The Labor Market Effects of Generative Artificial Intelligence and Job Loss Fears

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

July 21, 2026

Abstract

How do worker beliefs about layoff risk due to generative artificial intelligence compare to the realized labor market impact of the technology, and what shapes those beliefs? 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 in the mid-to-high 30% range among U.S. workers at the end of 2025. The average U.S. worker, regardless of whether they use the technology at work, believes there is a 20% chance of losing their job to generative AI within two years. Using a continuous-treatment difference-in-differences design, we find that these fears are not borne out, as job postings and layoffs in more exposed occupations show no statistically significant response to the diffusion of generative AI. Job loss fears instead track workers' direct experience of the technology. Fear rises with usage intensity (more hours per week) and is concentrated among those who deploy AI for core occupational functions such as programming rather than for drafting tasks such as writing. These results suggest that workers learn about AI's capacity to substitute for their labor by using it and form beliefs about displacement well ahead of any realized job loss.

*Hartley: UT-Austin School of Civic Leadership, hartleyj@stanford.edu, Jolevski: World Bank, fjolevski@worldbank.org, Melo: RAND Corporation, vitormelo95@gmail.com, Moore: University of Michigan, moorebd@umich.edu. Previously circulated under "The Labor Market Effects of Generative 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, Mihai Codreanu, Tyler Cowen, David Deming, Erik Hurst, Josh Gottlieb, Brian Kovak, 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 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 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 far less attention, however, is what workers themselves believe about AI-induced displacement, how those beliefs are formed, and how sharply they diverge from what the realized data show.

This paper pursues three interrelated questions. First, how widespread is generative AI adoption among U.S. workers? Second, what do workers believe about the risk of losing their jobs to the technology? Third, do generative AI's direct labor market effects square with workers' expectations?

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 25,100 total respondents. We document that GenAI adoption at work reached approximately 30% among U.S. workers by end of the first quarter of 2025 (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 25%. Estimates of U.S. workers' average annual subjective job loss probability (for any reason) are around 16–18%. In other words, workers fear AI-driven displacement at a rate that rivals their total perceived job loss risk from all sources. We also document that the more workers use GenAI on the job, the more they fear AI-induced job loss risk, even within occupational groups and demographic cells.

We then document that workers' substantial fears are not borne out in the data. Using a continuous-treatment difference-in-differences design that treats occupations as receiving different doses of AI exposure, we find no evidence that GenAI diffusion has reduced job openings or raised separation rates in exposed occupations. Point estimates for vacancy postings are negative and estimates for transitions into unemployment and non-employment are positive, but none is statistically distinguishable from zero, and the confidence intervals are tight enough to rule out anything close to the large-scale displacement that has animated public discussion. These nulls are robust to alternative sample restrictions, to a binary top-versus-bottom-quartile design, and to substituting an alternative occupational exposure measure. These findings on layoffs and vacancies align with contemporaneous administrative evidence from California unemployment insurance claims (Hyman et al., 2026) and from job postings data (Audoly et al., 2026).

Why, then, do fears remain so elevated? Our survey evidence points toward a learning-by-doing account, in which workers form beliefs about substitutability from direct experience with what the technology can do rather than from observation of realized labor market outcomes. Two findings support this reading. First, the relationship between usage intensity (more hours per week with AI) and fear is strong and holds within occupation groups and demographic cells. Second, the specific tasks for which workers deploy GenAI are systematically related to their fears. Those with above-median displacement fear are differentially likely to use AI for core occupational functions such as technology design, programming, mathematics, speaking, and equipment maintenance. Those with below-median fear report disproportionately using it for writing and learning strategies. Workers who use GenAI as a drafting and reasoning aid appear to experience it as a complement that accelerates their own output, whereas workers who watch it execute substantive technical work directly observe that the technology can perform a task central to their occupation.

Despite the fact such fears have not materialized, the fact that there is elevated anxiety about job loss is itself noteworthy. 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). Beliefs about displacement therefore carry real economic consequences whether or not they turn out to be accurate.

Our paper contributes to a growing literature on the labor market impacts of GenAI. 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), who document negative employment effects (particularly for younger workers). More broadly, our paper joins a literature on the real effects of subjective layoff uncertainty, (Stephens Jr, 2004; Hendren, 2017; Lehner et al., 2025).

The remainder of the paper proceeds as follows. Section 2 describes our survey design, documents AI adoption patterns, and characterizes workers' beliefs about displacement risk and the task-level correlates of those beliefs. Section 3 presents our empirical analysis of labor demand and job separations using Lightcast job postings data and the Current Population Survey. Section 4 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) GenAI 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 fielded in six waves across 25,146 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 GenAI use as well as other questions about education level, industry, income, and many other metrics.

Several 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 GenAI use, their central question (replicated in our survey) is fairly generic and begins with defining GenAI 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 GenAI 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 for more detailed questions about GenAI usage.[1]

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

Among those screened for AI use, we ask further questions about the intensive margin usage of the technology. For example, we ask "How many days each week do you use generative AI tools at work?" and "How many hours per week do you use generative AI tools at work?"

[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 GenAI adopters. From wave 6 (Mar. 2026) onwards, it was asked to all respondents. It was not asked in wave 1 (Dec. 2024).

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 (see Appendix Table A.2 for full list of options).

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 B 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) and the New York Fed Survey of Consumer Expectations (Hashim et al., 2026).

Internationally, adoption varies from Italy, Canada and Japan, which have less than 25% of their workforces using GenAI 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.2). 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.3). 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.4). Conditional on adoption, many users report frequent but intermittent use: roughly one-third of users report using GenAI every weekday (Appendix Figure A.7), 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

[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.13.

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

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

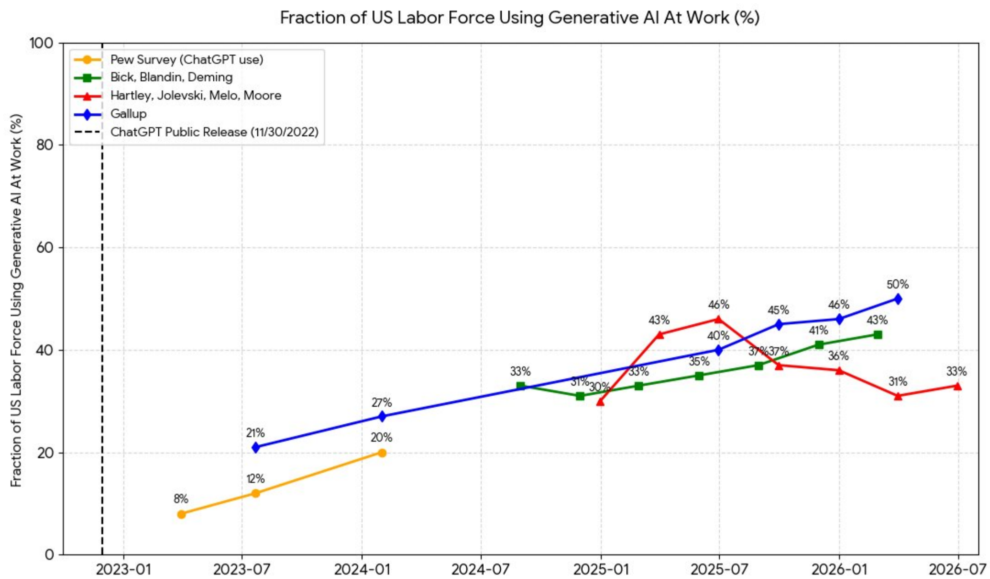

(b) Adoption at Work Across Countries in 2025

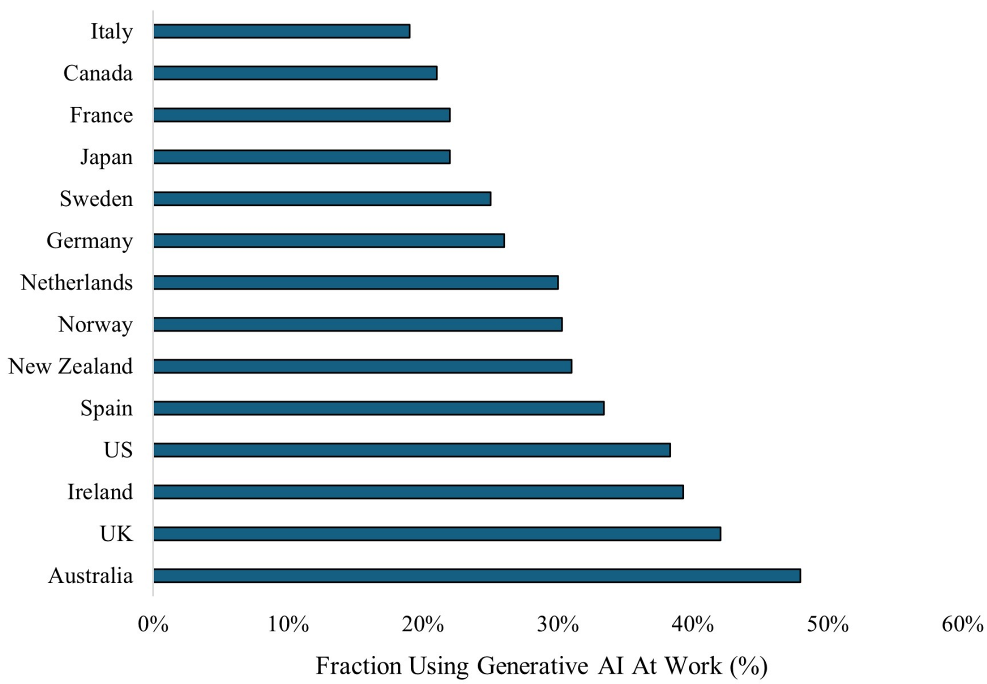

Note: Top panel shows adoption over time. First survey (orange) from Pew Charitable Trusts (2023-2024). Second survey (green) from Bick, Blandin, and Deming (2026) Third survey (red) from Hartley, Jolevski, Melo, and Moore (2026) Bottom panel shows international adoption in 2025 from Hartley, Jolevski,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

Figures A.2–A.11).

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.12). Specifically, we relate an occupation group's LLM exposure score from Eloundou et al. (2024) to that occupation's GenAI 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.

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, IBM, and Block have all cited AI as a reason for large layoffs since 2025. Firms report that they expect AI to reduce headcount (Yotzov et al., 2026) and 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). These concerns have even been recognized clinically: psychiatrists have identified a syndrome termed AI Replacement Dysfunction, marked by anxiety, depression, and existential fears about relevance, purpose, and future employability (McNamara and Thornton, 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 25% chance of AI-induced job loss, compared to 18% for those who do not use AI in their workplace. Figure 2a shows the distribution of these job loss probabilities for GenAI adopters at work (red) and non-adopters (blue). 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 realized 16% one-year job loss rate from the Survey of Income and Program Participation. 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.

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 our survey data on AI adoption and job loss perceptions and multiple AI exposure measures, we demonstrate the robustness of this relationship.

First, we show that workers who are more highly exposed to AI based on their occupation (independent of AI usage) are also more fearful of job loss. Those who are more highly exposed to AI based on the Eloundou et al. (2024) β score have a mean subjective job loss rate due to AI over the next two years of 21%, compared to just 16% for less-exposed workers.

Second, we show that workers who use 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 SOC code 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. The regression coefficient suggests that using AI for an extra ten hours per week increases a worker's self-reported likelihood of AI-induced job loss by 3.2 percentage points. 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 the more they use it.

