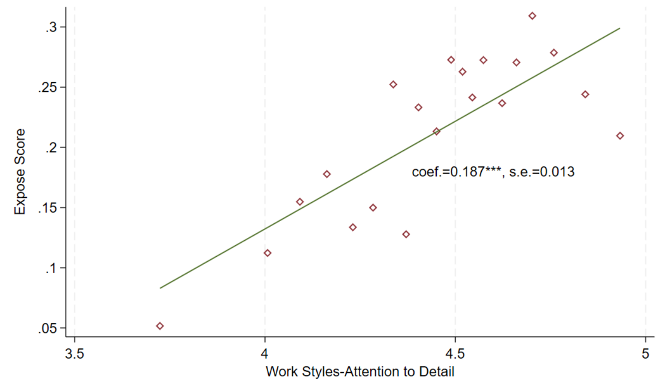

(h.5) Work Styles - Concern for Others

(h.6) Work Styles - Cooperation

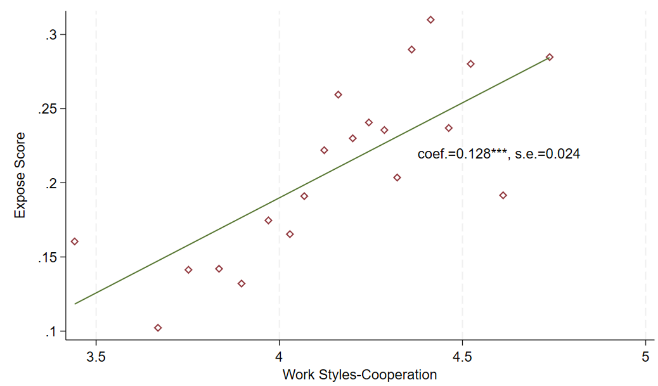

(h.7) Work Styles - Dependability

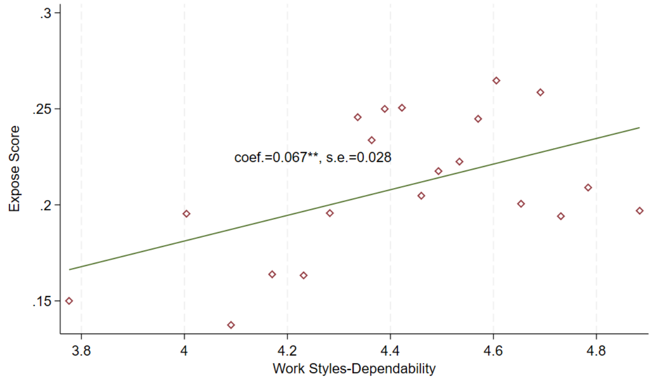

(h.8) Work Styles - Independence

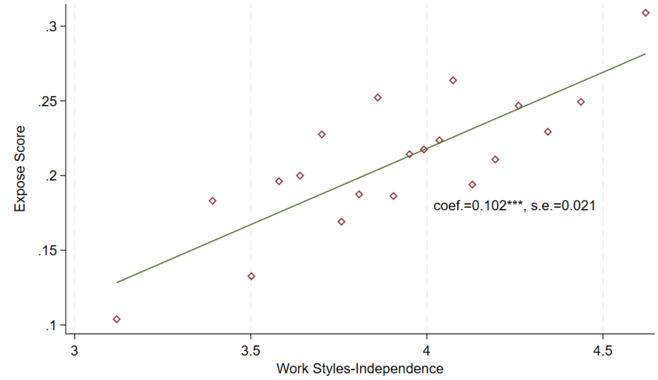

(h.9) Work Styles - Initiative

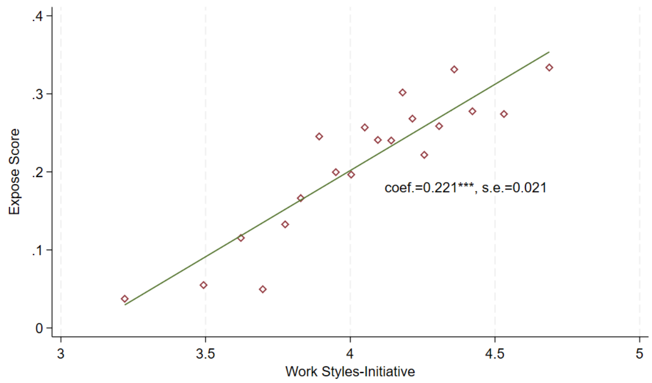

(h.10) Work Styles - Innovation

