The Labor Market Effects of Generative Artificial Intelligence

Jonathan S. Hartley, Brendan Moore, Filip Jolevski, and Vitor Melo*

March 26, 2026

Abstract

Do worker beliefs about layoff risk due to generative artificial intelligence shape the labor market impact of the technology itself? We field a survey measuring workers' perceptions of AI-related job loss risk and AI adoption in the labor market, which we measure to be 36% among U.S. workers at the end of 2025. We find that workers perceive substantial layoff risk from generative AI (particularly in highly-exposed occupations) even though realized layoffs remain limited and vacancies in these occupations have not declined. Motivated by this fact, we develop a model of on-the-job search in which switching jobs resets match tenure and exposes workers to higher layoff risk in new employment relationships. Consistent with the model, we find that workers in more AI-exposed occupations exhibit lower job-to-job mobility and reduced on-the-job search. These results suggest that perceptions of automation risk can dampen labor market dynamism even in the absence of large realized displacement.

*Hartley: Stanford University, Hoover, Mercatus, Macdonald-Laurier Institute, FREOPP, hartleyj@stanford.edu, Jolevski: World Bank, fjolevski@worldbank.org, Melo: Clemson University and Knee Regulatory Research Center at West Virginia University, vitormelo95@gmail.com, Moore: Stanford University, bdmoore@stanford.edu. Previously circulated under "The Labor Market Effects of Artificial Intelligence." The surveys in the study have been funded by Emergent Ventures Grant #7349 as well as the Archbridge Institute. The authors are grateful for excellent research assistance from Cody Dirks, Michael Li, and Omar Mirza for helpful comments from David Autor, Nick Bloom, Erik Brynjolfsson, Tyler Cowen, David Deming, Erik Hurst, Simon Jager, Liya Palagashvili, Eric Schmidt, and Jesse Shapiro. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

"Yet there is no country and no people, I think, who can look forward to the age of leisure and of abundance without a dread" -John Maynard Keynes, Economic Possibilities for our Grandchildren (1930)

1 Introduction

Few technological transitions and general purpose technologies (GPTs) in recent memory have generated as much debate about labor market consequences as the rapid diffusion of generative Artificial Intelligence (GenAI). Since the public release of ChatGPT in late 2022, economists, technologists, and policymakers have speculated about the extent to which large language models (LLMs) will fundamentally reshape the occupational landscape:eliminating jobs at scale, augmenting worker productivity, or both. A rapidly growing empirical literature documenting the labor market effects of LLMs tends to show limited overall effects of AI on outcomes such as employment and wages through 2025 (Humlum and Vestergaard, 2025), although some papers emphasize localized negative effects for young and entry-level workers (Brynjolfsson et al., 2025; Liu et al., 2025; Hosseini Maasoum and Lichtinger, 2025). What has received less attention, however, is a more subtle but potentially consequential channel: the role of beliefs about AI-induced displacement in shaping labor market behavior, independently of whether realized displacement occurs at all.

This paper pursues three interrelated questions. First, how widespread is generative AI adoption among U.S. workers as well as its direct labor market effects, and what do workers believe about the risk of losing their jobs to the technology? Second, do these beliefs themselves alter labor market behavior in ways that could reduce dynamism even absent large-scale layoffs?

To address these questions, we conduct a large-scale survey of U.S. workers, fielded quarterly across six waves from December 2024 through March 2026, with roughly 19,900 total respondents. We document that Generative AI adoption at work reached approximately 30% among U.S. workers by end of the first quarter of 2026 (a level consistent with contemporaneous survey evidence from Bick et al. 2026) before plateauing in the latter half of the year. Adoption is sharply stratified by education and income, with graduate-degree holders and workers earning above $200,000 annually adopting at rates approaching 50%, while workers without college degrees adopt at roughly half that rate.

Beyond adoption patterns, we uncover a striking divergence between perceived and realized displacement risk. Despite limited evidence of AI-driven layoffs in aggregate employment and vacancy data, American workers report remarkably high subjective probabilities of job loss: the average worker believes there is a one-in-five chance of being displaced by AI over the next two years. Among workers actively using generative AI on the job, this figure rises to 22%. Estimates of U.S. workers' average annual subjective job loss probability (for any reason) are around 13%. The implied AI-specific rate alone approaches that benchmark. We also document that the more workers use AI on the job, the more they fear AI-induced job loss risk, even holding fixed occupation and demographics.

This disconnect between perceived and realized risk is not, we argue, merely a curiosity. A well-established body of evidence demonstrates that workers alter consumption in anticipation of layoff risk, even before any actual displacement occurs (Stephens Jr, 2004; Hendren, 2017; Karahan et al., 2019). If workers in AI-exposed occupations believe that switching jobs exposes them to elevated displacement risk at a new employer (because tenure protections are lower and AI exposure is more uncertain in novel matches), then the threat of AI may suppress on-the-job search and job-to-job mobility even in the absence of observed displacement.

We formalize this logic in a model of on-the-job search with AI displacement risk, building on the Burdett and Mortensen (1998) framework. The key mechanism is straightforward: workers internalize that newly formed employment matches face higher separation rates than established ones, because tenure resets upon switching and new matches involve uncertain AI exposure. When perceived AI displacement risk rises, the expected value of a new match falls relative to staying put, dampening search intensity. The model yields three testable predictions: AI-exposed occupations should exhibit lower on-the-job search intensity, reduced job-to-job mobility, and greater wage dispersion as workers become less likely to move toward higher-paying opportunities.

Consistent with these predictions, our empirical results show that workers in more AI-exposed occupations search less on the job and transition between employers at lower rates. On the demand side, we find little evidence that AI adoption has meaningfully reduced job openings or employment levels in exposed occupations (point estimates are small and statistically indistinguishable from zero). However, wages have risen in high-exposure occupations following ChatGPT's release, consistent with AI functioning as a labor complement that raises productivity rather than a substitute that displaces workers in the near term. Taken together, these results suggest that AI's near-term labor market footprint operates more through a belief-driven suppression of dynamism than through direct displacement.

Our paper contributes to a growing literature on the labor market impacts of generative AI. We complement recent survey-based work by Bick et al. (2026), extending their analysis to the question of subjective displacement risk and its behavioral consequences. Our labor demand analysis builds on the event-study designs of Brynjolfsson et al. (2025) and Lichtinger and Maasoum (2025), whose findings of negative employment effects (particularly for younger workers) we discuss in the context of our own null results on employment and positive results on wages. More broadly, our paper joins a literature on the real effects of subjective uncertainty, including foundational work by Stephens Jr (2004), Hendren (2017), and Karahan et al. (2019) on how layoff risk shapes household behavior, and the more recent contribution of Lehner et al. (2025) on worker beliefs about displacement.

The remainder of the paper proceeds as follows. Section 2.1 describes our survey design, documents AI adoption patterns, and characterizes workers' beliefs about displacement risk. Section 3 presents our empirical analysis of labor demand, employment, and wages using Lightcast job postings data and BLS Occupational Employment Statistics. Section 4 develops and solves the on-the-job search model. Section 5 concludes.

2 Generative AI Adoption and Job Loss Risk Survey

2.1 Methodology

This paper's first contribution is to conduct a survey of workers (similar to the Pew and Bick, Blandin, and Deming (2026) Generative AI surveys) with IncQuery specifically asking about the use of LLMs in specific tasks and occupations. One goal is to determine whether AI automatable tasks identified by papers such as Eloundou et al. (2024), Brynjolfsson et al. (2018), Webb (2019), and Felten et al. (2023), are actually being used in the context of LLMs.

The survey is field in five waves across 19,902 respondents in a U.S. sample conducted quarterly from December 2024 through March 2026, among whom roughly 7,500 were generative AI users at work. The survey used a simple random sampling approach with a probabilistic draw of the sample from Dynata. The unit of analysis are workers in the U.S. The survey was self-administrated using a CAWI script. The questions used in our survey are similar to the survey conducted by Bick et al. (2026) (which instead uses Qualtrics), asking several questions about Generative AI use as well as other questions about education level, industry, income, and many other metrics.

Many of the questions from Bick et al. (2026) are repeated verbatim in our survey. Many of their questions focus on the extensive margin of AI use. With respect to time and frequency of Generative AI use, their central question (replicated in our survey) is fairly generic and begins with defining Generative AI and Large Language Models: "Generative AI is a type of artificial intelligence that creates text, images, audio, or video in response to prompts. Some examples of Generative AI include ChatGPT, Gemini, and Claude. A Large Language Model is A type of artificial intelligence (AI) model that uses machine learning to understand and generate human language."

This description of GenAI was followed by a series of questions: "Have you heard of generative AI before?", "What do you think is the probability (0-100%) that you will lose your job due to generative AI in the next two years?", "Have you used generative AI tools before?", and "Do you use generative AI for your job?" We use the last question as a screener more detailed questions about GenAI usage.[1]

For the purpose of maximizing our sample size of respondents adopting Generative AI, we ask demographic, occupation, and industry questions prior to the GenAI questions and then screen out respondents who do not use Generative AI at work while maintaining the demographic data of those removed in order to produce estimates of Generative AI use. We also screen out respondents under the age of 18, like Bick et al. (2026); however, we include those above age 65 as one difference between surveys. We also use several attention checks.

We also ask further questions about the intensive margin of GenAI use. For example, we ask questions such as "How many days each week do you use generative AI tools at work?" and "How

[1] For survey waves 2–5, the question about perceived job loss risk due to AI was only included after the screener and thus was only asked to those who used GenAI at work. From wave 6 (March 2026) onwards, it was asked to all respondents. It was not asked in wave 1 (Dec 2024).

many hours per week do you use generative AI tools at work?"

We also ask "In what tasks are you using generative AI/Large Language models at work?", where respondents select from a list of 35 options. Commonly selected tasks include writing, time management, troubleshooting, active learning, operations analysis, and critical thinking.

In addition to surveys in the U.S., we also implemented comparable surveys in Italy, Canada, France, Japan, Sweden, Germany, Netherlands, Norway, New Zealand, Spain, U.S., Ireland, UK, and Australia.

See Appendix C for our complete survey instrument.

2.2 Adoption Measures

We document that GenAI adoption at work among U.S. workers rose sharply in early 2025 and stabilized at a level in the mid-to-high 30% range over the remainder of the year (Figure 1a).[2][3] These levels are broadly consistent with the estimates of Bick et al. (2026), who report comparable adoption rates in their U.S. survey (shown in green in Figure 1a).

Internationally, adoption varies from Italy, Canada and Japan, which have less than 25% of their workforces using Generative AI to the US, Ireland, UK, and Australia which have adoption rates of about 40% or more (see Figure 1b).

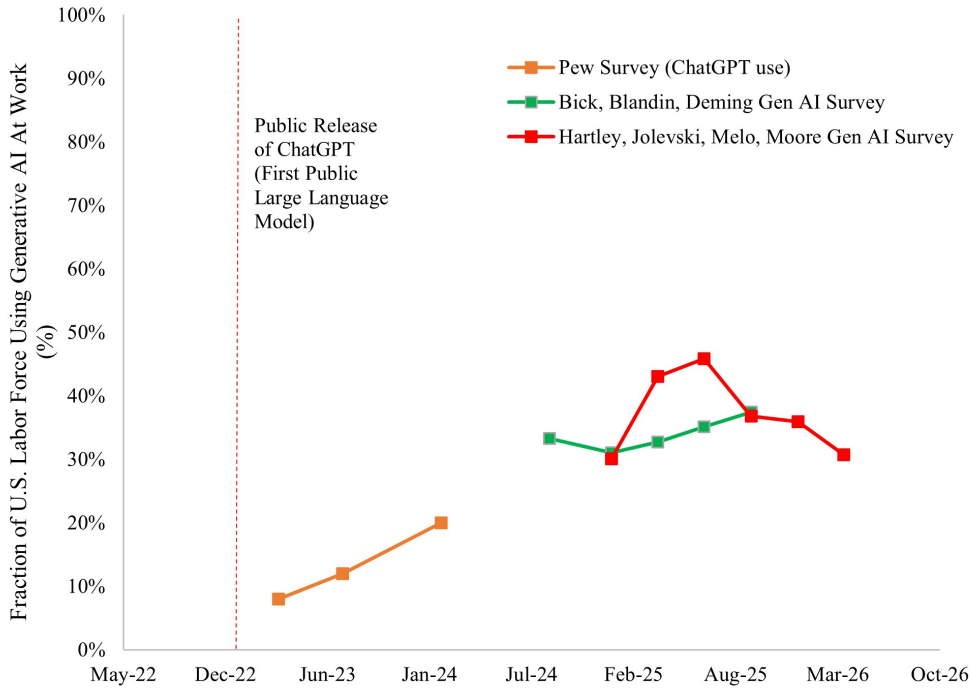
Adoption of GenAI at work exhibits several strong gradients. For example, educated workers are substantially more likely to use these tools at work. In our December 2024 wave, nearly 50% of workers with graduate degrees report using Generative AI at work, compared with 37% of college graduates and roughly 20% of workers with a high school degree or some college (Appendix Figure A.1). A similar pattern emerges across income groups: adoption is about 20% among workers earning below $50,000 but rises steadily with earnings, reaching nearly 50% among those earning above $200,000 annually (Appendix Figure A.2). Usage also varies across industries and occupations, with the highest adoption rates in information services, management, real estate, construction, and education, and substantially lower rates in sectors such as agriculture, mining, and government (Appendix Figure A.3). Conditional on adoption, many users report frequent but intermittent use: roughly one-third of users report using GenAI every weekday (Appendix Figure A.6), and the median user reports working with the technology five hours per week. Additional descriptive results by demographic group, industry, occupation, task type, and usage intensity are reported in Appendix Figures A.1–A.10).