Third, the specific tasks for which workers deploy genAI are systematically related to their displacement fears. Table 1 splits AI adopters at work into those with above- and below-median subjective AI displacement probability and compares the share of each group using AI for each task. Workers with above-median fear are differentially likely to use AI for technology design, speaking, equipment maintenance, mathematics, and programming (Panel A). Such tasks are more likely to be a core function of their job, hence higher job loss fears. By contrast, workers with below-median fear disproportionately use AI for writing, learning strategies, and critical thinking (Panel B). The writing gap is particularly illustrative, as nearly half of low-fear adopters use AI for writing, compared to roughly a third of high-fear adopters.

The task composition of AI use speaks to the inference workers draw about substitutability. Workers who use AI as a drafting and reasoning device are more likley to experience the technology as a complement that accelerates their own output, whereas workers who watch AI perform substantive technical tasks (writing code, designing systems) directly observe that the technology can execute a

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

(a) by GenAI Adoption

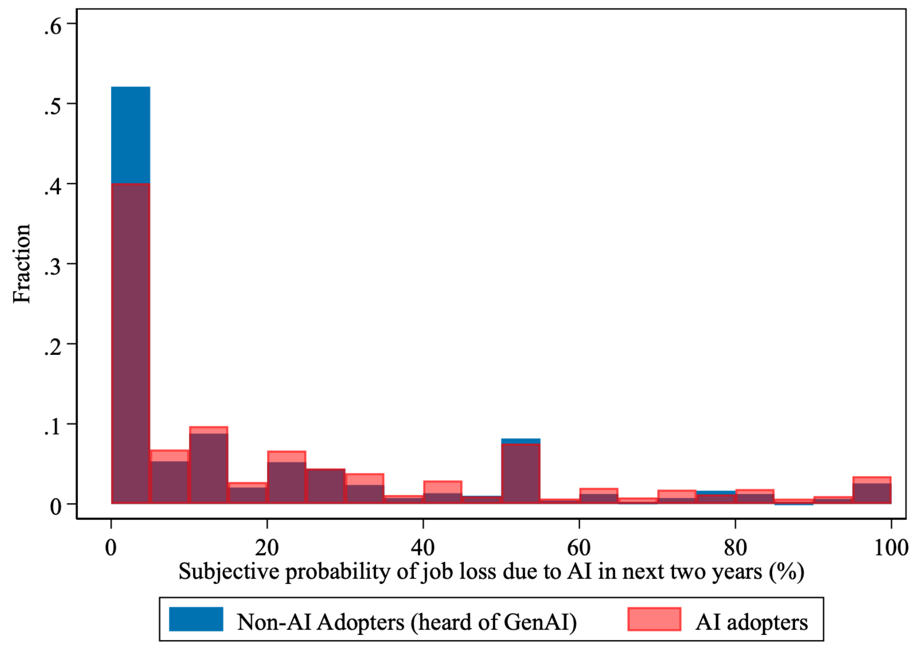

(b) by Predicted GenAI Exposure

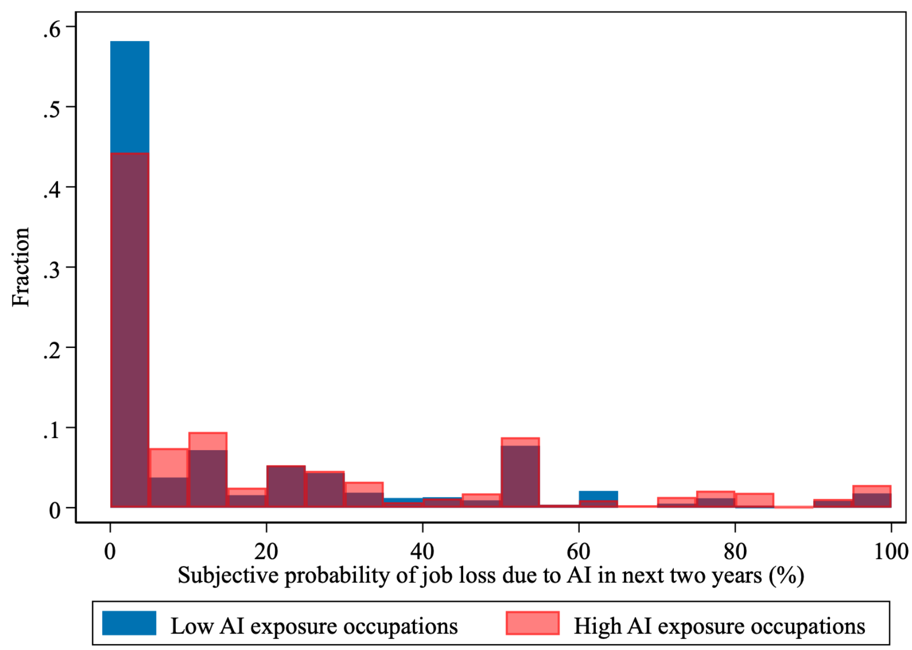

Source: Authors' survey. *Note*: Top panel shows histograms of subjective job loss perceptions due to generative AI in the next two years among non-AI adopters at work but who have heard of generative AI (blue) and AI users at work (red). Bottom panel is split by above- and below-median predicted AI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

task core to their occupation.

The Appendix provides further detail on which occupations are associated with tasks that we show have higher or lower job loss fears. Specifically, Appendix Table A.3 shows which occupations use AI most for writing (low-fear task) and Appendix Tables A.4 and A.5 do so for computer programming and speaking (high-fear tasks).

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

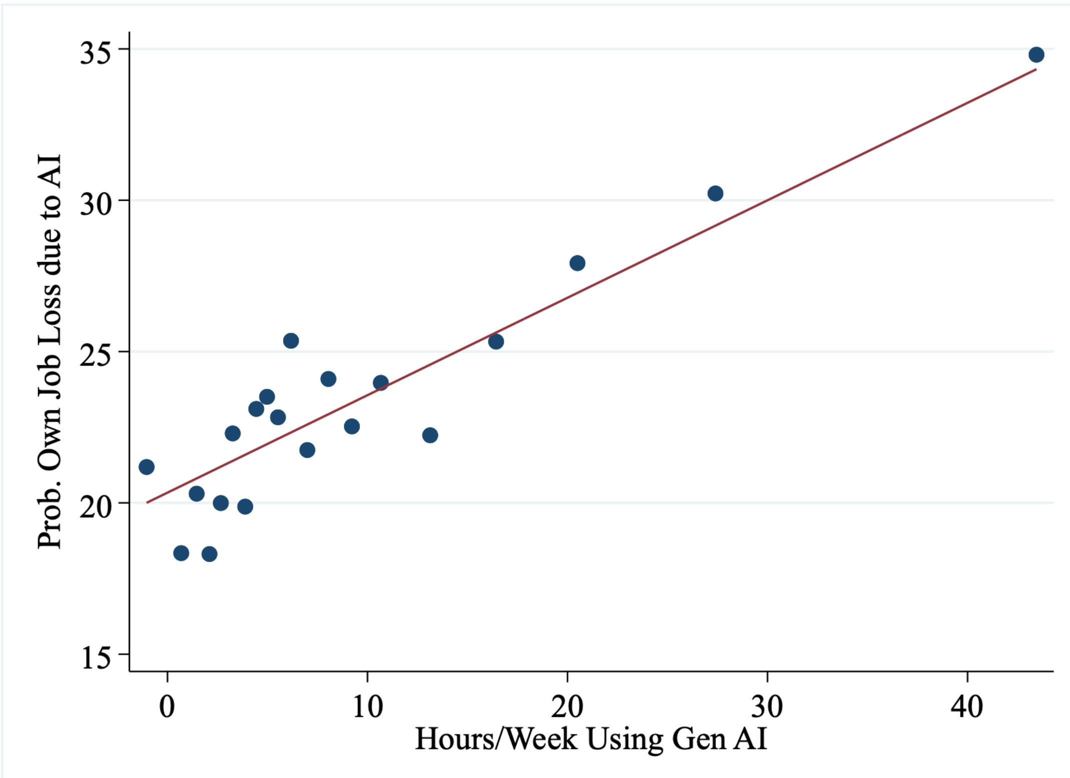

Note: Binscatter shows relationship between two-year subjective job loss rate due to AI with respect to the reported number of hours per week the respondent uses AI at work, among the sample that uses AI at work at all. Regression includes occupation group (2-digit SOC) fixed effects and controls for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

Table 1: GenAI Tasks at Work and AI Displacement Fears

Task	Share Using AI for Task Among		High-Fear
	High Job-Loss Fear	Low Job-Loss Fear	Difference
Panel A. Tasks Used Disproportionately by High-Fear Workers			
Technology Design	18.9%	12.9%	+6.0 pp
Speaking	13.4%	8.7%	+4.7 pp
Equipment Maintenance	9.6%	5.2%	+4.3 pp
Mathematics	14.3%	10.1%	+4.2 pp
Programming	16.9%	12.8%	+4.1 pp
Panel B. Tasks Used Disproportionately by Low-Fear Workers			
Critical Thinking	22.6%	23.9%	−1.3 pp
Learning Strategies	17.3%	21.3%	−4.0 pp
Writing	33.9%	48.5%	−14.6 pp

Notes: Sample restricted to workers who report using generative AI at work. "High (low) job-loss fear" indicates workers above (below) the median subjective probability of AI-induced job loss within two years. Columns report the share of each group selecting the listed task among the 35 task options in the survey (respondents may select multiple tasks). Panel A lists the tasks with the largest positive high-minus-low differences; Panel B lists those with the largest negative differences.

3 Effects of LLMs on the Labor Market

3.1 Data

Current Population Survey

We use the Current Population Survey (CPS) Basic Monthly files from January 2021 through December 2025 to study the effects of LLMs on employment and job-to-job flows. The CPS is a monthly household survey conducted by the U.S. Census Bureau for the Bureau of Labor Statistics (BLS), providing nationally representative data on employment status, occupation, and industry for the civilian population aged 16 and older.

Our main outcomes of interest are labor market transition rates, which we construct by linking respondents across consecutive monthly interviews using the CPS rotation group structure. Respondents are interviewed for four consecutive months, exit the sample for eight months, and return for another four months, allowing observation of month-to-month labor market transitions for matched workers. Because a labor market transition by definition occurs over two months, the first transitions we analyze occur between January and February 2021.

To study job loss, use the CPS to define two transition rates. 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 broad occupation level by aggregating individual-level transitions using CPS survey weights.

LightCast Job Openings

We analyze labor demand by using Lightcast's comprehensive job posting data. We aggregate job postings to year-month $\times$ 3-digit SOC (minor group) occupation cells, generating monthly counts of job postings for each occupation from February 2021 to February 2025. The 3-digit minor group is the most granular SOC level available in the Lightcast extract, and we harmonize all data sources in the difference-in-differences analysis to this level (94 occupation groups[5]).

AI Exposure and Predicted Adoption

Our preferred measure of occupational AI exposure comes from the β-score from Eloundou et al. (2024), which measures the capacity of large language models to reduce the time required to complete a given task by at least 50% while preserving quality. These scores are constructed from O*NET task descriptions and mapped to SOC occupations, providing a theoretically-grounded, pre-ChatGPT measure of occupational susceptibility to LLM automation that is plausibly exogenous to contemporaneous labor market conditions. We aggregate the occupation-level β-scores to the 3-digit SOC level using 2022 Occupational Employment Statistics employment weights.