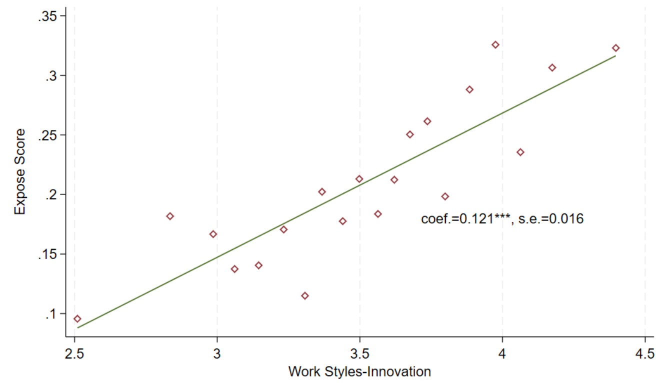

(h.11) Work Styles - Integrity

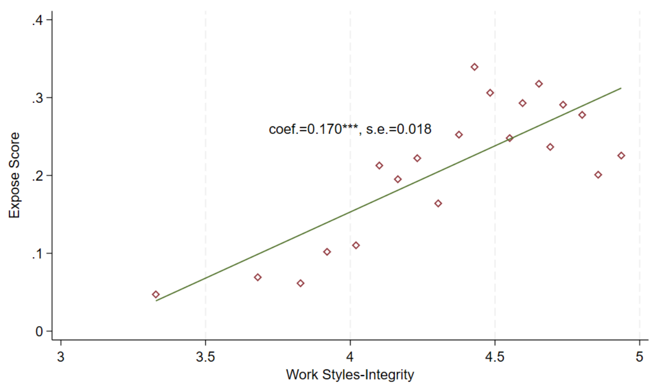

(h.12) Work Styles - Leadership

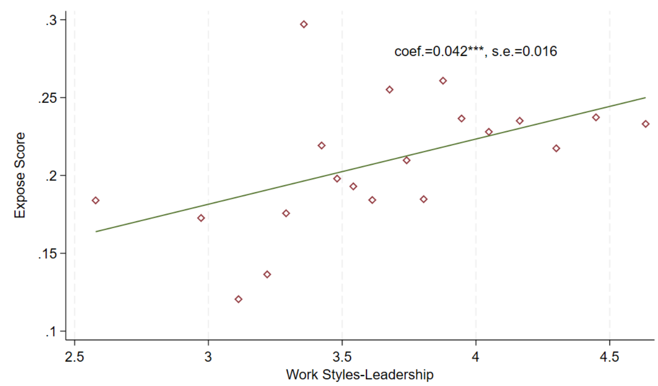

(h.13) Work Styles - Persistence

(h.14) Work Styles - Self Control

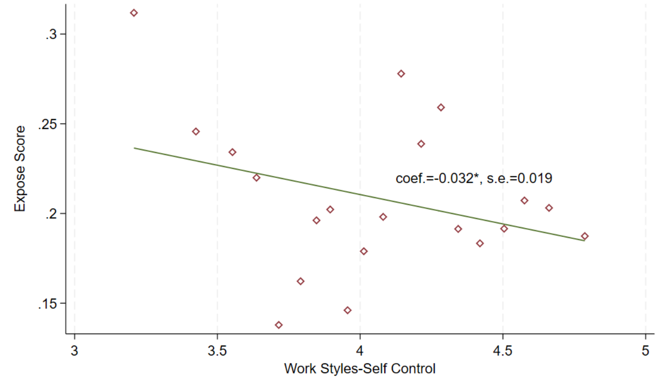

(h.15) Work Styles - Social Orientation

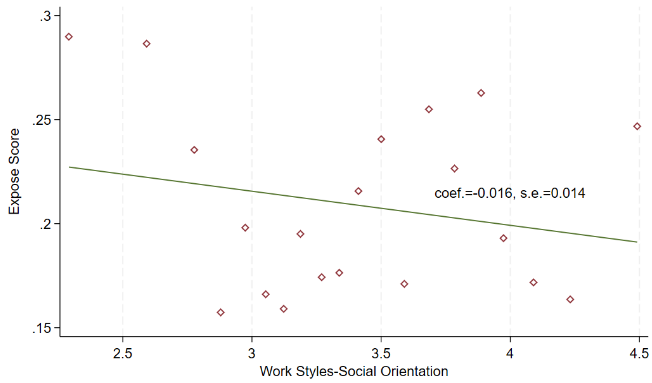

(h.16) Work Styles - Stress Tolerance

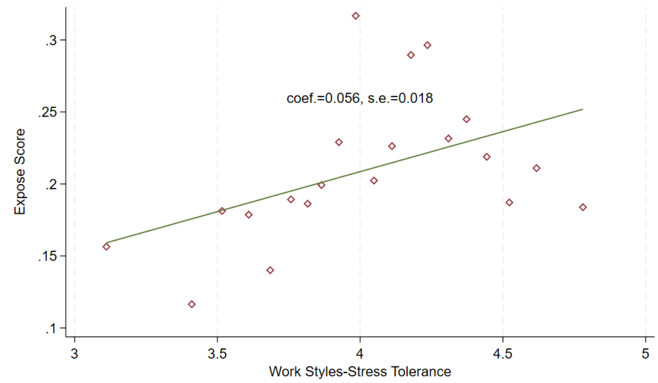