[2]To reweight our results, we took the November 2024 CPS (the closest time to when we began the survey) and restricted to those who were working or held jobs (i.e., not unemployed or NILF), aggregated into two education groups (< BA or >= BA), four age groups (18-34, 35-44, 45-54, 55+), and four races (White, Black, Hispanic, and other). We then collapse the data by using CPS survey weights (wtfinl) which is the variable which estimates how many total Americans a given respondent should represent how many Americans are in a particular age/sex/education/race cell and then divide by the overall working population to determine the nationally representative weights.

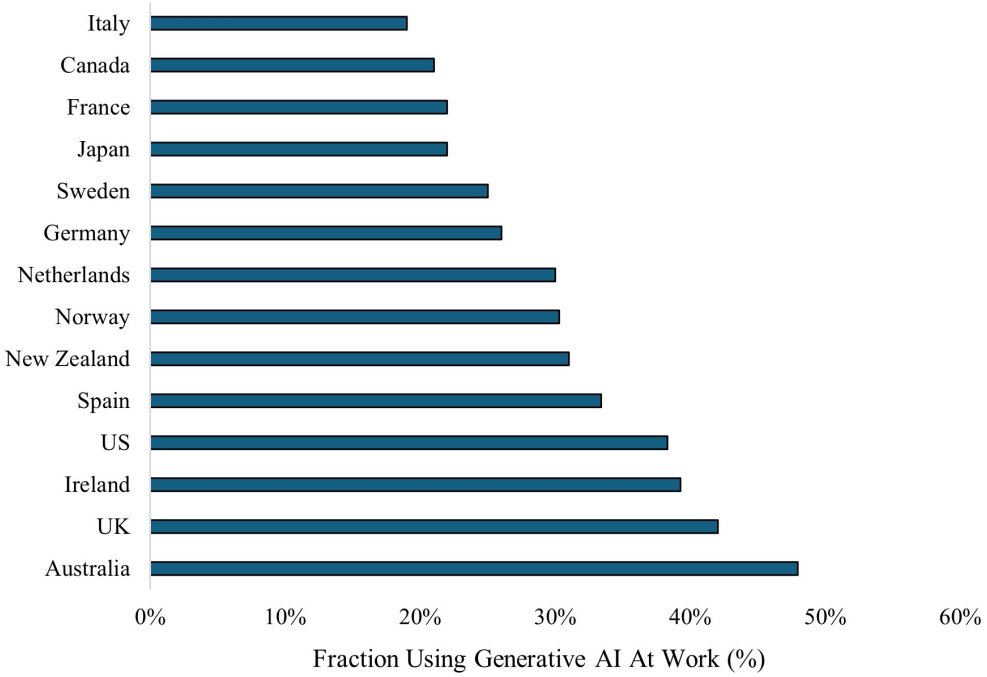
[3]The late-2025 decline in measured adoption is consistent with plateauing interest in GenAI tools, including declining Google search volume for "ChatGPT" and "Gemini", which we document in Appendix Figure A.12.

Figure 1: Fraction of Labor Force Using Generative AI Use At Work

(a) U.S. Labor Force Adoption Over Time



(b) Adoption at Work Across Countries in 2025



Note: Top panel shows adoption over time. First survey (orange) from Pew Charitable Trusts (2023-2024). Second survey (green) from Bick, Blandin, and Deming (2026) Third survey (red) from Hartley, Jolevski, Melo, and Moore (2026) Bottom panel shows international adoption in 2025 from Hartley, Jolevski, Melo, and Moore (2026).

We replicate the finding of Bick et al. (2026) that the occupations that were predicted to be most "exposed" to LLMs also tended to have higher adoption rates as well (Figure A.11). Specifically, we relate an occupation group's LLM exposure score from Eloundou et al. (2024) to that occupation's generative AI adoption at work in our survey.[4] We find AI exposure strongly predicts adoption among occupation groups with an estimated coefficient of 0.55 (similar to 0.67 in Bick et al. (2026)). For example, architecture & engineering and business & finance occupations have both high predicted AI exposure as well as high AI adoption in our survey, while blue collar and personal service occupations have both low predicted AI exposure and low adoption. However, certain occupational groups are more likely to adopt LLMs at work given their exposure measures, such as management roles, while other occupations (office & administration) under-adopt given their exposure. Such differences may be due to legal or workplace regulations, workplace norms around AI adoption, or individual workers' preferences when it comes to using such technologies.

We use this measure of "predicted AI adoption" – the fitted values from a granular occupation level regression of hypothetical exposure (Eloundou et al., 2024) on adoption (our survey) – throughout the rest of the paper, including for event study and difference-in-difference regressions in section 3.

2.3 Perceived Job Loss due to AI

In recent years, media commentators, academics, and technologists have increasingly warned that advances in AI may lead to significant labor market displacement (Amodei, 2026; Steinberger, 2026; Brynjolfsson et al., 2025). Large technology and telecommunications firms such as Amazon, Microsoft, Intel, Verizon, HP, Intel, IBM, Lenovo, and Block have all cited AI as a reason for large layoffs since 2025. Firms are increasingly examining how employees are adopting AI as part of their annual performance reviews, leading to concerns that companies are asking workers to train their replacements (Bloomberg News, 2025).

We document that American workers report heightened fear of job loss due to generative AI. The average American worker (regardless of AI adoption at work) believes there is a 20% chance they will be displaced in the next two years due to AI. Those who use AI at work report an average 22% chance of AI-induced job loss, compared to 17% for those who do not use AI in their workplace. Figure 2 shows the distribution of these job loss probabilities for all workers (blue) as well as for the sample of workers who report using GenAI at work (red). The main difference between the samples is that non-AI users are more certain of their job security, with 53% reporting a less-than-5% of AI-induced job loss, compared to just 40% for AI-adopters. A non-negligible share of AI-adopters (10%) assesses the chance of job loss at more than 50%, with some individuals believing with certainty that they will lose their job over the course of the next year.

To contextualize these large fears of job loss due to AI, we rely on the New York Fed's Survey of

[4]Following Bick et al. (2026), we use the β score from Eloundou et al. (2024), which describes how exposed an occupation is with the assistance of an additional technology.

Consumer Expectations (SCE) Labor Force Supplement, which regularly surveys American workers about their expectations of job loss for any reason. In 2024, U.S. workers reported an average subjective job loss rate over the following twelve months of 13%, nearly matching the implied 10% one-year job loss rate due to AI alone.[5] Such heightened job loss fears are significant given well-known workers change their behaviors of consumption (Stephens Jr, 2004; Hendren, 2017; Karahan et al., 2019) and job search (Lehner et al., 2025) in response to layoff risk. This perceived heightened layoff risk due to AI motivates our model in section 4.

Workers More Highly Exposed to AI Have Greater Displacement Fear

We document a new, stylized fact that workers more exposed to generative AI believe they are more likely to be displaced by the technology. Using a combination of our survey data on AI adoption and job loss perceptions as well as a battery of AI exposure measures, we demonstrate the robustness of this relationship.

First, we show that workers who report using AI regularly for more hours per week at their job also report higher subjective job loss risk due to the technology — a strong positive relationship that holds even conditioning on two digit occupation group fixed effects and demographic characteristics (Figure 3). In other words, within a given demographic group and broad occupational category (e.g. "Computer and Mathematical Occupation," "Legal Occupation," etc.), the more workers use generative AI on the job, the more they fear being displaced due to the technology. One reason for this intensive margin relationship between AI job loss fear and utilization could be that workers become more familiar with the technology's capabilities.

3 Effects of LLMs on Vacancies and Employment

3.1 Data

3.1.1 Current Population Survey

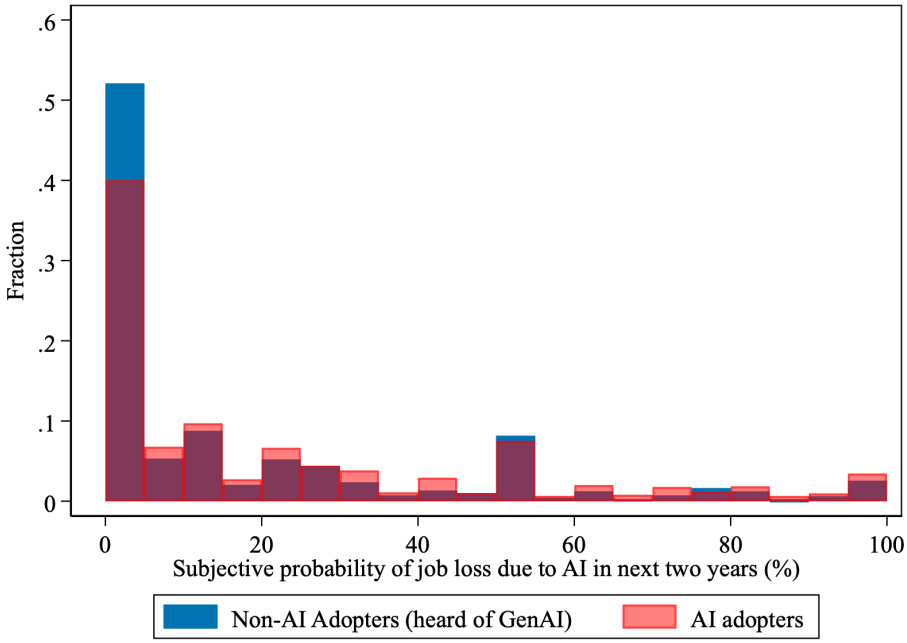
To measure labor supply dynamics and worker mobility, we use the Current Population Survey (CPS) Basic Monthly files from January 2019 through December 2025. The CPS is a monthly household survey conducted by the U.S. Census Bureau for the Bureau of Labor Statistics, providing nationally representative data on employment status, occupation, and industry for the civilian non-institutional population aged 16 and older.

Our primary outcomes of interest are labor market transition rates, which we construct by linking respondents across consecutive monthly interviews using the CPS rotation group structure. Individuals are interviewed for four consecutive months, leave the sample for eight months, and return for

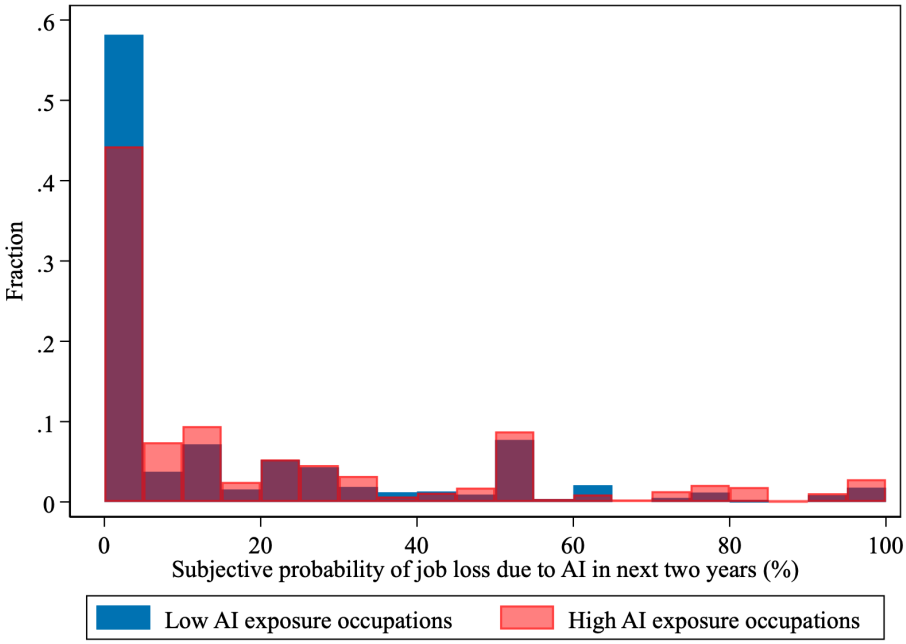
[5] We arrive at this calculation assuming workers expect a constant probability of job loss over the two year window they were asked to form their expectations. If workers instead anticipate rapid AI advancements and deployment at the latter end of the two-year window, this 10% imputation would be too high for a one-year horizon.