To interpret our regression coefficients in units of AI adoption rather than a more difficult-to-interpret AI exposure measure, we construct a measure of predicted AI adoption. Specifically, we regress our survey-based adoption rate on the aggregated β-score across 3-digit occupations, restricting to occupations with at least 20 survey respondents so that adoption rates are measured with adequate precision. Unsurprisingly, theoretical AI exposure strongly predicts AI adoption in our survey (coefficient of 0.338, t-statistic of 4.51). Our predicted adoption measure is the fitted value from this regression. Because predicted adoption is a linear transformation of the β-score, a regression of a labor market outcome on theoretical exposure and a regression of that outcome on predicted adoption exploit the same underlying exogenous variation, and therefore deliver identical t-statistics and statistical significance. Only the units of the coefficient differ: with theoretical exposure, coefficients are expressed in units of the Eloundou et al. (2024) exposure measure, whereas with predicted adoption, coefficients are expressed in units of the survey adoption rate.

We do not regress labor market outcomes directly on our survey-based adoption measure because adoption is likely endogenous to those outcomes. For example, GenAI adoption may rise precisely in occupations facing negative demand shocks, as firms push workers to adopt AI in order to produce more with less. Our predicted adoption measure obviates this concern by isolating only the variation in adoption attributable to theoretical exposure.

As a robustness check, we repeat our analysis using Anthropic's occupational AI exposure score. We prefer the Eloundou et al. (2024) measure as our baseline, however, because it is a theoretical measure of exposure, whereas Anthropic's score is constructed largely from observed usage. Predict-

[5] In baseline regressions, we drop occupations that are not recorded at least 20 times in our six survey waves. Hence our preferred specification, which includes occupation fixed effects, only accounts for 84 minor occupation groups

ing one adoption measure with another provides less credible exogenous variation than predicting adoption with a theoretical measure of susceptibility to GenAI.

3.2 Empirical Methodology

Our identification strategy tests how labor market outcomes respond to GenAI diffusion across occupations, using the public release of ChatGPT on November 30, 2022 as our treatment date. Our primary specification is a differences-in-differences with continuous treatment approach in the spirit of Callaway et al. (2024) that considers all occupations as receiving different doses of treatment. We study the response of labor market outcomes to occupational AI exposure in the months and quarters surrounding the ChatGPT rollout:

$$y_{ot} = \alpha_o + \delta_t + \beta\,(D_o \times \mathbf{1}[t \geq \text{Dec 2022}]) + \varepsilon_{ot}, \qquad D_o \in \{\, \text{Exposure}_o,\ \widehat{\text{Adoption}}_o \,\} \tag{1}$$

where y_{ot} is a labor market outcome for occupation o in period t, α_o are occupation fixed effects, and δ_t are time fixed effects, and $t^* =$ November 2022 is the omitted reference period. $\widehat{\beta}$ is an average causal response measuring the change in y_{ot} associated with a one-unit increase in dose D_o after ChatGPT's release, relative to the pre-period. On the right-hand side, our baseline variable is occupational exposure as defined by Eloundou et al. (2024). This exposure measure, defined before LLMs were widely adopted, should be orthogonal to subsequent labor market dynamics within an occupation, such as the path of vacancy postings and layoffs. Identification in the continuous-treatment setting requires a strong parallel trends assumption: absent ChatGPT's release, occupations at every level of exposure would have followed the same path of outcomes, and occupations did not select into higher exposure on the basis of their anticipated treatment effects.

For interpretability, we also regress our outcomes of interest on $\widehat{\text{Adoption}}_o$, which is occupation o's of AI adoption that would be predicted from a linear regression of adoption on AI exposure from Eloundou et al. (2024). This allows us to express our coefficents in terms of effects due to predicted adoption, which is more straightforward than percentage points of an abstract exposure measure. We use predicted rather than actual adoption to address a potential endogeneity concern: GenAI adoption may rise precisely in occupations facing negative demand shocks, as firms push workers to adopt AI in order to produce more with less. Estimates using actual adoption would therefore conflate the causal effect of AI with correlated demand shocks. Regressions are run at the 3-digit SOC level (omitting occupations that appeared less than 20 times across the six waves of our survey). Standard errors are clustered at the occupation level throughout.

We assess robustness along three dimensions. First, because our sample restriction is somewhat arbitrary, we re-estimate Equation (1) dropping occupations with fewer than 10 and fewer than 30 survey respondents rather than 20 (Appendix Table A.6). Second, we substitute the Massenkoff and McCrory (2026) (Anthropic) occupational exposure score for the Eloundou et al. (2024) β-score (Appendix Table A.7); we prefer the latter as our baseline because it is derived from task content rather than from realized AI usage, which provides more credible exogenous variation. Third, we estimate a binary-treatment version of the design comparing occupations in the top quartile of exposure to those in the bottom quartile (Appendix Table A.8).

3.3 Results

Table 2 presents our headline continuous-treatment difference-in-differences estimates from equation (1) for our three labor market outcomes: log job postings (Lightcast), the employment-to-unemployment (EU) transition rate, and the employment-to-nonparticipation (EN) transition rate (CPS). Panel A uses the Eloundou et al. (2024) exposure score directly as the dose; Panel B uses predicted adoption.

Across all six specifications, no coefficient is statistically distinguishable from zero. However, the pattern of point estimates is consistent in that vacancy effects are negative and separation effects are positive. Taking the Panel B estimates at face value, a ten-percentage-point increase in predicted AI adoption is associated with a 2.7% relative decline in job postings (s.e. 2.7) and a 0.14 percentage-point relative increase in the monthly EU rate (s.e. 0.12) which stands around 1.1%. The corresponding EN estimate is even smaller. Confidence intervals allow us to rule out relative posting declines larger than roughly 8% and EU increases larger than roughly 0.37 percentage points, per ten percentage points of predicted adoption. In the two and a half years following the release of ChatGPT, in other words, we can rule out anything close to the substantial displacement that motivated popular concerns over AI documented in section 2.3.

Vacancies

We find no evidence that GenAI exposure has reduced job openings. The continuous-treatment estimates in the first column of Table 2 imply that a ten-percentage-point increase in predicted AI adoption is associated with a 2.7% relative decline in postings, statistically indistinguishable from zero, with a confidence interval that rules out relative declines larger than roughly 8%. This complements similar findings also using Lightcast by Audoly et al. (2026).

This null is robust to our sample restrictions and choice of specification. Varying the survey-response threshold for including an occupation from 20 to either 10 or 30 leaves the posting coefficient negative and insignificant, and if anything smaller in magnitude (Appendix Table A.6). A binary-treatment design comparing occupations in the top quartile of exposure to those in the bottom quartile likewise yields a small negative and insignificant estimate (Appendix Table A.8).

The one specification that departs from this pattern uses the Massenkoff and McCrory (2026) (Anthropic) exposure score in place of the Eloundou et al. (2024) β-score. In that case, the vacancy posting coefficient is negative and statistically significant at the 1% level (Appendix Table A.7), implying a large 11% relative decline in postings per ten percentage points of a relative increase in predicted adoption. We are reluctant to read this evidence along as a sign of negative labor market because, as discussed in Section 3.2, the Anthropic score is constructed from observed patterns of AI usage

Table 2: Continuous Treatment Difference-in-Differences: AI Exposure and Labor Market Outcomes

	Vacancies	Employment Transitions	
	log(postings)	Emp→Unemp (*EU*)	Emp→NILF (*EN*)
Panel A. Exposure (Eloundou et al., 2024)			
$\text{Exposure}_o \times Post_t$	−0.090	0.465	0.113
	(0.091)	(0.404)	(0.243)
Panel B. Adoption Predicted by Occupational Exposure			
$\widehat{\text{Adoption}}_o \times Post_t$	−0.266	1.374	0.334
	(0.268)	(1.194)	(0.717)
Occupation & Month FE	✓	✓	✓
Occupation Cells	85	85	85
Observations	4,165	4,930	4,930

Notes: Each column-panel cell reports a separate regression of equation (1) at the 3-digit SOC × month level. Standard errors clustered at the 3-digit SOC level in parentheses. $Post_t$ indicates months from December 2022 onward. *EU* and *EN* transition rates are expressed in monthly percentage points. The sample is restricted to 3-digit SOCs occupations with at least 20 survey respondents. Panel A uses the employment-weighted Eloundou et al. (2024) β-score aggregated to 3-digit SOC; Panel B uses predicted adoption, the fitted value from a cross-sectional regression of the survey adoption rate on the Panel A exposure measure. Sample periods: February 2021–February 2025 for Lightcast postings; February 2021–November 2025 for CPS transitions.

rather than from task content, so it is precisely the measure most vulnerable to the endogeneity concern that motivates our predicted-adoption design: AI adoption may be most prevalent where firms are already under pressure to cut back on labor costs.

Layoffs

We find no evidence that workers in AI-exposed occupations are being laid off at higher rates, as measured either by transitions into unemployment or non-employment. Figure 4 presents event-study estimates for the *EU* and *EN* transition rates using a binary treatment indicator for occupations that are above-median employment-weighted predicted AI adoption.[6] Across both panels, coefficients are small and statistically indistinguishable from zero throughout the sample period. We find no evidence of differential pre-trends prior to the ChatGPT release, supporting the parallel trends assumption underlying our design.

The continuous-treatment estimates in Table 2 further illustrate these findings. The estimated effect on the monthly *EU* rate is 0.47 percentage points per unit of the Eloundou et al. (2024) β-score (s.e. 0.40), and the *EN* estimate is even smaller. Both are statistically indistinguishable from zero. These findings are consistent with Hyman et al. (2026) and other work that has found limited effect of AI exposure on layoffs in the United States at the aggregate level.

[6] We estimate *EU* and *EN* rates in the event-study at a quarterly frequency for clearer exposition.

The EU estimate remains positive and insignificant across every robustness exercise: at both alternative survey-response thresholds, under the binary top-versus-bottom-quartile design, and using the Massenkoff and McCrory (2026) exposure measure. The EN estimate is less stable in sign but never approaches significance in any specification.

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 gap between perceived and realized displacement risk is itself a first-order fact about the current labor market, as the anxiety documented in Section 2.1 is running well ahead of anything visible in separations data.

4 Conclusion

This paper documents a sharp divergence between what U.S. workers believe about GenAI-induced displacement and what has actually occurred in the labor market through 2025. Using a large-scale, multi-wave survey of over 25,000 U.S. workers fielded between December 2024 and March 2026, we document three main findings.

First, GenAI adoption at work rose sharply in early 2025, reaching approximately 30% of the U.S. labor force by the end of the first quarter of 2025 before plateauing in the mid-to-high 30% range 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, American workers report substantial subjective job loss risk attributable to GenAI. 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 slightly exceeds the baseline two-year job loss rate from all causes combined as measured by the SIPP. Workers in more AI-exposed occupations and those who use AI more intensively on the job report the highest displacement fears. Further, workers who deploy GenAI for core technical functions such as programming and technology design report markedly higher fears than those who use it for writing and learning strategies. Together these patterns suggest that direct experience with the technology's capabilities amplifies rather than alleviates concern.

Third, we find no evidence that these fears have yet been realized. Job postings in more AI-exposed occupations show no statistically significant decline following the public release of ChatGPT, and layoffs (transitions from employment into unemployment and out of the labor force) show no significant response either. These nulls hold across alternative sample restrictions, a binary top-versus-bottom-quartile design, and an alternative occupational exposure measure.