(i.1) Work Values - Achievement

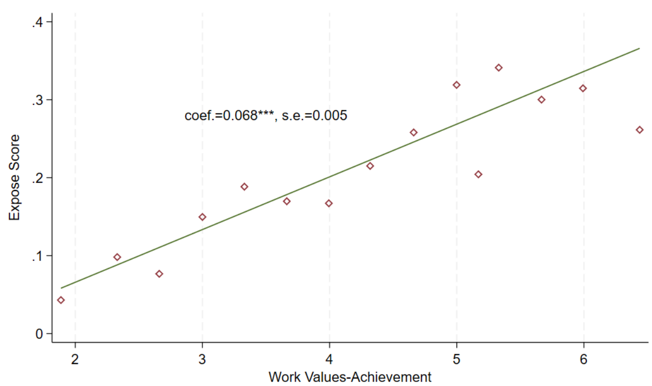

(i.2) Work Values - Independence

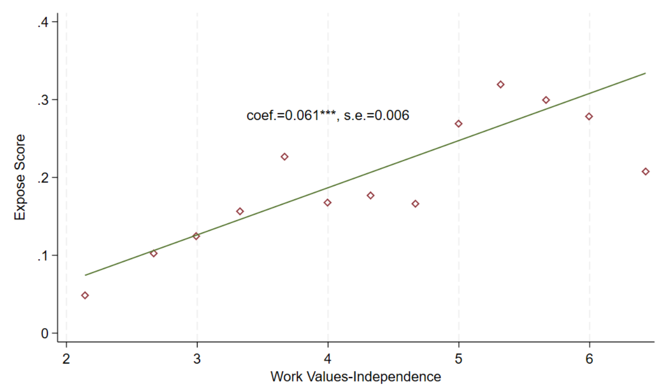

(i.3) Work Values - Recognition

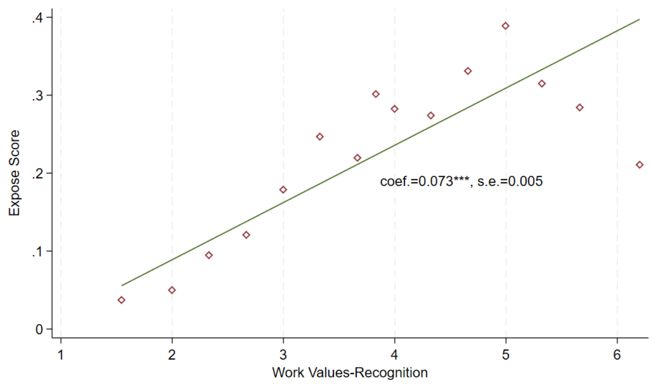

(i.4) Work Values - Relationship

(i.5) Work Values - Support

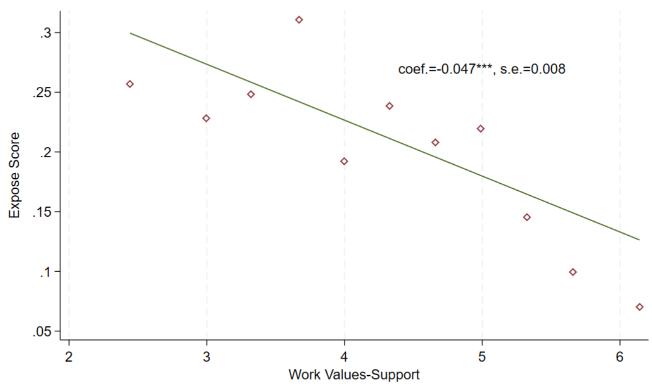

(i.6) Work Values - Working Conditions

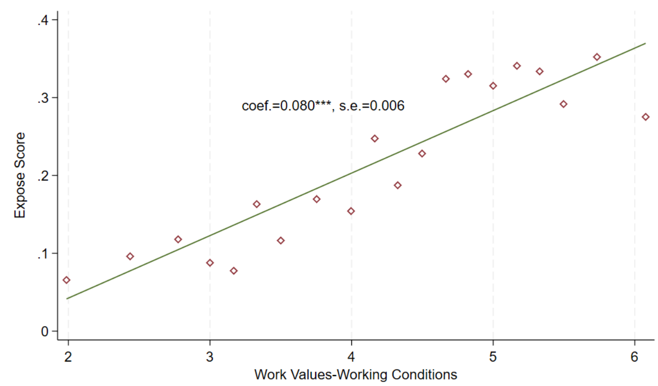

Table B.1.