Figure 2: Job Loss Perceptions Due to Generative AI During the Next Two Years

(a) by GenAI Adoption

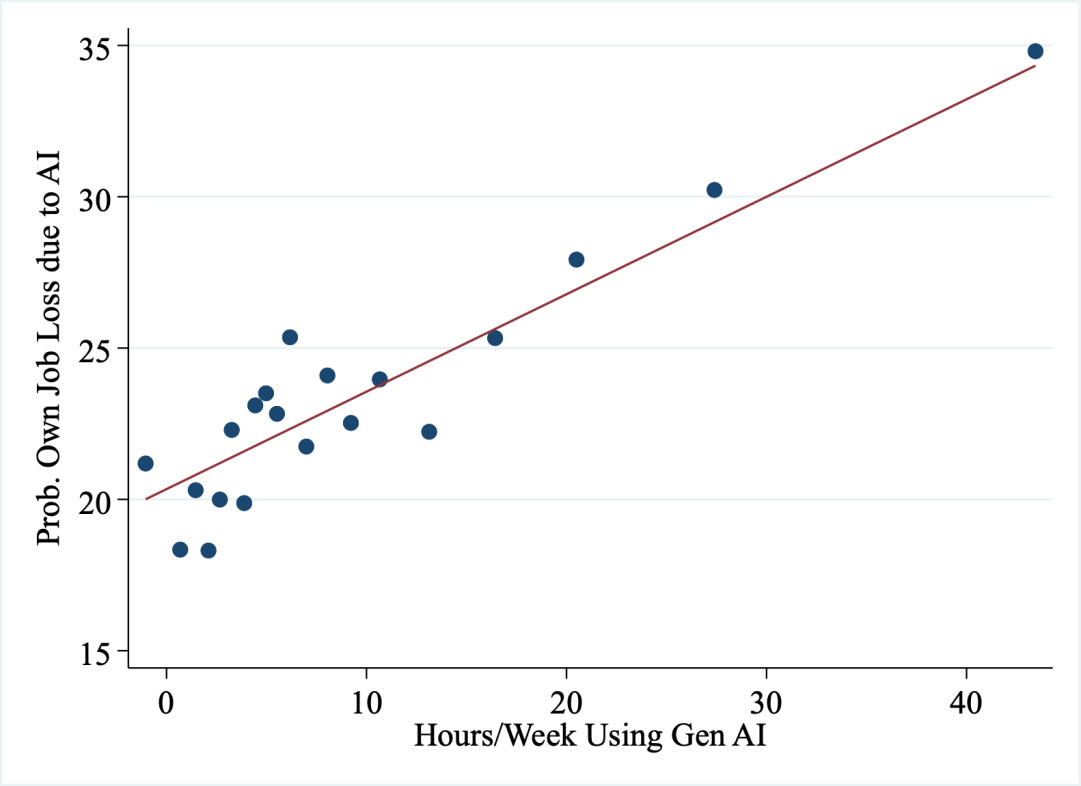


(b) by Predicted GenAI Exposure



Source: Authors' survey. *Note*: Top panel shows histograms of subjective job loss perceptions due to generative AI in the next two years among non-AI adopters at work but who have heard of generative AI (blue) and AI users at work (red). Bottom panel is split by above- and below-median predicted AI exposure. Survey sample is weighted population weights by age-sex-education-race from the Current Population Survey.

Figure 3: Intensive Margin Relationship Between AI Usage at Work & AI Job Loss Fear



Note: Binscatter shows relationship between two-year subjective job loss rate due to AI with respect to the reported number of hours per week the respondent uses AI at work, among the sample that uses AI at work at all. Regression includes occupation group (2-digit SOC) fixed effects and controls for gender, race, age, and education. From regression, main coefficient (slope) of interest equals 0.322 ($t = 9.33$) and intercept equals 37.1.

another four months, allowing us to observe month-to-month labor market transitions for matched individuals. We follow the matching procedure of Shimer (2012), linking individuals across months using age, sex, and race identifiers, and dropping observations with implausible transitions.

We define three transition rates. The job-to-job (J2J) transition rate captures the share of employed workers in month t who report a different employer in month $t + 1$ while remaining employed throughout. The employed-to-unemployed (EU) transition rate captures separations into unemployment, and the employed-to-nonemployment (EN) rate captures transitions out of the labor force entirely, including both unemployment and labor force exit. We construct these rates at the 3-digit SOC occupation level by aggregating individual-level transitions using CPS survey weights.

To measure on-the-job search intensity, we use the CPS supplement questions on job search activity among the employed, available in selected monthly supplements. Specifically, we use responses to questions about whether employed workers searched for a new job in the prior four weeks as our extensive margin measure of on-the-job search.

Our identification strategy mirrors the labor demand analysis. We interact each occupation's predicted AI exposure score from Eloundou et al. (2024) with an indicator for the post-ChatGPT period

(December 2022 onward), controlling for occupation fixed effects and calendar quarter fixed effects throughout. Standard errors are clustered at the 3-digit SOC occupation level. The pre-period spanning January 2019 through November 2022 provides a baseline against which to evaluate changes in mobility patterns. We exclude the period March 2020 through June 2020 from all specifications to avoid contaminating estimates with pandemic-driven labor market disruptions.

3.1.2 LightCast/Burning Glass Institute Job Openings

We analyze labor demand by using Lightcast's comprehensive job posing data. We aggregate job postings to year-month, 4-digit SOC occupation, and 4 digit NAICS industry level, generating monthly counts of job postings for each unique occupation-industry combination for a period from February 2021 to February 2025. We combine this data with our survey responses to build a panel dataset of 593,280 observations.

3.1.3 Occupational Employment Statistics (OES) Jobs and Wages

We complement our job posting analysis with employment and wage data from the Bureau Of Labor Statistics' Occupational Employment and Wage Statistics (OEWS). The OEWS provides annual measures of (i) total employment, (ii) median hourly wages, and (iii) median annual salary for across different SOC occupations. We aggregate these codes to the 4-digit SOC code level to maintain consistency across our analyses and combine with the same survey results used in our job posting analyses.

3.2 Empirical Methodology

Our identification strategy aims to test the impact of generative AI usage on labor demand, total employment, and changes in both annual and hourly wages. We use the public release of Open AI's ChatGPT release on November 30, 2022 as our time-based treatment of generative AI usage. We implement both a continuous treatment and difference-in-difference methodology to compare labor market outcomes before and after the public release of ChatGPT.

For job posting analysis using the monthly Lightcast data, the treatment period begins in December of 2022. For our employment and wage analysis using the OEWS data, we modify the treatment period to be 2023 to accommodate the annual structure of the data. This maintains the same underlying identification strategy while adopting the annual frequency of the OEWS data.

We estimate the impact of generative AI usage on labor market demand using two complementary specifications. Our continuous treatment regression exploits the full variation in AI usage across occupations:

$$\text{Postings}_{oit} = \alpha + \beta \left(\text{GenAIUsage}_o \times \text{Post}_t\right) + \gamma_{oi} + \delta_t + \varepsilon_{oit}, \tag{1}$$

Where total_postingsoit represents the number of job postings for occupation o in industry i during month t. Our key variable is AI usage, where in the control period is given a value of 0, and in the treatment period is given the continuous usage in a given occupation. γ_{oi} represents occupation-industry level fixed effects and δ_t represent time-level fixed effects. We cluster standard errors and the occupation level. Our binomial regression follows the same equation above, but where treatment is replaced with the interaction of the usage- and time-level treatments.

We further estimate the impact of generative AI usage on employment levels and wages using the OEWS data with similar specifications to our job posting analyses. Our continuous treatment regression is as follows:

$$\text{Outcome}_{ot} = \alpha + \beta \left(\text{GenAIUsage}_o \times \text{Post}_t\right) + \gamma_o + \delta_t + \varepsilon_{ot}, \tag{2}$$

Where outcome can be either (1) total employment, (2) median hourly wage, or (3) median annual wage for occupation o in time t. The key explanatory variable is the same specification as in our job posting analysis, interacting post-2023 with our continuous AI usage measure. We include occupation and year fixed effects, clustering standard errors at the occupation level. Note, unlike out job posing analysis, there is not industry identifiers in the OEWS data, and we therefore cannot include them in our fixed effects. Our binomial regression follows the same format as the job posting binomial regression.

3.3 Results

Job Separations

Figure 4 presents event-study estimates of the effect of predicted AI exposure on employment-to-unemployment (EU) and employment-to-nonemployment (EN) transition rates, using the public release of ChatGPT in November 2022 as the treatment date. The estimates are based on a difference-in-differences design comparing occupations above and below the median level of employment-weighted predicted AI adoption.

Across both panels, coefficients are small in magnitude and statistically indistinguishable from zero throughout the sample period. Importantly, we find no evidence of differential pre-trends prior to the ChatGPT release, supporting the parallel trends assumption underlying our design. In the post-period, point estimates remain centered around zero with confidence intervals that rule out economically meaningful increases in separation rates. These results hold for both the EU transition rate—the standard margin of job loss into unemployment—and the broader EN rate that includes transitions out of the labor force entirely.

The null result on separations is perhaps the paper's starkest empirical finding. Despite the heightened displacement fears documented in Section 2.1, workers in AI-exposed occupations are not actually losing their jobs at higher rates than their counterparts in less exposed occupations. This divergence between perceived and realized displacement risk is central to the interpretation we

develop in Section 4: it is the belief, not the realization, that is doing the behavioral work.

Job-to-Job Transitions

While separation rates show no response to AI exposure, job-to-job (J2J) mobility tells a notably different story. Figure 5 plots event-study estimates of the effect of predicted AI exposure on the J2J transition rate, again relative to the ChatGPT release date. In contrast to the separation results, point estimates in the post-period are positive and, while imprecisely estimated, suggest that workers in more AI-exposed occupations have if anything become *more* likely to switch employers following the diffusion of generative AI, not less.

This pattern is consistent with the reallocation mechanism in our model. Recall that the model predicts workers in AI-exposed occupations will increase on-the-job search when they perceive their current match as vulnerable to displacement—provided that alternative occupations offer a sufficiently attractive outside option. The positive post-period J2J coefficients are consistent with workers acting on that incentive, moving toward employers or roles they perceive as less exposed to automation risk. The pre-period estimates show no clear trend, lending further credibility to the design.

We interpret the combination of stable EU rates and elevated J2J transitions as evidence that the primary near-term labor market response to generative AI diffusion has been reallocation rather than displacement. Workers are reshuffling across employers and, likely, across tasks—but they are not being pushed into unemployment at higher rates. Whether this reallocation is efficient, in the sense of moving workers toward higher-productivity matches, or instead reflects precautionary churning driven by misperceived risk, is a question we cannot fully resolve with the current data. The positive wage effects documented in Section 3.4 are at least consistent with the former interpretation.

Vacancies

Across specifications, we find only suggestive evidence that Generative AI exposure significantly negatively affects job openings. Point estimates are small in magnitude and statistically indistinguishable from zero. Table 2 shows limited effect of predicted AI exposure on job postings. The monthly LightCast estimates in Table 1 show no discernible pre-trends prior to December 2022, consistent with the event study in Figure 6, supporting the validity of our difference-in-differences design.

Table 1: Comparison of Means Before and After the Event (Log Scale)

Variable	**N**	**Mean (Before)**	**Mean (After)**	**Approx. Change (%)**
ln(Total Postings)	508,789	6.2146	6.6846	47.00
ln(Total Employment)	32,380	6.9078	7.6009	69.31
ln(Median Hourly Wage)	32,380	3.9220	4.1027	18.07
ln(Median Annual Salary)	33,081	11.5617	11.7361	17.44

Note: The analysis covers a 48-month period from March 2021 to February 2025. The 'before' and 'after' periods are defined relative to December 2022.