Taken together, these results suggest that the near-term labor market footprint of GenAI has been felt more in what workers expect than in what they have experienced. This gap matters for how pol-

Figure 4: Effect of Predicted AI Exposure on Job Separations

(a) Employment-to-Unemployment (EU) Transition Rate

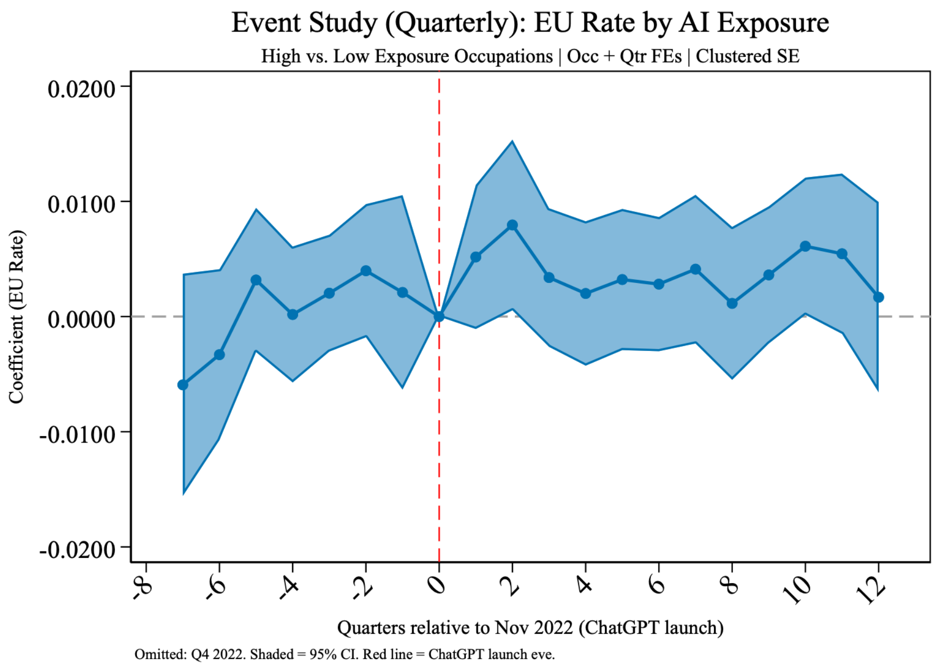

(b) Employment-to-Nonparticipation (EN) Transition Rate

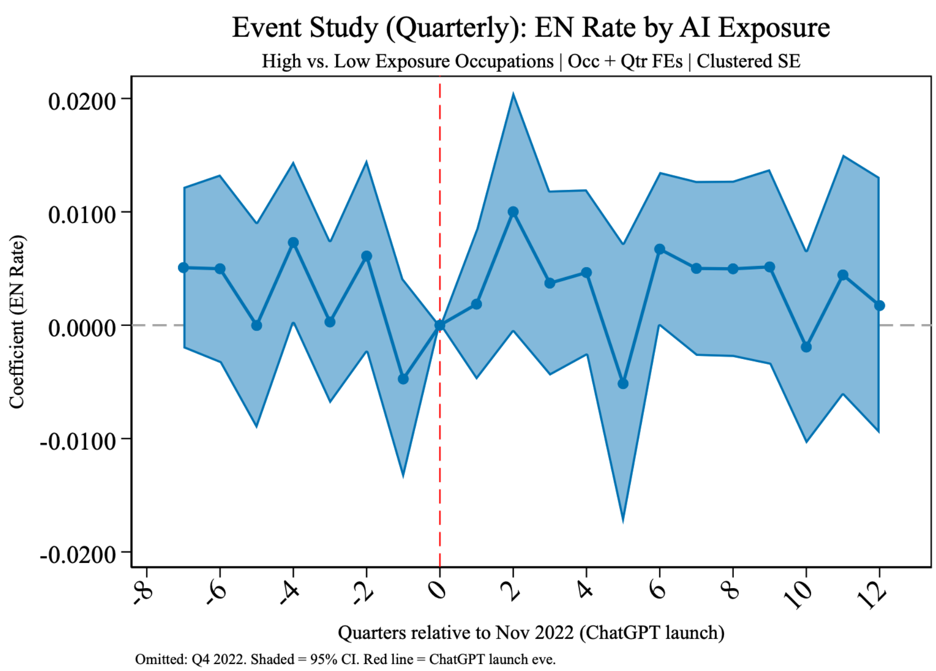

Note: Dependent variable on top panel is EU transition rate and on bottom panel is EN (NILF) transition rate. Control group are occupations with below median level of employment-weighted predicted AI adoption (0.33). Treatment group includes occupations with above median level of employment-weighted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

icymakers and firms interpret aggregate labor market data in an era of rapid AI diffusion, since anticipated displacement can shape household behavior well before any displacement occurs. Whether the elevated fears we document reflect workers correctly forecasting displacement that has not yet arrived, or instead a systematic misreading of personal experience with a technology that is more likely to reshape jobs than to replace them, is a question we cannot resolve with the current data.

References

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

AUDOLY, R., M. GUERIN, AND G. TOPA (2026): "Do Job Postings Show Early Labor-Market Effects of AI?" *Liberty Street Economics*, liberty Street Economics Blog.

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.

CALLAWAY, B., A. GOODMAN-BACON, AND P. H. C. SANT'ANNA (2024): "Difference-in-Differences with a Continuous Treatment," Working Paper 32117, National Bureau of Economic Research.

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

FELTEN, E., M. RAJ, AND R. SEAMANS (2023): "How will language modelers like ChatGPT affect occupations and industries?" *arXiv preprint arXiv:2303.01157*.

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

HASHIM, A., G. KOSAR, AND W. VAN DER KLAAUW (2026): "Use of Gen AI in the Workplace and the Value of Access to Training," *Liberty Street Economics*, liberty Street Economics Blog.

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.

HYMAN, B., T. VON WACHTER, D. HERRERA, R. MOGHADAM, J. MORRIS, S. MOTGHARE, AND K. SEGAL (2026): "Tracking AI-Related Job Loss Using Unemployment Insurance Claims Data in California," Report, California Policy Lab, University of California.

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*.

MASSENKOFF, M. AND P. MCCRORY (2026): "Labor market impacts of AI: A new measure and early evidence," .

MCNAMARA, S. N. AND J. E. THORNTON (2025): "Artificial Intelligence Replacement Dysfunction (AIRD): A Call to Action for Mental Health Professionals in an Era of Workforce Displacement," *Cureus*, 17.

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*.

YOTZOV, I., J. M. BARRERO, N. BLOOM, P. BUNN, S. J. DAVIS, K. M. FOSTER, A. JALCA, B. H. MEYER, P. MIZEN, M. A. NAVARRETE, ET AL. (2026): "Firm data on AI," Tech. rep., National Bureau of Economic Research.

Online Appendix Materials

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 GenAI (Figure 4). Nearly 50% of those in the sample with a graduate degree use GenAI. 37% of those with a college degree as their highest level education attained use GenAI. Roughly 20% of those who are only high school graduates or have some college use AI.

With respect to industries, there is a varied picture (Figure 5). Industries most likely to use AI at work include "Information Services" and "Management of Companies" (using GenAI 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.

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 GenAI at work. Nearly 50% of workers making above $200,000 annually use AI at work.

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

ChatGPT remains the most popular GenAI 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 GenAI at work at some point since tools became public.

However, the intensive margin of using GenAI tools at work also tells a different story. Conditional on using AI at work, about 33% of workers use the technology five days per week at work (every weekday). Roughly 12% of adopters users use such tools at work only one day at work. About 17% and 18% of adopters use the tools at work two and three days per week respectively. The number of hours at work each week using GenAI 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 GenAI tools every weekday (or several weekdays), the total hours at work each week using the tools is reported to be in most cases less than 15 hours, suggesting there is intermittent use of AI throughout the workday. On the whole, these results combined suggest that GenAI 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: Average Number of Hours of Generative AI Use At Work Over Time

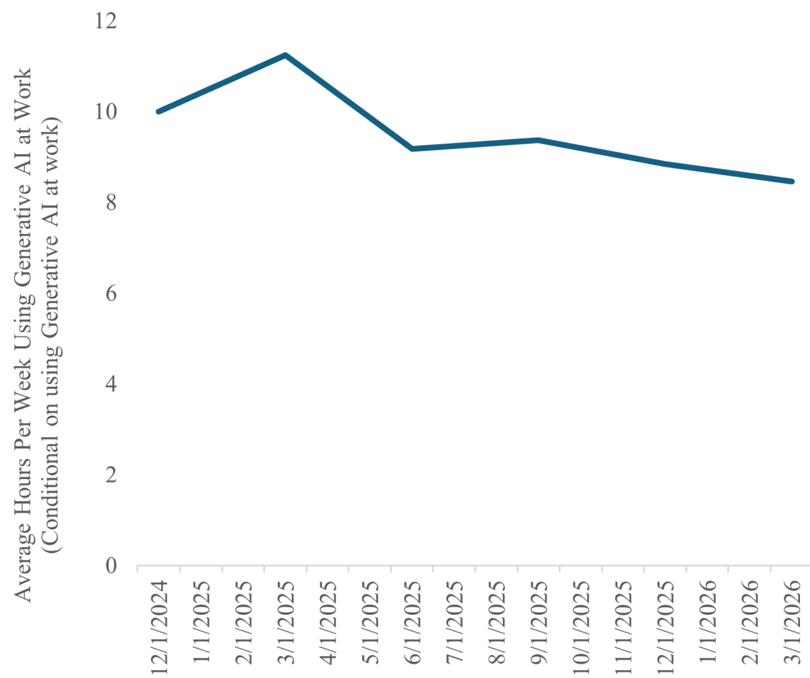

Source: Authors' survey.

Figure A.2: Generative AI Use By Level of Education

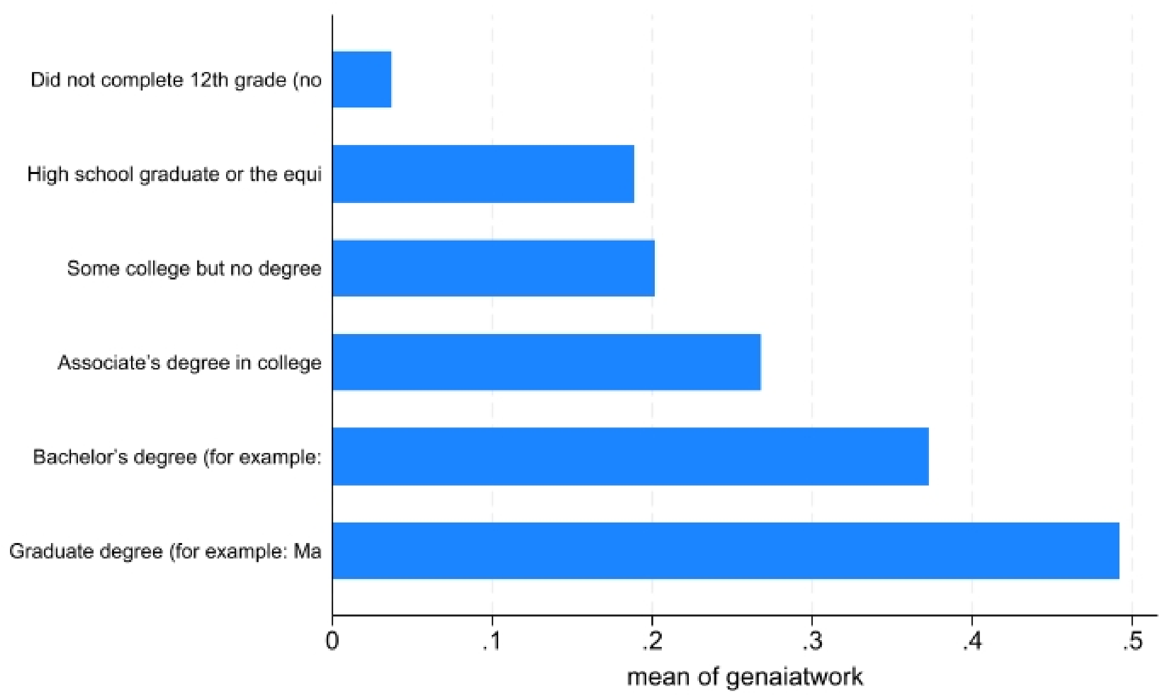

Source: Authors' survey.