O*NET-Task Exposure

This Table presents the GPT 3.5 rating of occupation-task. We train GPT 3.5 to score the expose of 18,189 distinct occupation-task to Generative AI. E0 means no exposure, E1 means direct exposure, E2 means exposure by LLM-powered applications, and E3 means exposure given image capabilities. Prompts used are illustrated in Appendix C. Occupation-task description comes from O*NET 27.1.

Tag	E0		E1		E2		E3	
	13691(75.27%)		3274(18.00%)		1149(6.32%)		75(0.41%)	
Confidence	High	Medium	High	Medium	High	Medium	High	Medium
	13691	0	3273	1	702	447	68	7

Appendix C

GPT Prompts for Exposure Scoring

System prompt = "Consider the most powerful OpenAI large language model (LLM). This model can complete many tasks that can be formulated as having text input and text output where the context for the input can be captured in 2000 words. The model also cannot draw up-to-date facts (those from ¡1 year ago) unless they are captured in the input. Assume you are a worker with an average level of expertise in your role trying to complete the given task. You have access to the LLM as well as any other existing software or computer hardware tools mentioned in the task. You also have access to any commonly available technical tools accessible via a laptop (e.g. a microphone, speakers, etc.). You do not have access to any other physical tools or materials. You are a helpful research assistant who wants to label the given tasks according to the rubric below. Equivalent quality means someone reviewing the work would not be able to tell whether a human completed it on their own or with assistance from the LLM. If you aren't sure how to judge the amount of time a task takes, consider whether the tools described exposed the majority of subtasks associated with the task.
Exposure rubric:
E1 - Direct exposure: Label tasks E1 if direct access to the LLM through an interface like ChatGPT or the OpenAI playground alone can reduce the time it takes to complete the task with equivalent quality by at least half. This includes tasks that can be reduced to: - Writing and transforming text and code according to complex instructions, - Providing edits to existing text or code following specifications, - Writing code that can help perform a task that used to be done by hand, - Translating text between languages, - Summarizing medium-length documents, - Providing feedback on documents, - Answering questions about a document, - Generating questions a user might want to ask about a document, - Writing questions for an interview or assessment, - Writing and responding to emails, including ones that involve refuting information or engaging in a negotiation (but only if the negotiation is via written correspondence), - Maintain records of written data, - Prepare training materials based on general knowledge, or - Inform anyone of any information via any written or spoken medium.
E2 - Exposure by LLM-powered applications: Label tasks E2 if having access to the LLM alone may not reduce the time it takes to complete the task by at least half, but it is easy to imagine additional software that could be developed on top of the LLM that would reduce the time it takes to complete the task by half. This software may include capabilities such as: - Summarizing documents longer than 2000 words and answering questions about those documents, - Retrieving up-to-date facts from the Internet and using those facts in combination with the LLM capabilities, - Searching over an organization's existing knowledge, data, or documents and retreiving information, - Retrieving highly specialized domain knowledge, - Make recommendations given data or written input, - Analyze written information to inform decisions, - Prepare training materials based on highly specialized knowledge, - Provide counsel on issues, and - Maintain complex databases.
E3 - Exposure given image capabilities: Suppose you had access to both the LLM and

a system that could view, caption, and create images as well as any systems powered by the LLM (those in E2 above). This system cannot take video as an input and it cannot produce video as an output. This system cannot accurately retrieve very detailed information from image inputs, such as measurements of dimensions within an image. Label tasks as E3 if there is a significant reduction in the time it takes to complete the task given access to a LLM and these image capabilities: - Reading text from PDFs, - Scanning images, or - Creating or editing digital images according to instructions. The images can be realistic but they should not be detailed. The model can identify objects in the image but not relationships between those options
E0 - No exposure: Label tasks E0 if none of the above clearly decrease the time it takes for an experienced worker to complete the task with high quality by at least half. Some examples: - If a task requires a high degree of human interaction (for example, in-person demonstrations) then it should be classified as E0. - If a task requires precise measurements then it should be classified as E0. - If a task requires reviewing visuals in detail then it should be classified as E0. - If a task requires any use of a hand or walking then it should be classified as E0. - Tools built on top of the LLM cannot make any decisions that might impact human livelihood (e.g.hiring, grading, etc.). If any part of the task involves collecting inputs to make a final decision (as opposed to analyzing data to inform a decision or make a recommendation) then it should be classified as E0. The LLM can make recommendations. - Even if tools built on top of the LLM can do a task, if using those tools would not save an experienced worker significant time completing the task, then it should be classified as E0. - The LLM and systems built on top of it cannot do anything that legally requires a human to perform the task. - If there is existing technology not powered by an LLM that is commonly used and can complete the task then you should mark the task E0 if using an LLM or LLM-powered tool will not further reduce the time to complete the task. When in doubt, you should default to E0."