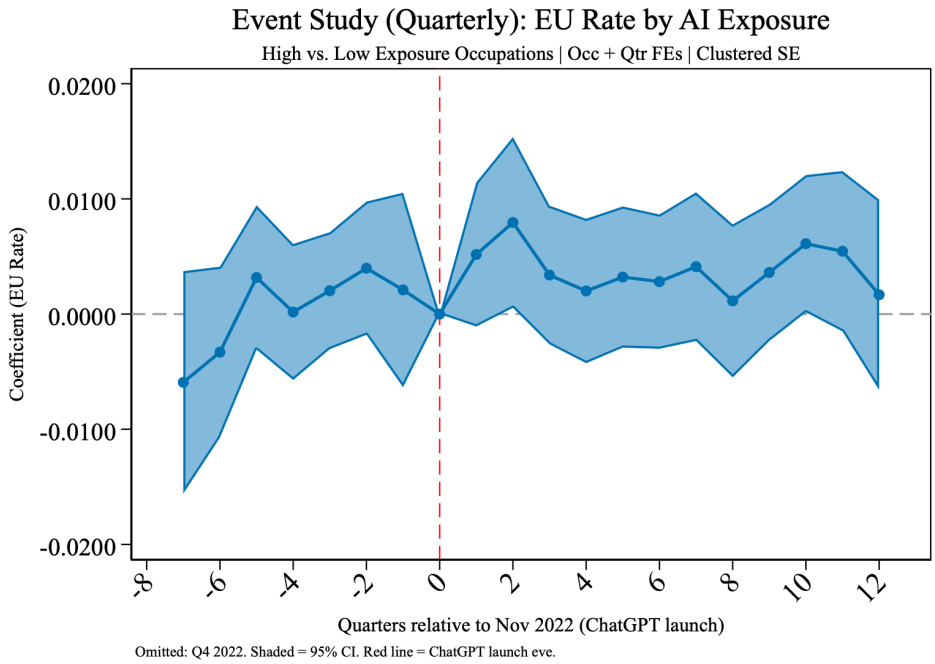
Table 2: Predicted AI Exposure: Difference-in-Difference Results

	(1) **log(postings)**	**(2)** **log(postings)**	**(3)** **log(postings)**
Post × Treatment (DiD)	-0.040	0.058	-0.069
	(0.068)	(0.107)	(0.067)
Treatment Group	0.767	0.868	—
	(0.568)	(0.615)	—
Post-Period	-0.267***	-0.367***	—
	(0.053)	(0.089)	—
Controls	No	Yes	Yes
Observations	3,086	3,086	3,086
R^2	0.031	0.058	0.980

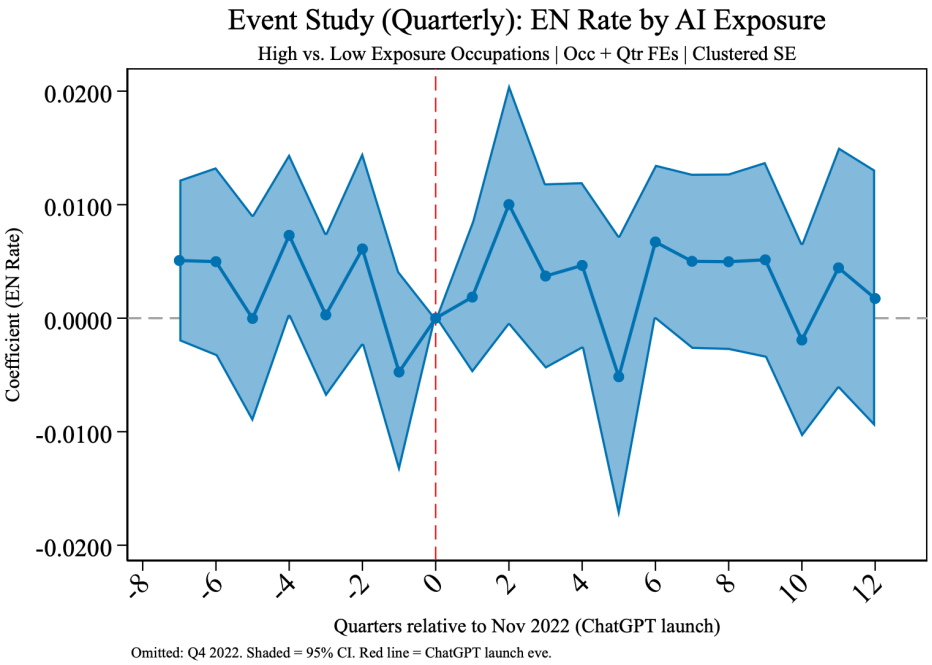
Notes: Standard errors in parentheses. In Model 3, Treatment and Post indicators are absorbed by fixed effects (—).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 4: Effect of Predicted AI Exposure on Job Separations

(a) U.S. Labor Force Adoption Over Time

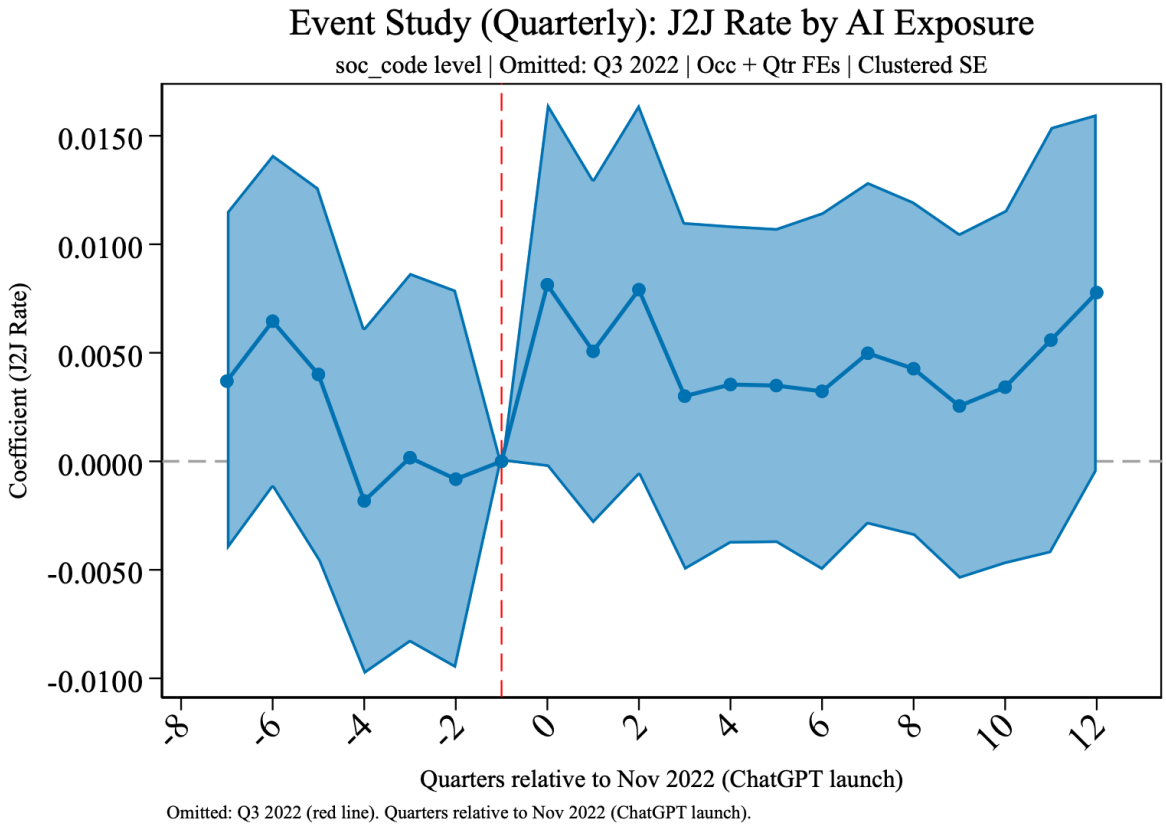


(b) Adoption at Work Across Countries in 2025



Note: Dependent variable on top panel is *EU* transition rate and on bottom panel is *EN* (NILF) transition rate. Control group are occupations with below median level of employment-weighted predicted AI adoption (0.33). Treatment group includes occupations with above median level of employment-weighted predicted AI adoption.

Figure 5: Effect of Predicted AI Exposure on Job-to-Job Switching



4 Model of On-the-Job Search with AI-Induced Reallocation

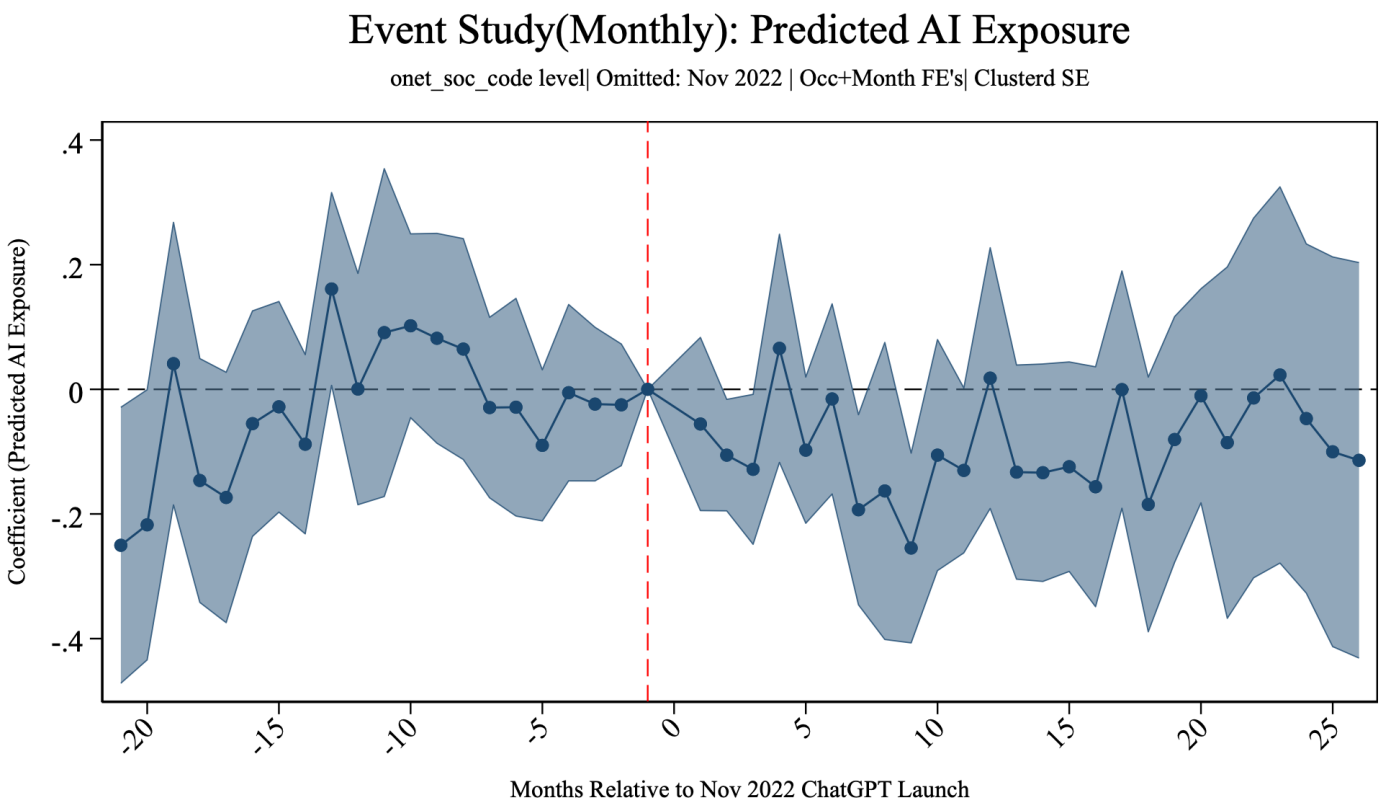
We develop a simple extension of the Burdett and Mortensen (1998) on-the-job search framework in which workers face perceived displacement risk from generative artificial intelligence (AI). In this setup, workers in AI-exposed occupations view their current match as especially vulnerable to AI-related displacement and can partially escape that risk by moving to a new employer. The result is that higher perceived AI risk can increase on-the-job search and job-to-job mobility.

4.1 Environment

Time is continuous. A continuum of workers of measure m and a continuum of firms participate in the labor market. Firms post wages and workers search sequentially.

Workers can be unemployed (U) or employed (E). Unemployed workers receive flow utility b. Employed workers earn wage w.

Figure 6: Effect of Predicted AI Exposure on Job Vacancies



Note: Dependent variable is log(vacancies). Event study regression run at monthly level with Nov 2022 as the omited period. Occupation and month fixed effects.

Job offers arrive according to Poisson processes:

$$\lambda_0 = \text{arrival rate of offers while unemployed,} \tag{3}$$

$$\lambda_1 s_i = \text{arrival rate of offers while employed,} \tag{4}$$

where $s_i \in [0,1]$ is worker i's search intensity. Workers discount the future at rate r.

4.2 AI Displacement Risk and Sectoral Reallocation

Workers differ in their exposure to generative AI across occupations. Let $j \in \{x, u\}$ denote whether a worker in occupation-type j is employed in an AI-exposed (x) or less exposed (u) occupation.