Figure A.3: Generative AI Use By Income

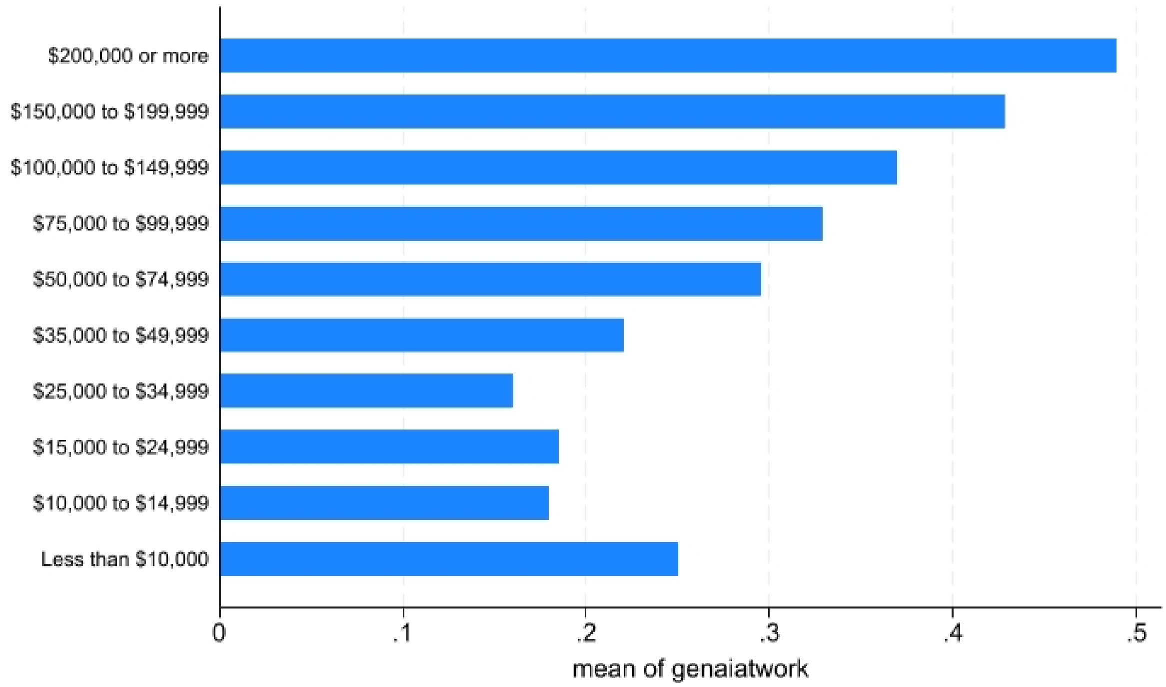

Source: Authors' survey.

Figure A.4: Generative AI Use By Industry

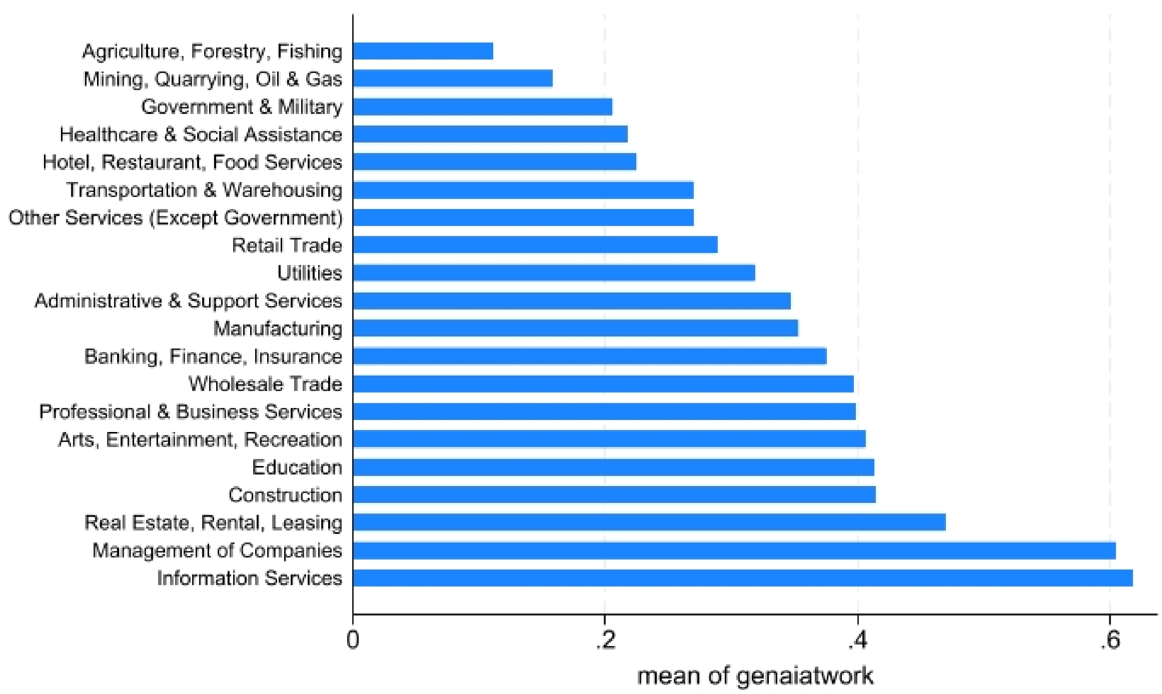

Source: Authors' survey.

Figure A.5: Generative AI Use By Gender

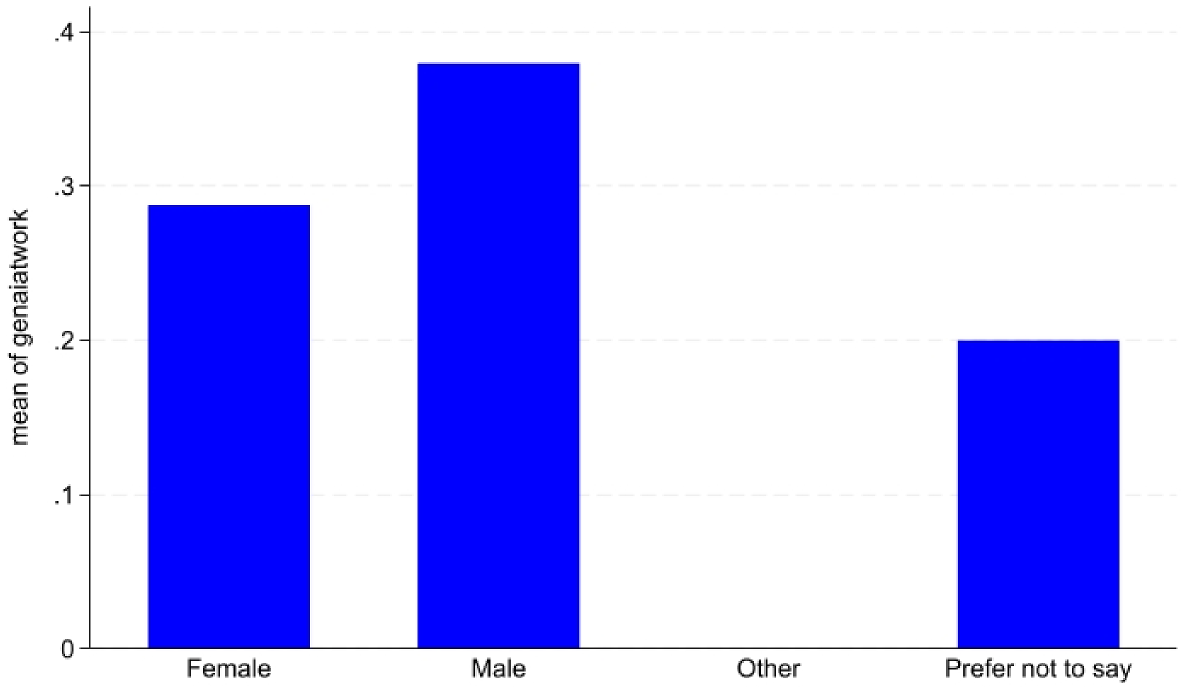

Source: Authors' survey.

Figure A.6: Generative AI Use By Race

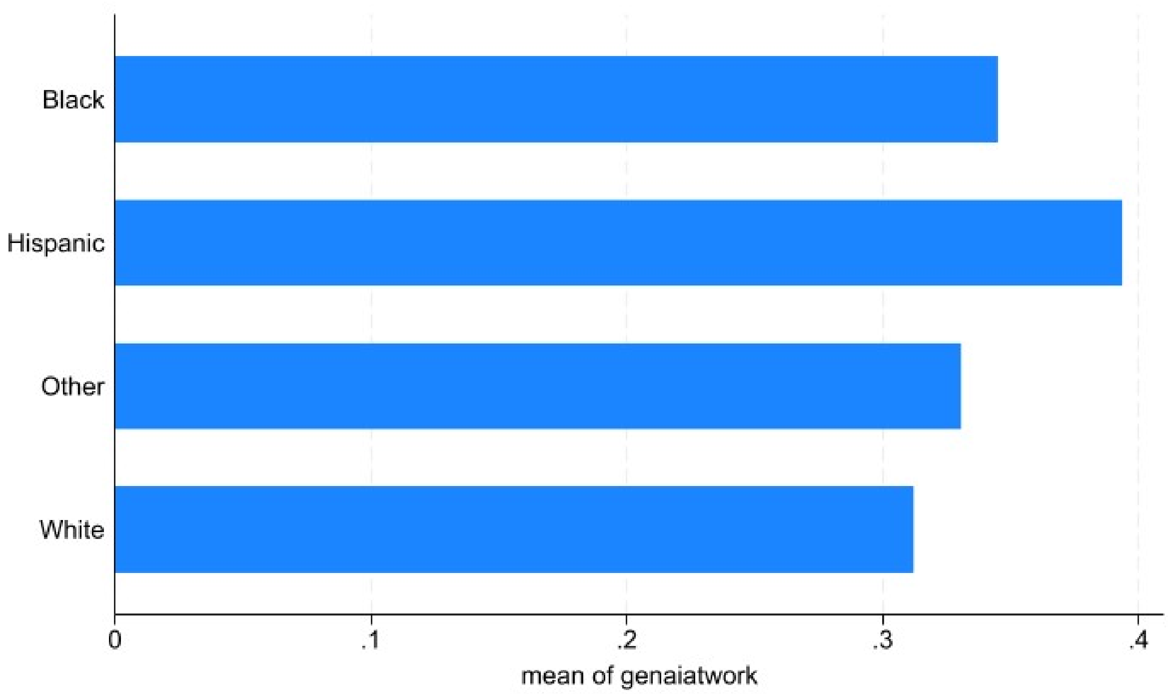

Source: Authors' survey.