User prompt = "Read the following occupation title and one of the tasks that belong to that occupation. Then do three things. 1: Reason step by step to decide which of the labels (E0/E1/E2/E3) from the exposure rubric you were given applies to the task's exposure to LLM. Report. Give an explanation for which label applies and report the label that you think fits best. Do not say zero or N/A.
2: Report only the label that you determined for the task, which should match the label in step 1. Do not reply N/A.
3: Given the amount of speculation required in step 1, describe your certainty about the estimate–either high, moderate, or low.
The task to label follows:
Occupation: Inspectors, Testers, Sorters, Samplers, and Weighers. Task: Adjust, clean, or repair products or processing equipment to correct defects found during inspections."
Assistant prompt = "1: Label (E0/E1/E2/E3): E0. Explanation: The model does not have access to any kind of physicality, and more than half of the task (adjusting, cleaning and repairing equipment) described requires hands or other embodiment.
2: E0.

3: High confidence."

User prompt = "Read the following occupation title and one of the tasks that belong to that occupation. Then do three things. 1: Reason step by step to decide which of the labels (E0/E1/E2/E3) from the exposure rubric you were given applies to the task's exposure to LLM. Report. Give an explanation for which label applies and report the label that you think fits best. Do not say zero or N/A.
2: Report only the label that you determined for the task, which should match the label in step 1. Do not reply N/A.
3: Given the amount of speculation required in step 1, describe your certainty about the estimate–either high, moderate, or low.
The task to label follows:
Occupation: Computer and Information Research Scientists. Task: Apply theoretical expertise and innovation to create or apply new technology, such as adapting principles for applying computers to new uses."
Assistant prompt = "1: Label (E0/E1/E2/E3): E1. Explanation: The model can learn theoretical expertise during training as part of its general knowledge base, and the principles to adapt can be captured in the text input to the model.
2: E1.
3: Medium confidence."

User prompt = "Read the following occupation title and one of the tasks that belong to that occupation. Then do three things. 1: Reason step by step to decide which of the labels (E0/E1/E2/E3) from the exposure rubric you were given applies to the task's exposure to LLM. Report. Give an explanation for which label applies and report the label that you think fits best. Do not say zero or N/A.
2: Report only the label that you determined for the task, which should match the label in step 1. Do not reply N/A.
3: Given the amount of speculation required in step 1, describe your certainty about the estimate–either high, moderate, or low.
The task to label follows:
Occupation: Video Game Designers. Task: Create gameplay prototypes for presentation to creative and technical staff and management."
Assistant prompt = "1: Label (E0/E1/E2/E3): E2. Explanation: The model could aid in generating both textual descriptions and basic visual elements, and with additional development, software could leverage these capabilities to create more comprehensive gameplay prototypes. However, the complexity and quality of the visual prototypes would depend on the capabilities of the given image tools and the development effort involved.
2: E2.
3: Medium confidence."

User prompt = "Read the following occupation title and one of the tasks that belong to that occupation. Then do three things. 1: Reason step by step to decide which of the labels (E0/E1/E2/E3) from the exposure rubric you were given applies to the task's exposure to

LLM. Report. Give an explanation for which label applies and report the label that you think fits best. Do not say zero or N/A.
2: Report only the label that you determined for the task, which should match the label in step 1. Do not reply N/A.
3: Given the amount of speculation required in step 1, describe your certainty about the estimate–either high, moderate, or low.
The task to label follows:
Occupation: Executive Secretaries and Executive Administrative Assistants. Task: Schedule dining reservations."
Assistant prompt = "1: Label (E0/E1/E2/E3): E2. Explanation: Automation technology already exists for this (e.g. Resy) and it's unclear what an LLM offers on top of using that technology (no-diff). That said, you could build something that allows you to ask the LLM to make a reservation on Resy for you.
2: E2.
3: High confidence."