Matches are destroyed at exogenous rates δ_j, which may differ across occupation-type (exposed vs. less exposed) but are constant over time. These represent the true separation rates governing employment transitions.

Workers, however, form beliefs about job stability. Let $\widehat{\delta_j}$ denote the perceived separation rate in occupation j. Consistent with our survey evidence in Figure 2b, we assume that workers perceive jobs in AI-exposed occupations as more vulnerable to displacement:

$$\widehat{\delta}_x > \widehat{\delta}_u. \tag{5}$$

Crucially, perceived separation risk may diverge from realized separation risk. In particular, workers may overestimate the risk of job loss in AI-exposed occupations:

$$\hat{\delta}_x > \delta_x. \tag{6}$$

The model can also account for job loss misperceptions across both sectors (i.e. workers perceive uniformly too-high job loss rates $\widehat{\delta_j} > \delta_j$), and the implications of this formulation remain the same so long as workers misperceive a *differentially* large layoff risk for the AI-exposed sector (namely, that $\widehat{\delta}_x - \delta_x > \widehat{\delta}_u - \delta_u$). Workers can reallocate across occupations through on-the-job search. Job offers arrive from occupation $k \in \{x, u\}$ with probability π_k, and wages are drawn from occupation-specific distributions $F_k(w)$.

4.3 Worker Value Functions

Let V_U denote the value of unemployment, and $V_j(w)$ the value of employment in occupation $j \in \{x, u\}$.

The value of unemployment is

$$rV_U = b + \lambda_0 \sum_{k \in \{x,u\}} \pi_k \int \max\{V_k(z), V_U\}\, dF_k(z) - \lambda_0 V_U. \tag{7}$$

Workers make decisions based on perceived separation risk. Accordingly, value functions are defined with respect to perceived transition rates. The perceived value of employment in occupation j is

$$rV_j(w) = w + \lambda_1 s_i \sum_{k \in \{x,u\}} \pi_k \int \max\{V_k(z) - V_j(w), 0\}\, dF_k(z) + \widehat{\delta}_j\big(V_U - V_j(w)\big). \tag{8}$$

4.4 Search Intensity Choice

Workers choose search intensity $s_i \in [0,1]$ subject to convex search costs $c(s_i)$, with $c'(s_i) > 0$ and $c''(s_i) > 0$.

Search determines the arrival rate of outside offers while employed. Given perceived values, workers choose s_i to maximize the expected gains from on-the-job search:

$$\max_{s_i} \ \lambda_1 s_i \sum_{k \in \{x,u\}} \pi_k \int \max\{V_k(z) - V_j(w), 0\}\, dF_k(z) - c(s_i). \tag{9}$$

The first-order condition is

$$c'(s_i) = \lambda_1 \sum_{k \in \{x,u\}} \pi_k \int \max\{V_k(z) - V_j(w), 0\}\, dF_k(z). \tag{10}$$

Thus, optimal search intensity depends on the perceived surplus from switching jobs. When the perceived value of alternative occupations rises relative to the value of remaining in one's current occupation $V_j(w)$, workers increase on-the-job search.

4.5 Implications of Higher Perceived AI Risk

We now consider how changes in perceived separation risk affect worker behavior.

Consider a worker currently employed in an AI-exposed occupation x. An increase in perceived AI-related displacement risk raises $\widehat{\delta}_x$, reducing the perceived value of remaining in that occupation, $\frac{\partial V_x(w)}{\partial \widehat{\delta}_x} < 0$, since higher perceived separation risk increases the likelihood of transitioning to unemployment. On the other hand, the perceived value of employment in less exposed occupations (u) is unaffected, meaning $\frac{\partial V_u(z)}{\partial \widehat{\delta}_x} = 0$.

Therefore, an increase in perceived AI-related displacement risk raises the perceived surplus from reallocation:

$$\frac{\partial}{\partial \widehat{\delta}_x}\Big[V_u(z) - V_x(w)\Big] > 0. \tag{11}$$

From the search first-order condition in Section 4, this increase in surplus raises the expected gains from search and implies

$$\frac{\partial s_i}{\partial \widehat{\delta}_x} > 0. \tag{12}$$

Hence, workers in AI-exposed occupations respond to higher perceived displacement risk by increasing on-the-job search and reallocating toward less exposed occupations.

Importantly, these behavioral responses arise from perceived changes in job stability and need not reflect changes in actual separation rates. Even if true separation rates δ_x remain unchanged, increases in perceived risk $\widehat{\delta}_x$ can generate increased job-to-job mobility and occupational reallocation.

4.6 Testable Predictions

The model yields three empirical predictions.

First, workers in more AI-exposed occupations should exhibit higher on-the-job search intensity after the release of ChatGPT.

Second, job-to-job mobility should rise in more AI-exposed occupations, especially where workers perceive stronger opportunities to move into more AI-complementary roles.

Third, wage growth may be strongest among movers if job-to-job transitions serve as a margin of reallocation toward firms and tasks that use AI more productively.

This interpretation suggests that perceived AI displacement risk need not suppress labor market dynamism; under plausible conditions, it can instead accelerate worker reallocation.

5 Conclusion

This paper investigates whether worker beliefs about generative AI-induced displacement shape labor market behavior, even in the absence of large-scale realized job loss. Using a large-scale, multi-wave survey of nearly 20,000 U.S. workers fielded between December 2024 and March 2026, we document three main findings.

First, generative AI adoption at work reached approximately 36% of the U.S. labor force by late 2025, with adoption strongly stratified by education and income. Workers with graduate degrees and those earning above $200,000 annually adopt at rates approaching 50%, while workers without college degrees adopt at roughly half that rate. These adoption patterns closely track the theoretical AI-exposure scores of Eloundou et al. (2024), with occupations predicted to be most exposed to large language models also exhibiting the highest realized adoption rates.

Second, despite limited aggregate evidence of AI-driven displacement—job openings and employment levels in high-exposure occupations show no statistically significant decline following the public release of ChatGPT—American workers report substantial subjective job loss risk attributable to generative AI. The average worker believes there is roughly a one-in-five chance of being displaced by AI over the next two years, a figure that approaches the baseline annual job loss rate from all causes combined as measured by the New York Fed's Survey of Consumer Expectations. Workers in more AI-exposed occupations and those who use AI more intensively on the job report the highest displacement fears, suggesting that familiarity with the technology's capabilities amplifies rather

than alleviates concern.

Third, we show that these beliefs have measurable consequences for labor market behavior. Workers in more AI-exposed occupations exhibit lower rates of on-the-job search and reduced job-to-job mobility following the diffusion of generative AI, consistent with the predictions of our on-the-job search model. In that model, workers who perceive their current match as vulnerable to AI displacement rationally reduce search intensity when switching jobs resets tenure protections and introduces uncertainty about AI exposure in a new employment relationship. By contrast, wages have risen in high-exposure occupations, consistent with generative AI functioning as a labor complement that raises worker productivity rather than a near-term substitute.

Taken together, these results suggest that the near-term labor market footprint of generative AI operates at least as much through a belief channel as through direct displacement. Perceptions of automation risk can dampen labor market dynamism—reducing the reallocation that ordinarily matches workers to their most productive uses—even when realized displacement remains limited. This distinction has implications for how policymakers and firms should interpret aggregate labor market data in an era of rapid AI diffusion. Stable employment and vacancy figures are not, by themselves, evidence that workers have adjusted smoothly to the technology; they may instead reflect a suppression of the job-switching behavior through which labor markets typically absorb technological change.

Several important questions remain. Our survey evidence covers a period in which generative AI capabilities were still evolving rapidly, and worker perceptions may update as the technology matures and its labor market effects become better understood. Whether the belief-driven dampening of labor market dynamism documented here is transitory—dissipating as uncertainty resolves—or persistent is an open question that future work with longer panels will be better positioned to address. Additionally, while our model and empirical results emphasize a precautionary motive for reduced search, alternative mechanisms, including reduced outside options due to AI-driven compression of skill premia, warrant further investigation. We leave these questions to future research.

References

AMODEI, D. (2026): "The Adolescence of Technology," *Essay on Anthropic CEO's website*.

BICK, A., A. BLANDIN, AND D. J. DEMING (2026): "The rapid adoption of generative AI," *Management Science*.

BLOOMBERG NEWS (2025): "KPMG Staff to Be Rated on AI Usage in Yearly Performance Reviews," Accessed March 11, 2026.

BRYNJOLFSSON, E., B. CHANDAR, AND R. CHEN (2025): "Canaries in the Coal Mine?: Six Facts about the Recent Employment Effects of Artificial Intelligence," .

BRYNJOLFSSON, E., T. MITCHELL, AND D. ROCK (2018): "What can machines learn and what does it mean for occupations and the economy?" in *AEA papers and proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 108, 43–47.

BURDETT, K. AND D. T. MORTENSEN (1998): "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39, 257–273.

ELOUNDOU, T., S. MANNING, P. MISHKIN, AND D. ROCK (2024): "GPTs are GPTs: Labor market impact potential of LLMs," *Science*, 384, 1306–1308.

HARTLEY, J., F. JOLEVSKI, V. MELO, AND B. MOORE (2026): "The labor market effects of generative artificial intelligence," *Available at SSRN*.

HENDREN, N. (2017): "Knowledge of future job loss and implications for unemployment insurance," *American Economic Review*, 107, 1778–1823.

HOSSEINI MAASOUM, S. M. AND G. LICHTINGER (2025): "Generative ai as seniority-biased technological change: Evidence from us resume and job posting data," *Available at SSRN 5425555*.

HUMLUM, A. AND E. VESTERGAARD (2025): "Large language models, small labor market effects," Tech. rep., National Bureau of Economic Research.

KARAHAN, F., B. MOORE, AND L. PILOSSOPH (2019): "Expecting the unexpected: Job losses and household spending," Tech. rep., Federal Reserve Bank of New York.

LEHNER, L., I. TALESARA, AND A. WICKARD (2025): "Worker Beliefs about Layoff Risk," .

LICHTINGER, J. AND C. MAASOUM (2025): "Seniority-Biased Technological Change: Generative AI and the Labor Market," *Available at SSRN*.

LIU, Y., H. WANG, AND S. YU (2025): "Labor Demand in the Age of Generative AI: Early Evidence from the US Job Posting Data," *Available at SSRN 5504741*.

SHIMER, R. (2012): "Reassessing the ins and outs of unemployment," *Review of Economic Dynamics*, 15, 127–148.

STEINBERGER, M. (2026): "Mass Hysteria. Thousands of Jobs Lost. Just How Bad Is It Going to Get?" *The New York Times*, opinion. Accessed: 2026-03-10.

STEPHENS JR, M. (2004): "Job loss expectations, realizations, and household consumption behavior," *Review of Economics and Statistics*, 86, 253–269.

WEBB, M. (2019): "The impact of artificial intelligence on the labor market," *Available at SSRN 3482150*.

Online Appendix Materials

A Additional Figures and Tables 23

B Model Appendix 34

C Survey Instrument 35

A Additional Figures and Tables

With respect to education, we find in our December 2024 survey wave (consistent with the surveys of Pew and Bick, Blandin, and Deming (2024)) more educated workers are more likely to use Generative AI (Figure 4). Nearly 50% of those in the sample with a graduate degree use Generative AI. 37% of those with a college degree as their highest level education attained use Generative AI. Roughly 20% of those who are only high school graduates or have some college use Generative AI.

With respect to industries, there is a varied picture (Figure 5). Industries most likely to use Generative AI at work include "Information Services" and "Management of Companies" (using Generative AI tools more than 60% of the time) and "Real Estate, Rental, Leasing", "Construction", and "Education" (using such tools more than 40% of the time). Meanwhile workers in industries such as "Agriculture, Forestry, Fishing", "Mining, Quarrying, Oil & Gas", and "Government & Military" use it far less.

Generative AI use is increasing with incomes. However, the increase begins to happen at income levels around above $50,000 where approximately 20% of individuals below this amount say they have used Generative AI at work. Nearly 50% of workers making above $200,000 annually use Generative AI at work.