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

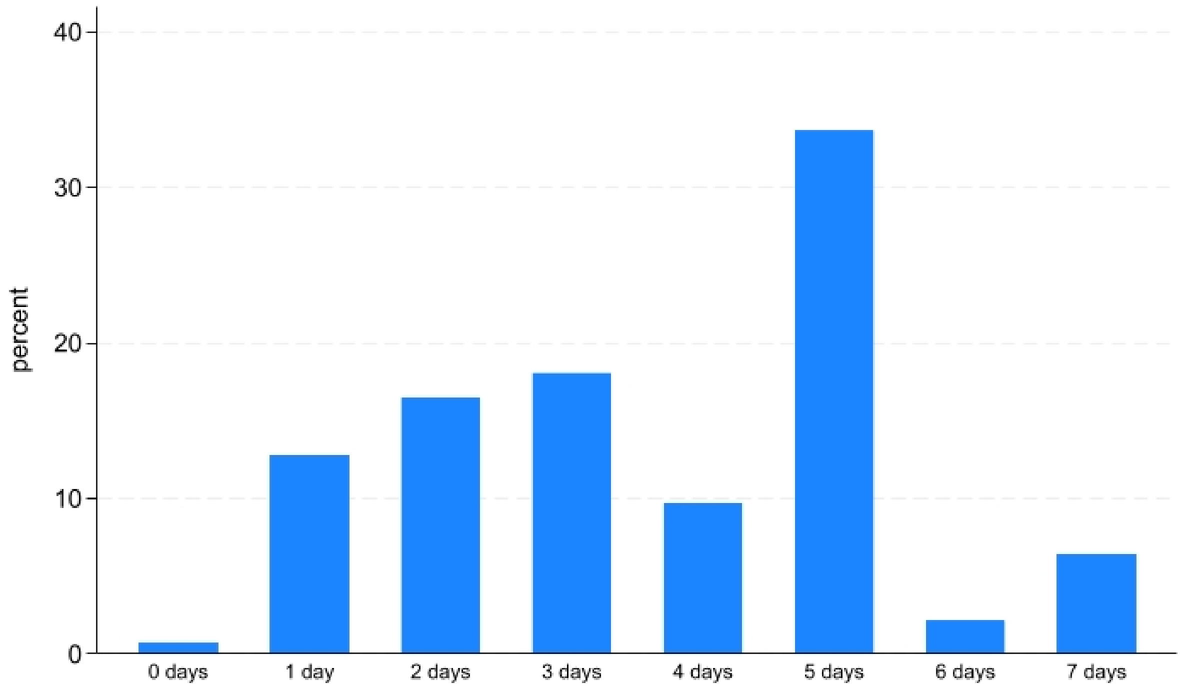

Source: Authors' survey.

Figure A.8: Number of Hours At Work Each Week Using Generative AI Tools

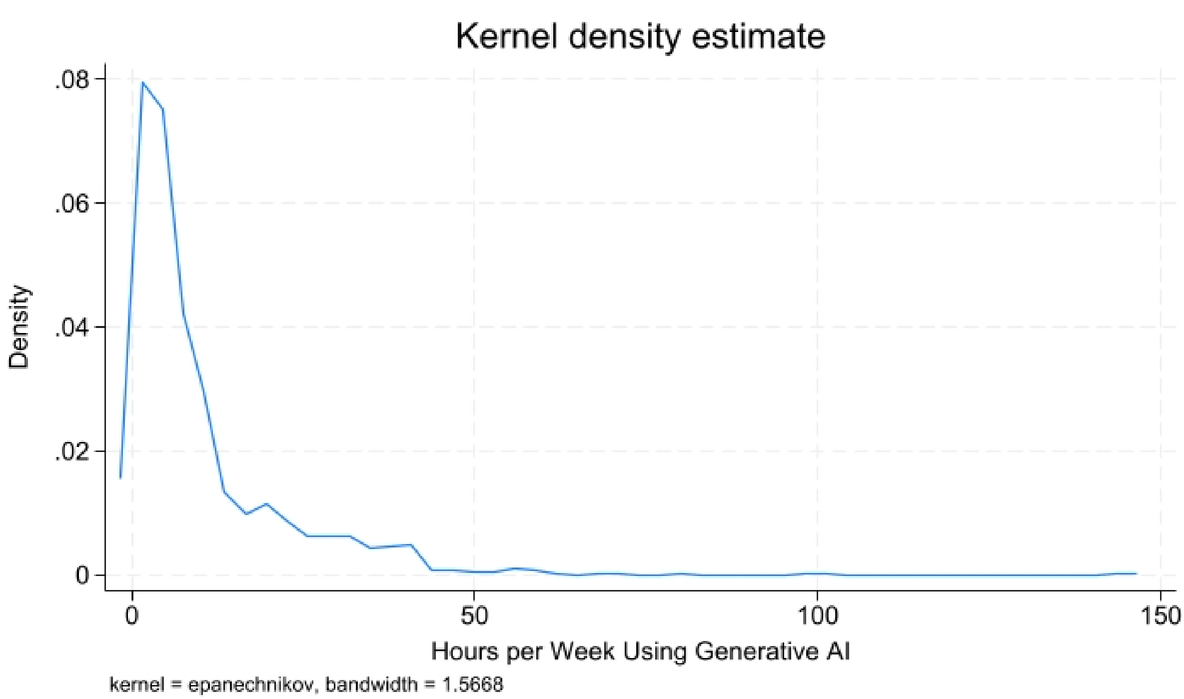

Source: Authors' survey.

Figure A.9: Average number of minutes to complete a task with and without Generative AI

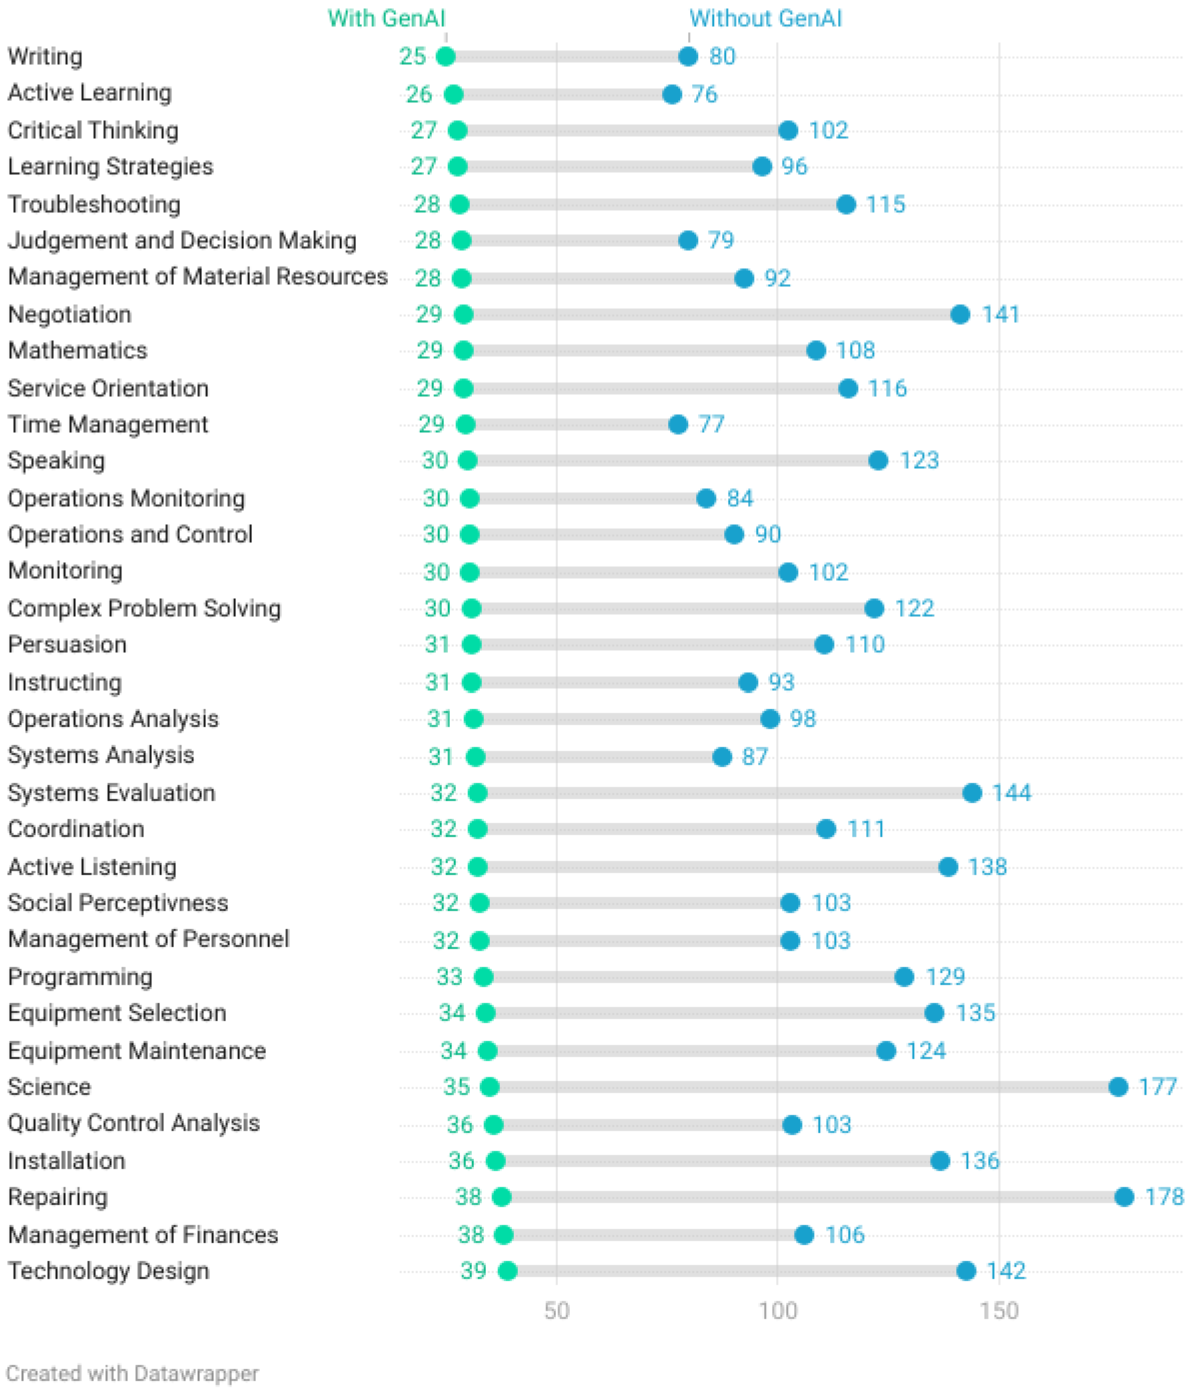

Source: Authors' survey.

Figure A.10: Productivity increase in completing various tasks using Generative AI

Source: Authors' survey.

Figure A.11: Uses of Generative AI to complete tasks

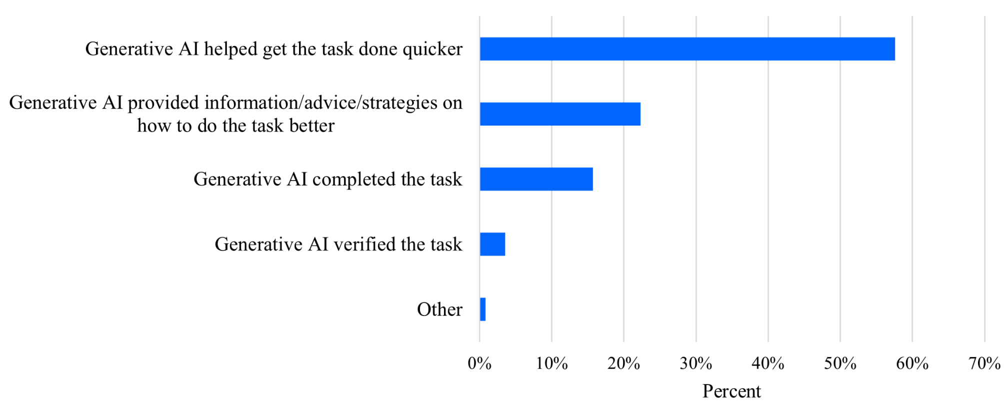

Source: Authors' survey.

Figure A.12: GenAI Adoption at Work vs. Predicted Exposure by Occupation

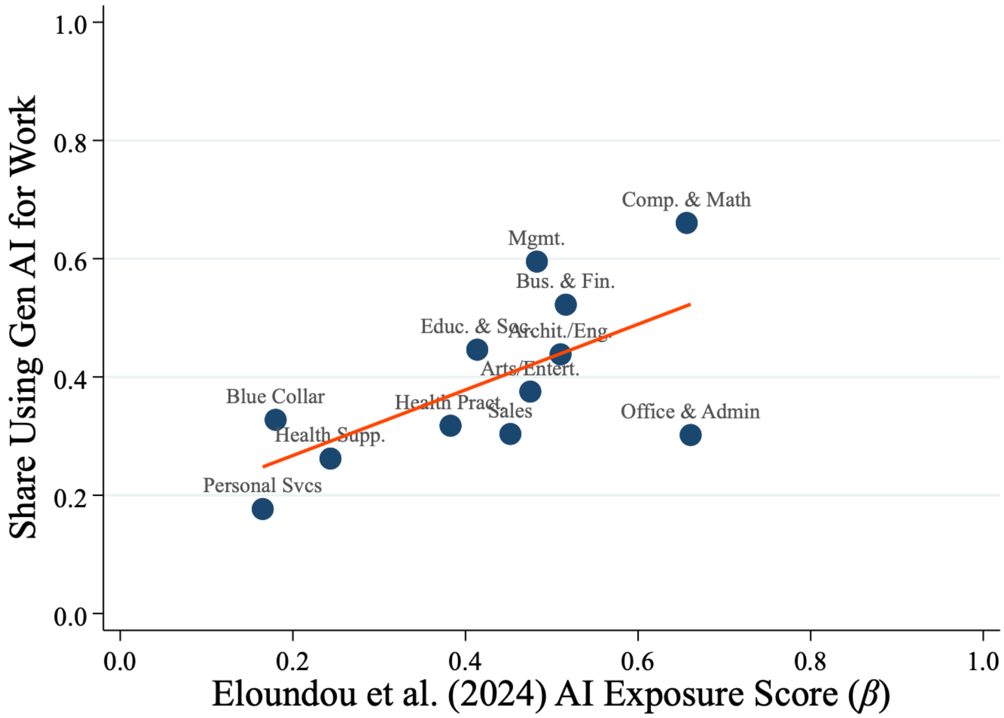

Source: Authors' survey & Eloundou, Manning, Mishkin, and Rock (2024) LLM β-exposure measure. Coefficient: 0.55; Intercept: 0.16.

Figure A.13: ChatGPT and Gemini Google Search Volume in the United States

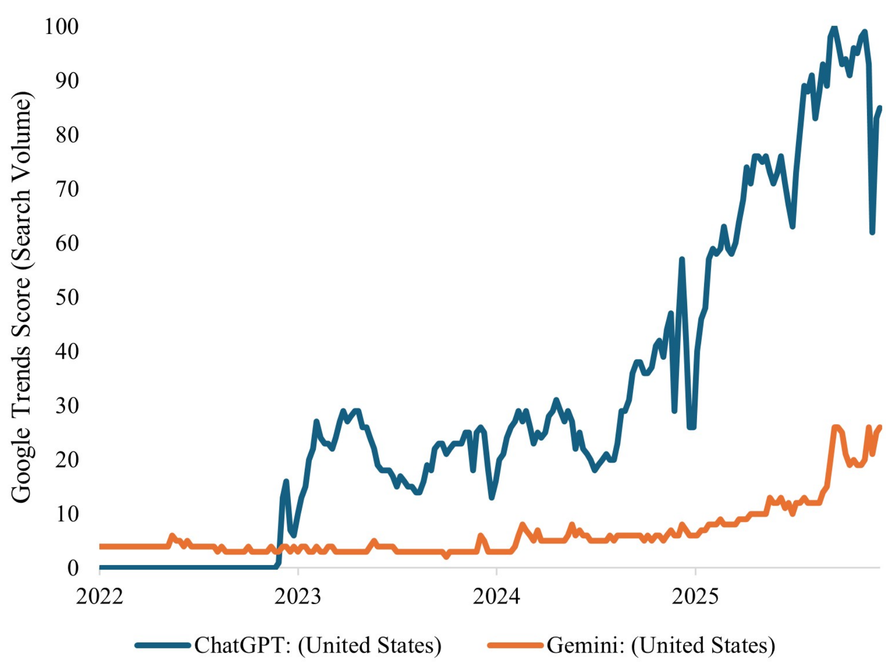

Source: Google Trends.

Figure A.14: Evolution of Key Labor Market Indicators

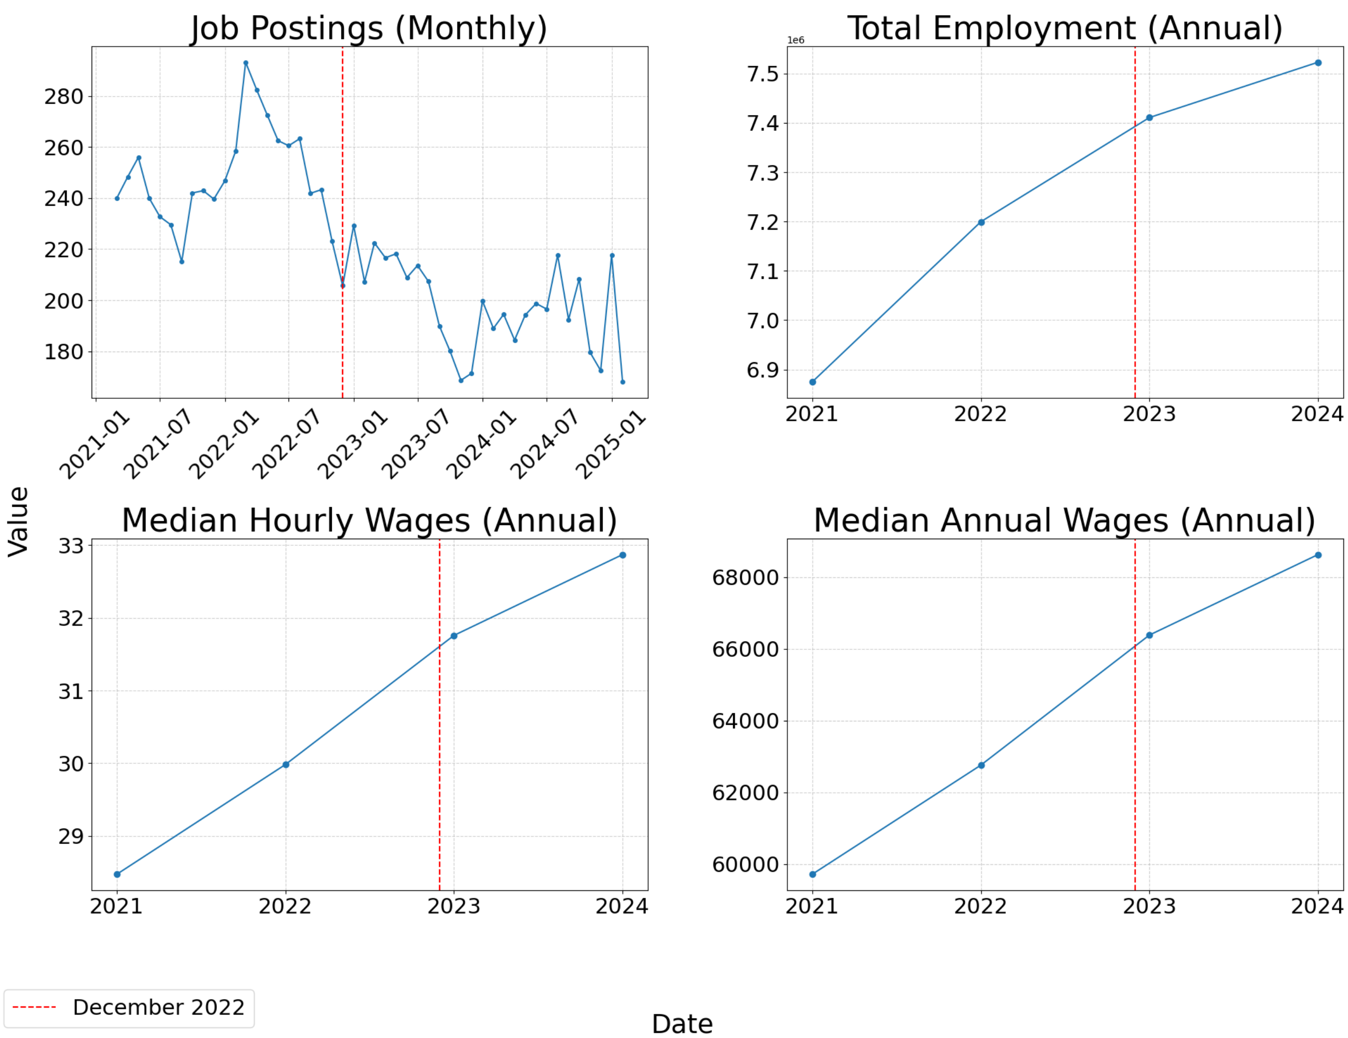

Notes: The panel displays trends for four key labor market indicators. The top-left plot shows the monthly evolution of job postings (Source: Lightcast). The remaining plots show the annual evolution of total employment, median hourly wages, and median annual salary (Source: OES). All data represent 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

Table A.2: Task Options for Survey Question on Generative AI Use at Work

1.	Active Listening	19.	Equipment Selection
2.	Writing	20.	Installation
3.	Speaking	21.	Programming
4.	Mathematics	22.	Operations Monitoring
5.	Science	23.	Operation and Control
6.	Critical Thinking	24.	Equipment Maintenance
7.	Active Learning	25.	Troubleshooting
8.	Learning Strategies	26.	Repairing
9.	Monitoring	27.	Quality Control Analysis
10.	Social Perceptiveness	28.	Judgment and Decision Making
11.	Coordination	29.	Systems Analysis
12.	Persuasion	30.	Systems Evaluation
13.	Negotiation	31.	Time Management
14.	Instructing	32.	Management of Financial Resources
15.	Service Orientation	33.	Management of Material Resources
16.	Complex Problem Solving	34.	Management of Personnel Resources
17.	Operations Analysis	35.	Other (please specify)
18.	Technology Design		

Notes: This table lists the response options for the survey question "In what tasks are you using generative AI/large language models at work?" Options are listed in the order in which they were shown to respondents. The question was presented as a checkbox, so respondents could select multiple options. Respondents selecting "Other" were asked to specify the task in a free-text field.