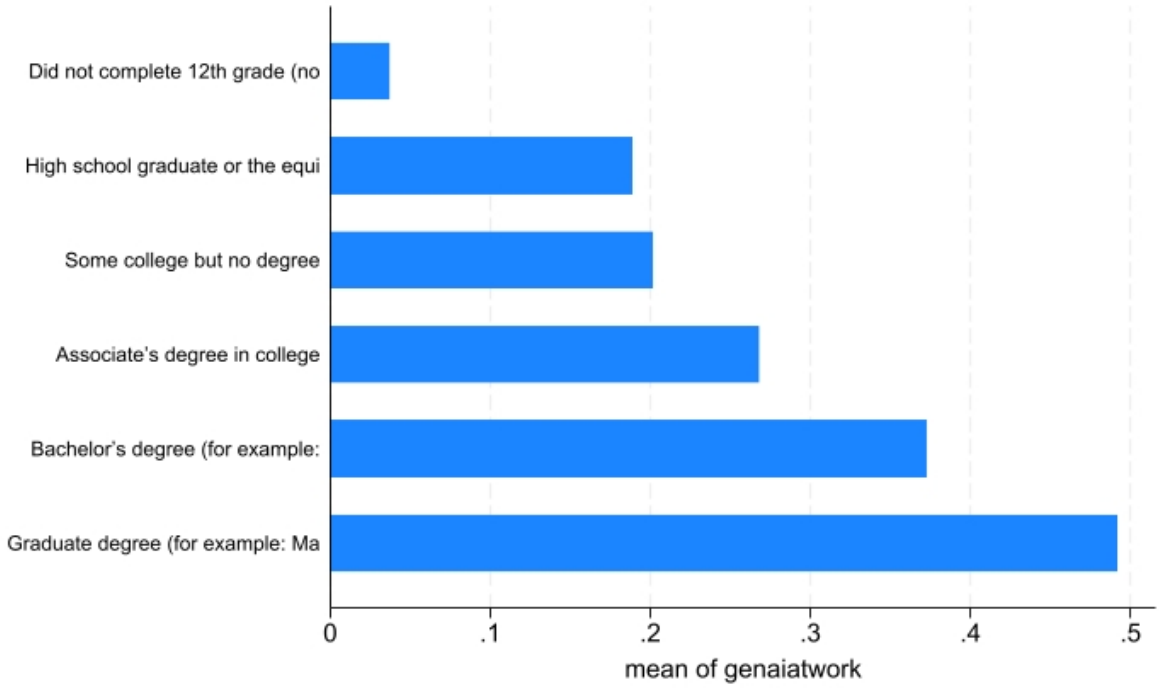
With respect to gender, a greater proportion of men claim to use Generative AI at work than women in the survey. 38.0% of men report to use Generative AI at work in the sample while 27.8% of women report to use Generative AI at work. With respect to race, we find that Hispanics are most likely to use Generative AI, followed by Black, Other, and White individuals in the sample.

ChatGPT remains the most popular Generative AI tool, closely followed by Gemini. The next most popular tools are GitHub Copilot and Scribe.

On the extensive margin, we found that 30.1% of survey respondents have used Generative AI at work at some point since Generative AI tools became public.

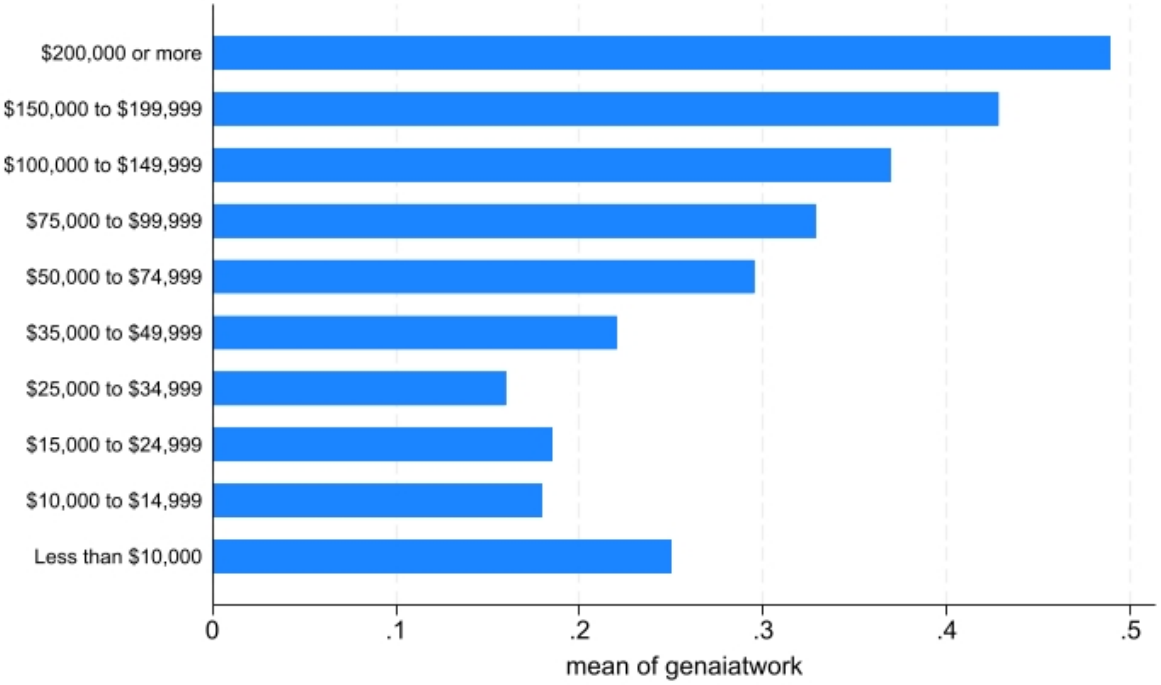
However, the intensive margin of using Generative AI tools at work also tells a different story. Conditional on using Generative AI at work, about 33% of workers use Generative AI five days per week at work (every weekday). Roughly 12% of Generative AI users use such tools at work only 1 day at work. About 17% and 18% of Generative AI users use Generative AI tools at work two and three days per week respectively. The number of hours at work each week using Generative AI tools tells a somewhat different story (the distribution of hours* using Generative AI tools at work however is perhaps less than the number of days worked would imply). While many workers may use Generative AI tools every weekday (or several weekdays), the total hours at work each week using Generative AI tools is reported to be in most cases less than 15 hours, suggesting there is intermittent use of Generative AI throughout the workday. On the whole, these results combined suggest that Generative AI is in part used by those with highest incomes and most education and is used in many cases each day at work, albeit intermittently.

Figure A.1: Generative AI Use By Level of Education



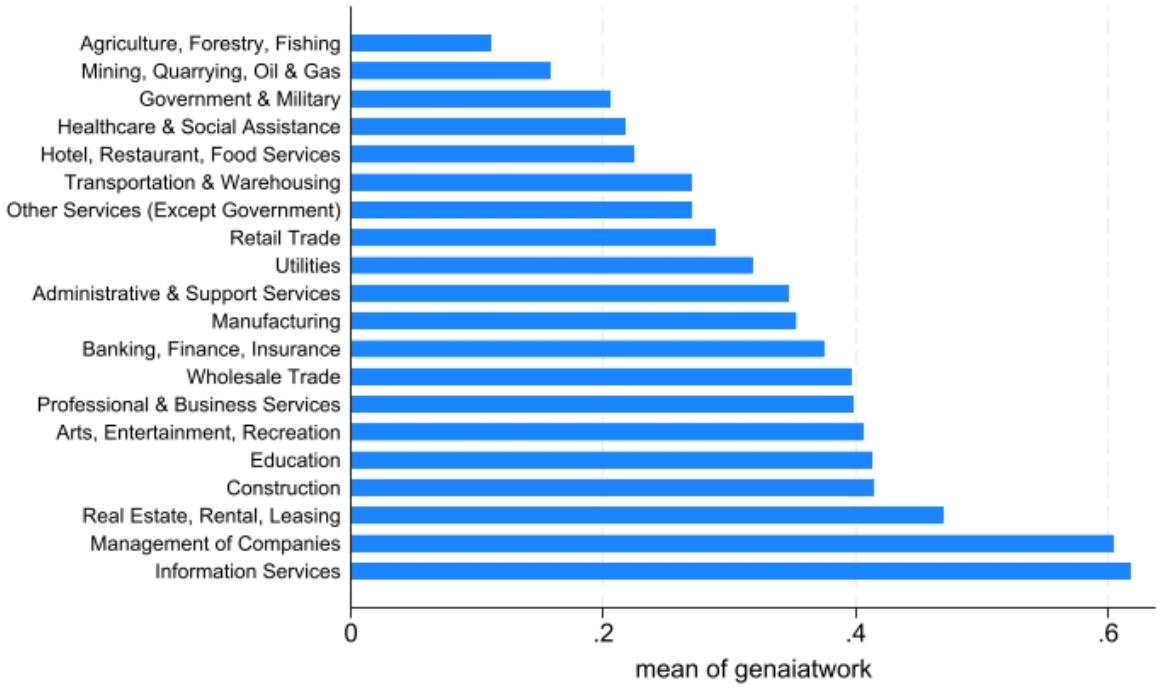
Source: Authors' survey.

Figure A.2: Generative AI Use By Income



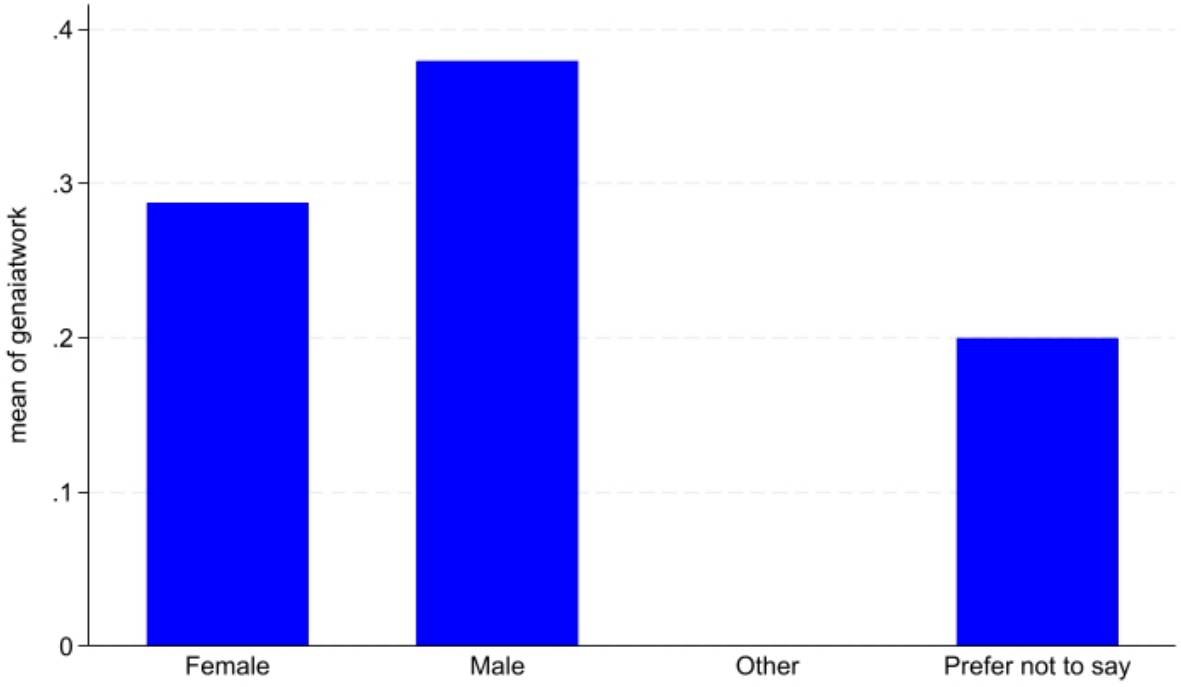
Source: Authors' survey.

Figure A.3: Generative AI Use By Industry



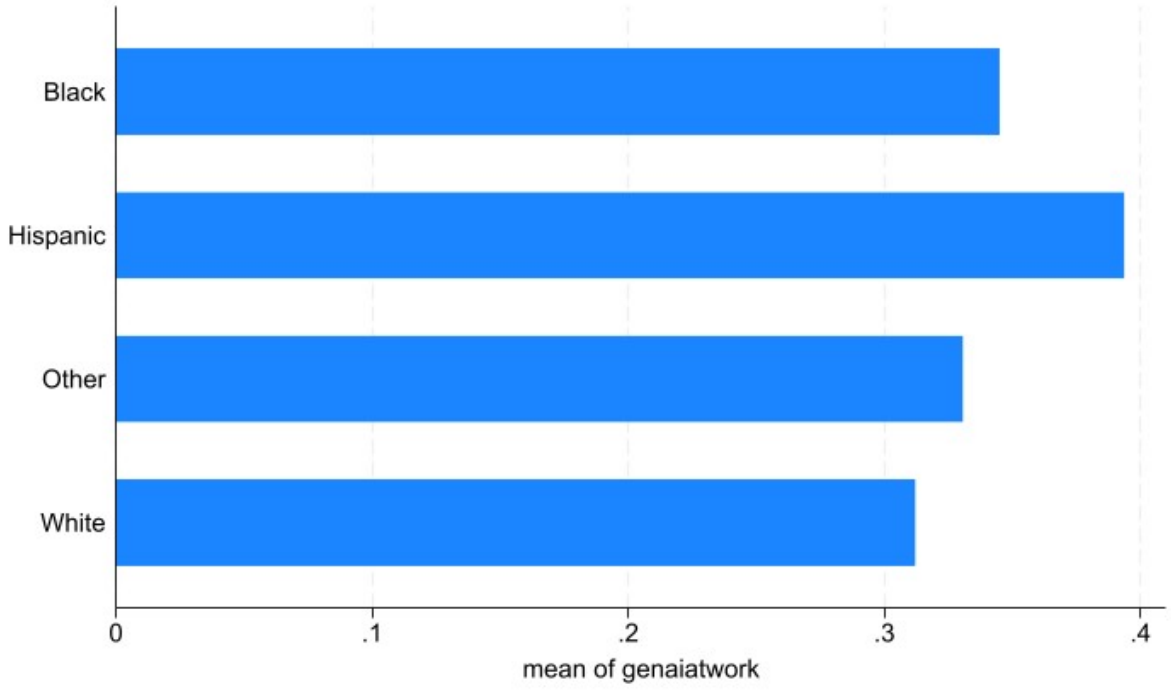
Source: Authors' survey.

Figure A.4: Generative AI Use By Gender



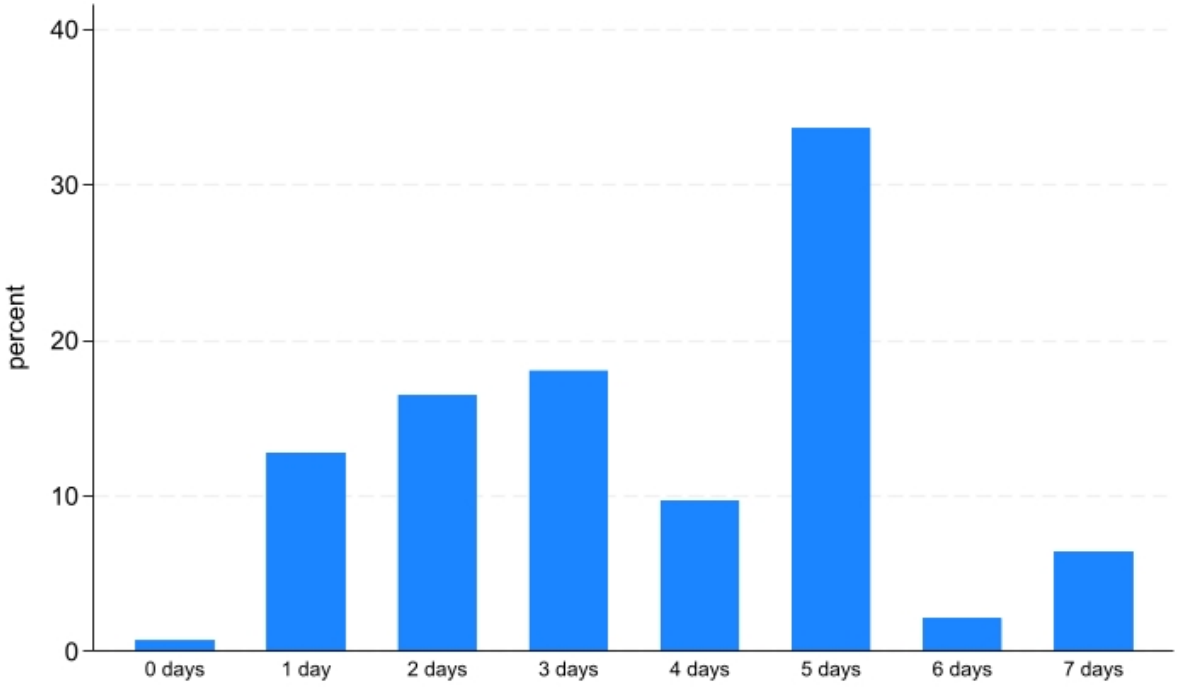
Source: Authors' survey.

Figure A.5: Generative AI Use By Race



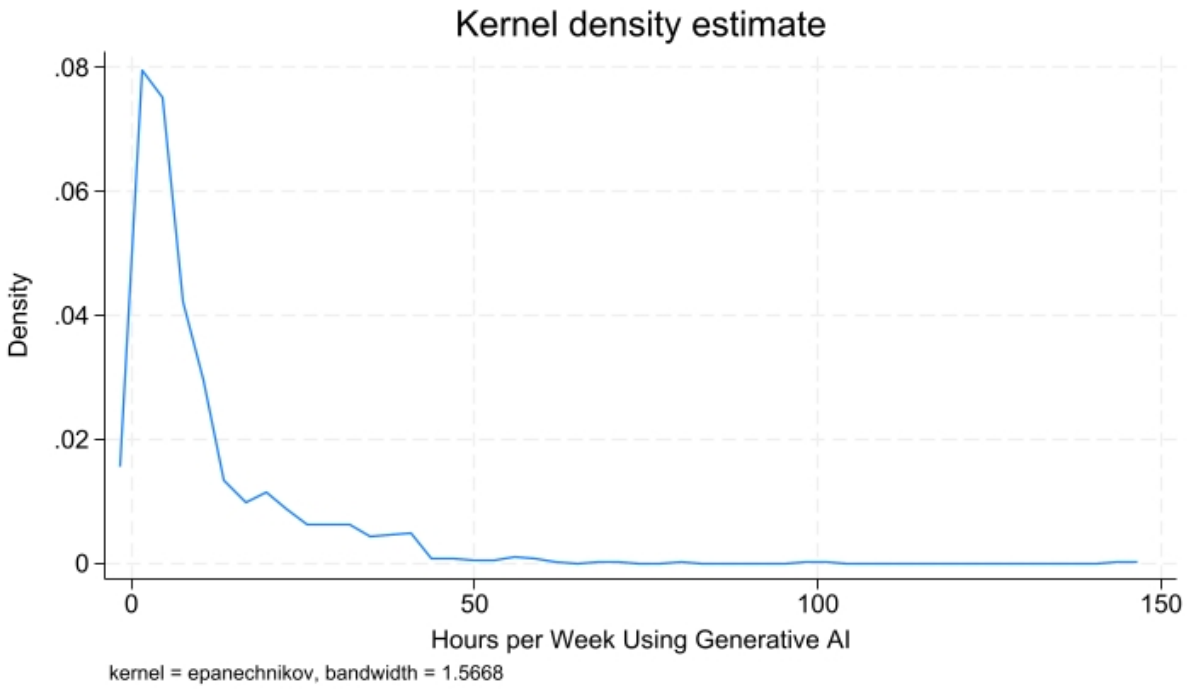
Source: Authors' survey.

Figure A.6: Number of Days At Work Each Week Using Generative AI Tools



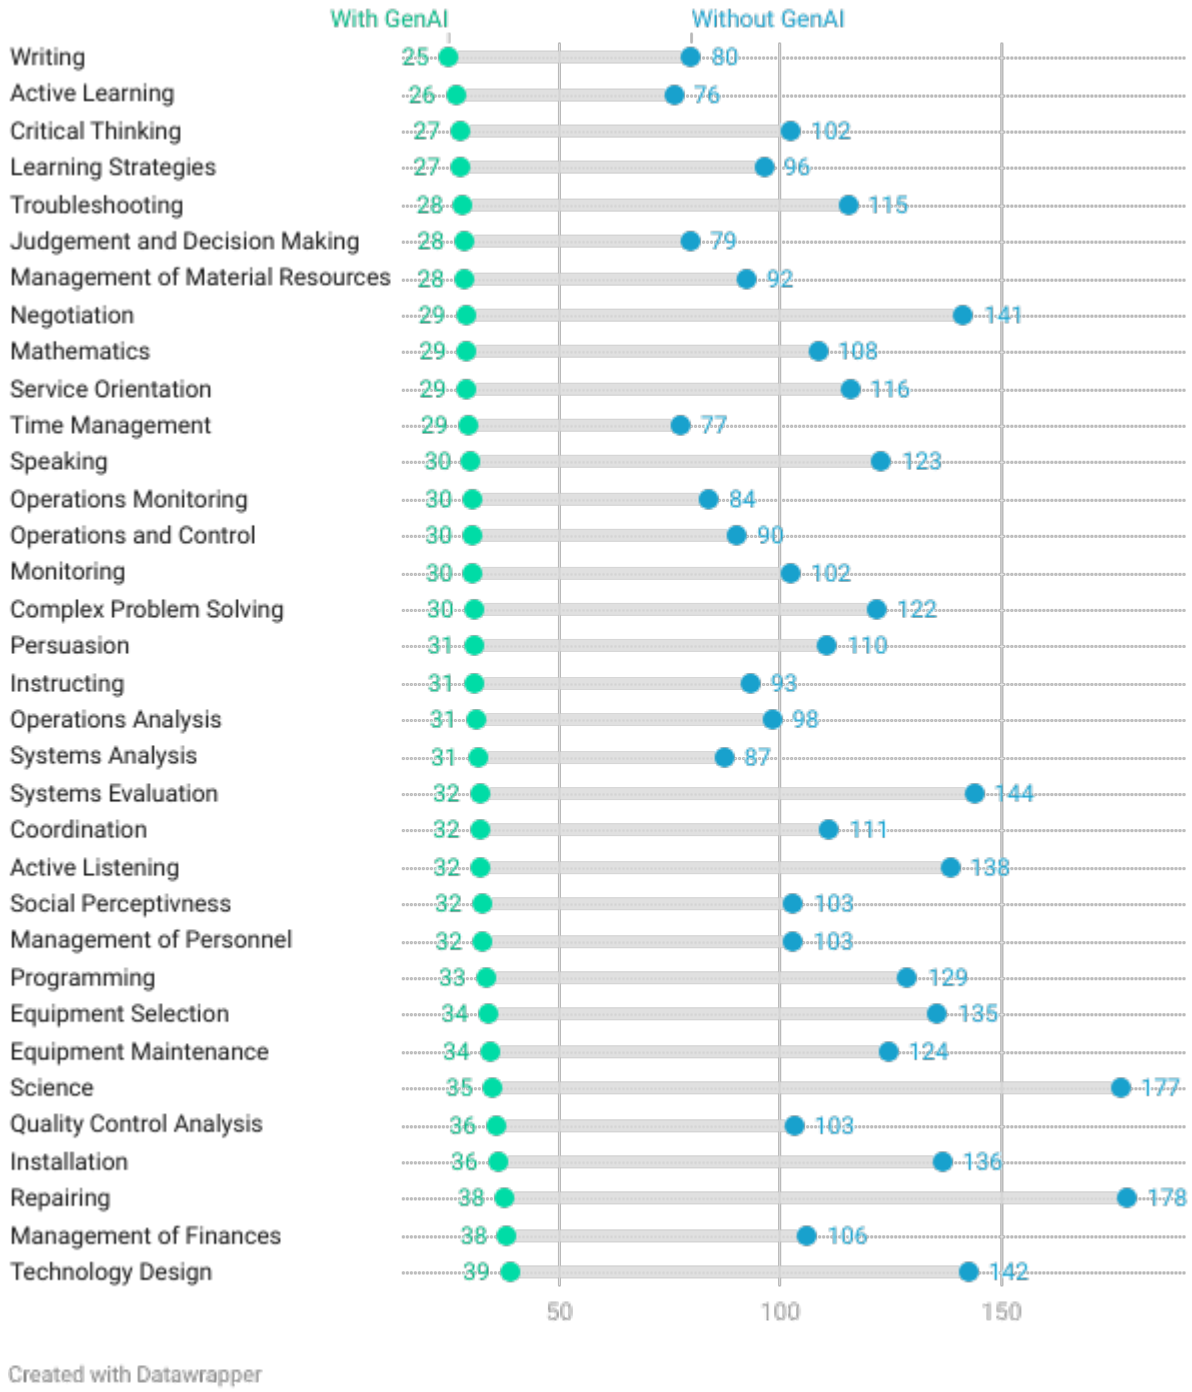
Source: Authors' survey.

Figure A.7: Number of Hours At Work Each Week Using Generative AI Tools



Source: Authors' survey.

Figure A.8: Average number of minutes to complete a task with and without Generative AI



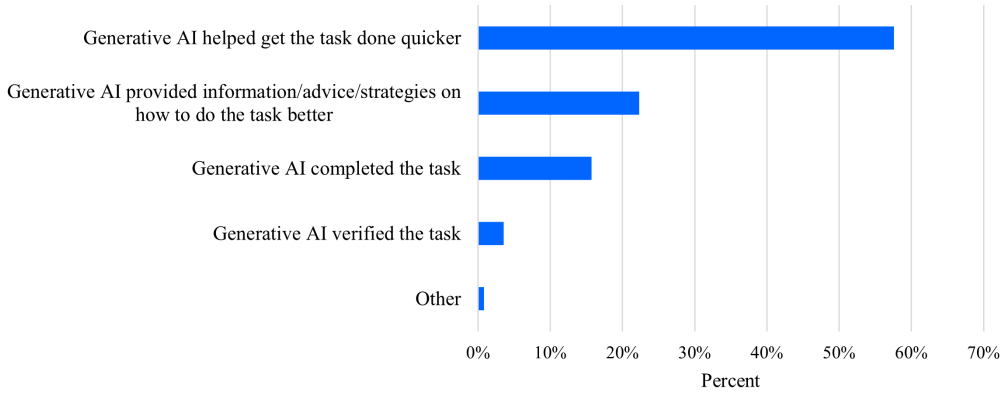
Source: Authors' survey.

Figure A.9: Productivity increase in completing various tasks using Generative AI



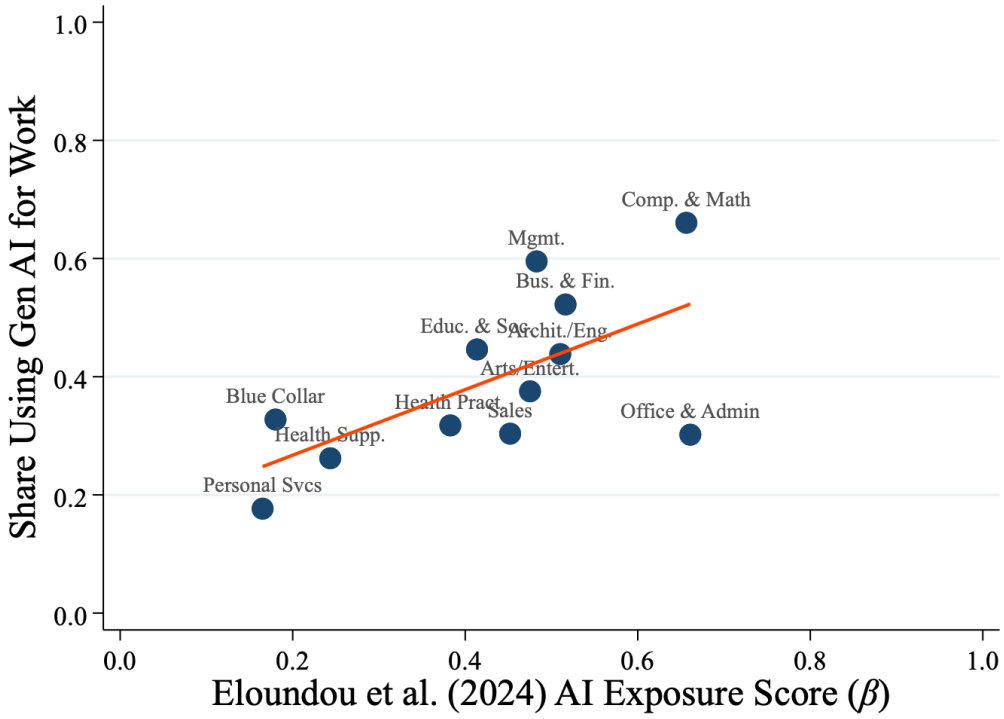
Source: Authors' survey.

Figure A.10: Uses of Generative AI to complete tasks



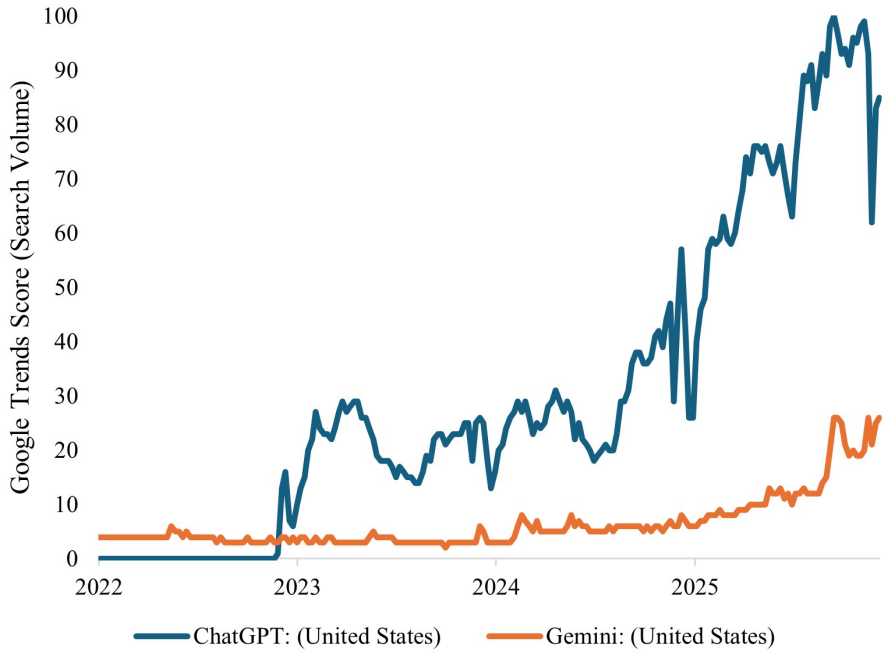
Source: Authors' survey.

Figure A.11: GenAI Adoption at Work vs. Predicted Exposure by Occupation



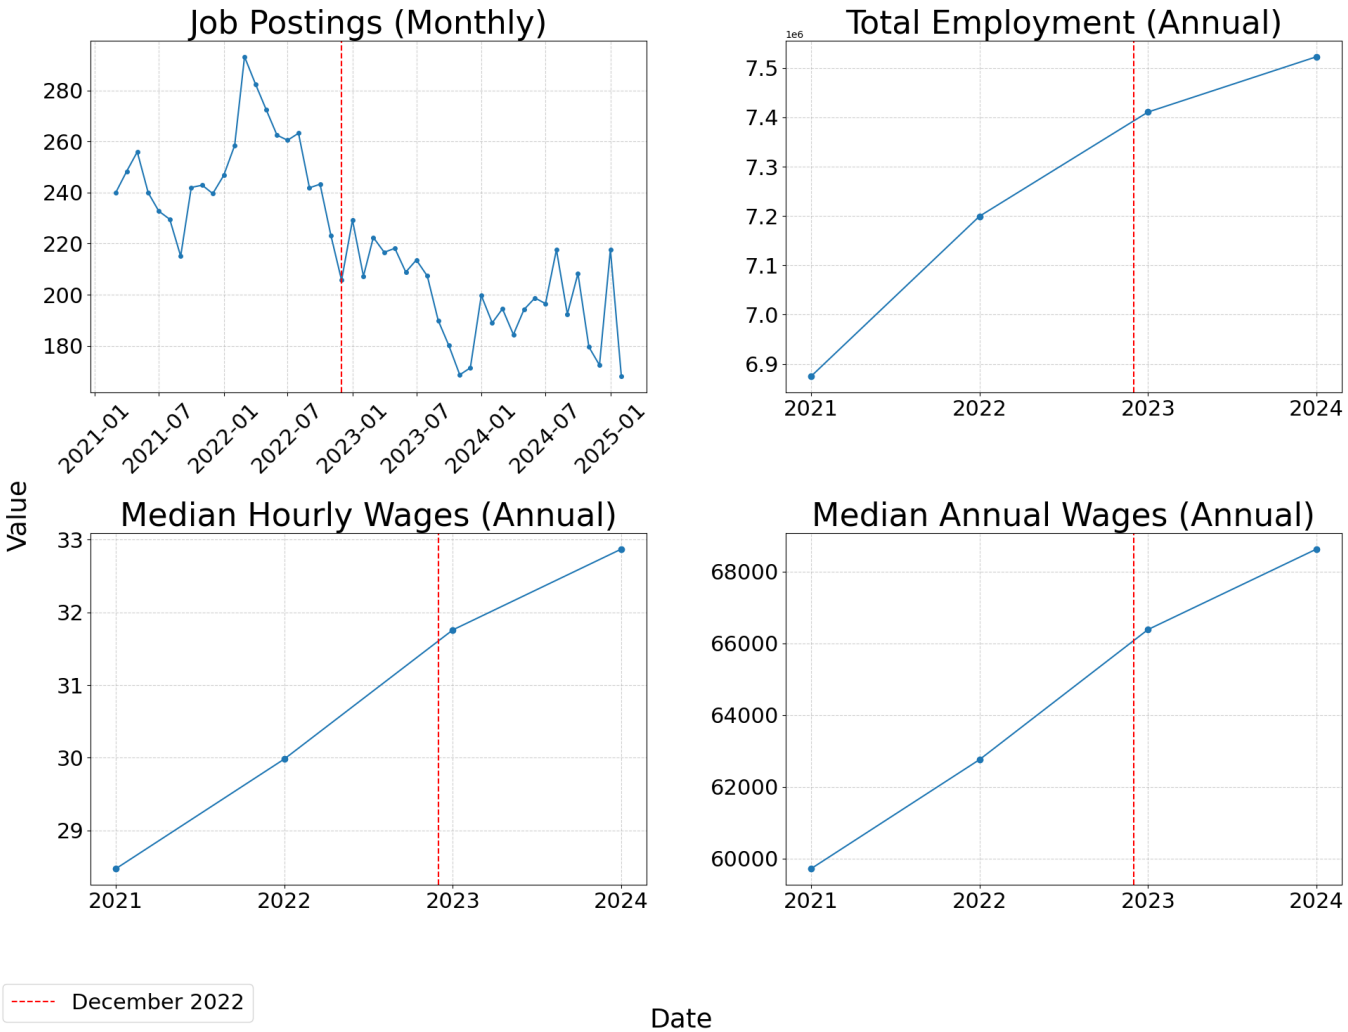
Source: Authors' survey & Eloundou, Manning, Mishkin, and Rock (2024) LLM β-exposure measure. Coefficient: 0.55; Intercept: 0.16.

Figure A.12: ChatGPT and Gemini Google Search Volume in the United States



Source: Google Trends.

Figure A.13: Evolution of Key Labor Market Indicators



Notes: The panel displays trends for four key labor market indicators. The top-left plot shows the monthly evolution of job postings (Source: Lightcast). The remaining plots show the annual evolution of total employment, median hourly wages, and median annual salary (Source: OES). All data represent the average across the 48 SOC-3 occupational groups in the sample. The vertical dashed line indicates the event date of December 2022.

Table A.1: AI Impact Exposure by Occupation Group (Ranked)

Occupation Group	**AI Exposure (% Yes)**
Supervisors of Construction and Extraction Workers	74.07
Computer Occupations	67.89
Advertising, Marketing, Promotions, Public Relations, and Sales Managers	67.57
Business Operations Specialists	66.81
Top Executives	66.81
Operations Specialties Managers	66.43
Other Management Occupations	62.50
Lawyers, Judges, and Related Workers	61.29
Postsecondary Teachers	55.74
Business Operations Specialists	52.56
Other Sales and Related Workers	51.72
Preschool, Elementary, Middle, Secondary, and Special Education Teachers	50.78
Supervisors of Office and Administrative Support Workers	47.83
Financial Specialists	47.71
Sales Representatives, Wholesale and Manufacturing	44.44
Healthcare Diagnosing or Treating Practitioners	44.19
Supervisors of Sales Workers	43.48
Health Technologists and Technicians	41.86
Supervisors of Production Workers	40.00
Assemblers and Fabricators	39.29
Home Health and Personal Care Aides; and Nursing Assistants, Orderlies, and Psychiatric Aides	38.60
Other Installation, Maintenance, and Repair Occupations	37.50
Other Educational Instruction and Library Occupations	36.17
Construction Trades Workers	32.50
Other Office and Administrative Support Workers	31.11
Other Healthcare Practitioners and Technical Occupations	31.03
Information and Record Clerks	30.43
Counselors, Social Workers, and Other Community and Social Service Specialists	29.73
Other Healthcare Support Occupations	29.31
Sales Representatives, Services	28.57
Retail Sales Workers	26.09
Financial Clerks	26.00
Material Moving Workers	22.73
Food and Beverage Serving Workers	20.83
Legal Support Workers	20.00

Continued on next page

Table A.1: AI Impact Exposure by Occupation Group (continued)

Occupation Group	**AI Exposure (% Yes)**
Motor Vehicle Operators	17.24
Other Personal Care and Service Workers	15.00
Metal Workers and Plastic Workers	13.64
Cooks and Food Preparation Workers	12.50
Secretaries and Administrative Assistants	11.43
Other Protective Service Workers	9.52

B Model Appendix

C Survey Instrument