Table A.3: Occupations Most Frequently Using AI for: Writing

Occupation	GenAI Adoption	Perceived Job Loss Risk Among Adopters	*N*
Education and Childcare Administrator	78.6%	27.6%	14
Secretary or Administrative Assistant	76.9%	27.7%	26
Librarian and Media Collection	72.7%	5.0%	11
Marketing Agent	71.4%	38.4%	14
Health Teacher, Postsecondary	70.6%	4.1%	17
Real Estate Broker & Sales Agent	68.4%	17.6%	38
Lawyer and Judicial Law Clerk	66.7%	10.9%	69

Note: List excludes occupations described as a "miscellaneous" type and occupations that appear less than ten times in the survey over six waves. Rightmost column *N* describes how often an occupation appears in the survey.

Table A.4: Occupations Most Frequently Using AI for: Programming

Occupation	GenAI Adoption	Perceived Job Loss Risk Among Adopters	*N*
Software & Web Developer or Programmer	55.7%	26.4%	228
Data Scientist	53.3%	14.0%	15
Database & Network Administrator/Architect	45.5%	12.3%	22
Electrician	40.0%	4.5%	10
Radio/Telecommunication Equipment Installer	40.0%	30.8%	10
Research Scientist	36.0%	37.1%	25
Computer, ATM, & Office Machine Repairers	35.0%	29.5%	20

Note: List excludes occupations described as a "miscellaneous" type and occupations that appear less than ten times in the survey over six waves.

Table A.5: Occupations Most Frequently Using AI for: Speaking

Occupation	GenAI Adoption	Perceived Job Loss Risk Among Adopters	*N*
Electrician	40.0%	4.5%	10
Dentist	33.3%	16.7%	12
Insurance Sales Agent	33.3%	23.7%	15
Social Sciences Teacher, Postsecondary	28.6%	26.7%	21
Home Health and Personal Care Aide	27.3%	27.4%	33
Clinical Laboratory Technologist and Technician	25.0%	19.4%	12
Cashier	21.7%	33.7%	23

Note: List excludes occupations described as a "miscellaneous" type and occupations that appear less than ten times in the survey over six waves.

Table A.6: Continuous Treatment DID: AI Exposure and Labor Market Outcomes with Different Thresholds for Dropping Occupations with Small Survey Response Rates

	Vacancies		Employment Transitions			
	log(postings)		*EU*		*EN*	
Panel A. Exposure (Eloundou et al., 2024)						
$\text{Exposure}_o \times Post_t$	-0.047	-0.073	0.312	0.477	-0.112	0.170
	(0.087)	(0.090)	(0.429)	(0.406)	(0.307)	(0.246)
Panel B. Adoption Predicted by Occupational Exposure						
$\widehat{\text{Adoption}}_o \times Post_t$	-0.139	-0.206	0.920	1.341	-0.331	0.477
	(0.256)	(0.253)	(1.266)	(1.142)	(0.908)	(0.693)
Occupation & Month FE	✓	✓	✓	✓	✓	✓
Drop Occupation if Fewer than N	10	30	10	30	10	30
Occupation Cells	90	80	90	80	90	80
Observations	4,403	3,919	5,219	3,919	5,219	3,919

Notes: Each column-panel cell reports a separate regression of equation (1) at the 3-digit SOC $\times$ month level. Standard errors clustered at the 3-digit SOC level in parentheses. $Post_t$ indicates months from December 2022 onward. *EU* and *EN* transition rates are expressed in monthly percentage points. The sample is restricted to 3-digit SOCs occupations with at least 10 or 30 survey respondents, depending on the column. Panel A uses the employment-weighted Eloundou et al. (2024) β-score aggregated to 3-digit SOC; Panel B uses predicted adoption, the fitted value from a cross-sectional regression of the survey adoption rate on the Panel A exposure measure. Sample periods: February 2021–February 2025 for Lightcast postings; February 2021–November 2025 for CPS transitions.

Table A.7: Continuous Treatment DID: AI Exposure and Labor Market Outcomes using Anthropic Measure

	Vacancies	Employment Transitions	
	log(postings)	Emp→Unemp (*EU*)	Emp→NILF (*EN*)
Panel A. Exposure (Massenkoff and McCrory (2026))			
$\text{Exposure}_o \times Post_t$	−0.297	0.634	−0.031
	(0.103)	(0.637)	(0.400)
Panel B. Adoption Predicted by Occupational Exposure			
$\widehat{\text{Adoption}}_o \times Post_t$	−1.118	2.385	−0.115
	(0.389)	(2.398)	(1.505)
Occupation & Month FE	✓	✓	✓
Occupation Cells	85	85	85
Observations	4,158	4,929	4,930

Notes: Each column-panel cell reports a separate regression of equation (1) at the 3-digit SOC × month level. Standard errors clustered at the 3-digit SOC level in parentheses. $Post_t$ indicates months from December 2022 onward. *EU* and *EN* transition rates are expressed in monthly percentage points. The sample is restricted to 3-digit SOCs occupations with at least 20 survey respondents. Panel A uses the Massenkoff and McCrory (2026) (Anthropic) score aggregated to 3-digit SOC; Panel B uses predicted adoption, the fitted value from a cross-sectional regression of the survey adoption rate on the Panel A exposure measure. Sample periods: February 2021–February 2025 for Lightcast postings; February 2021–November 2025 for CPS transitions.

Table A.8: Binary Treatment DID: AI Exposure and Labor Market Outcomes

	Vacancies	Employment Transitions	
	log(postings)	Emp→Unemp (EU)	Emp→NILF (EN)
Exposure$_o$ $\times$ $Post_t$	-0.0449	0.202	0.144
	(0.0355)	(0.194)	(0.106)
Occupation & Month FE	✓	✓	✓
Occupation Cells	85	85	85
Observations	4,158	4,929	4,930

Notes: Each column-panel cell reports a separate regression of equation (1) at the 3-digit SOC $\times$ month level. Standard errors clustered at the 3-digit SOC level in parentheses. $Post_t$ indicates months from December 2022 onward. EU and EN transition rates are expressed in monthly percentage points. The sample is restricted to 3-digit SOCs occupations with at least 20 survey respondents. Treatment group includes occupations in the top 25% of Eloundou et al. (2024) β exposure, while control includes occupations in the bottom 25%. Sample periods: February 2021–February 2025 for Lightcast postings; February 2021–November 2025 for CPS transitions.

B Survey Instrument

This appendix reproduces the survey instrument as fielded on IncQuery for the 2026Q1 wave; earlier waves are nearly identical (see Section 2.1 for differences in question placement across waves). Question numbers are platform identifiers and are therefore not sequential; questions appear below in display order. Color coding gray = question format, blue = coding instructions and display logic, red = termination logic, green = reader notes.

Screening Questions

1. Thank you for agreeing to participate. Honest responses and your full attention are very important for research such as this. Are you able to give your undivided attention to answer to the best of your ability to give the most accurate answers to this survey?
Multiple choice | Required | Single-select

Yes

No **[TERMINATE]**

2. What is your age?
Number | Required | Min: 0 | Max: 130
If answer < 18, **TERMINATE**. Flag qc_age if answer > 100.

3. What is your gender?
Multiple choice | Required | Single-select

Male

Female

Other

Prefer not to say

4. What is the highest level of school you have completed or the highest degree you have received?
Multiple choice | Required | Single-select

Did not complete 12th grade (no high school degree)

High school graduate or the equivalent (for example: GED)

Some college but no degree

Associate's degree in college

Bachelor's degree (for example: BA, AB, BS)

Graduate degree (for example: Master's, Professional, or Doctorate degree)

5. What is your zip code?
Open-ended | Required | Single line | Length: 5
Set state to state for zip code. Set region to US census region.

6. Do you consider yourself to be of Hispanic origin?
Multiple choice | Required | Single-select

Yes

No

7. Here is a list of categories of race and origin. Please select all that apply to you. You may choose more than one.
Multiple choice | Required | Multi-select

White

Black or African American

American Indian or Alaska Native

Asian

Native Hawaiian or Other Pacific Islander

Hispanic Origin

Other

8. What was the total income of your household in the last year (2024) before any taxes or deductions? Please include all sources of income from everyone who currently lives with you.
Multiple choice | Required | Single-select

Less than $10,000

$10,000 to $14,999

$15,000 to $24,999

$25,000 to $34,999

$35,000 to $49,999

$50,000 to $74,999

$75,000 to $99,999

$100,000 to $149,999

$150,000 to $199,999

$200,000 or more

Industry and Occupation

9. Which of the following best describes your employment status?
Multiple choice | Required | Single-select

Employed full-time (30 hours or more a week)

Employed part-time (less than 30 hours a week)

Contractor / Freelancer / Temporary Employee

Self-employed

Retired

Unemployed

Temporarily laid off

None of the above

Show the following block if Q9 is any of “Employed full-time,” “Employed part-time,” “Contractor / Freelancer / Temporary Employee,” “Self-employed”:

10. In this job, what type of employer do you work for?
Multiple choice | Required | Single-select

Government (including state or local governments, a public school, university or hospital)

Private-sector, for-profit company

Non-profit organization (including charitable organizations)

Self-employed

Work in a business owned by someone else in this household

11. What kind of business or industry is this job?
Multiple choice | Required | Single-select

Agriculture, Forestry, Fishing and Hunting

Mining, Quarrying, and Oil and Gas Extraction

Utilities

Construction

Manufacturing

Wholesale Trade

Retail Trade

Transportation and Warehousing

Information Services (including Publishing, Media, or Tech)

Banking, Finance, or Insurance

Real Estate, or Rental and Leasing

Professional, Technical, or Business Services

Education

Health Care and Social Assistance

Arts, Entertainment, and Recreation

Hotel, Accommodation, Restaurant, or Food Services

Other Services (except Government)

Government, including Military

Management of Companies and Enterprises

Administrative and Support and Waste Management and Remediation Services

57. What best describes your occupation in terms of day-to-day responsibilities?
Open-ended | Required | Multi-line
Set occupation_types to a mapping from each occupational group to its detailed occupation list (22 major groups, ~570 detailed occupations, following the BLS Standard Occupational Classification). Omitted here for brevity.

58. What is your occupational group?
Multiple choice | Required | Single-select | Randomize

Management Occupation

Business and Financial Operations Occupation

Computer and Mathematical Occupation

Architecture and Engineering Occupation

Life, Physical, and Social Science Occupation

Community and Social Service Occupation

Legal Occupation

Educational Instruction and Library Occupation

Arts, Design, Entertainment, Sports, and Media Occupation

Healthcare Practitioners and Technical Occupation

Healthcare Support Occupation

Protective Service Occupation

Food Preparation and Serving Related Occupation

Building and Grounds Cleaning and Maintenance Occupation

Personal Care and Service Occupation

Sales and Related Occupation

Office and Administrative Support Occupation

Farming, Fishing, and Forestry Occupation

Construction and Extraction Occupation

Installation, Maintenance, and Repair Occupation

Production Occupation

Transportation and Material Moving Occupation

94. What type of occupation?
Multiple choice | Required | Single-select
Dynamic choices: occupation_types for the occupational group selected in Q58.

AI Usage

Set hover definitions shown to respondents: (a) "**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." (b) "A **large language model** is a type of artificial intelligence (AI) model that uses machine learning to understand and generate human language."

12. Have you heard of generative AI before?
Multiple choice | Required | Single-select

Yes

No **[TERMINATE]**

Job risk block; randomize question order:

84. What do you think is the probability (0-100%) that you will lose your job due to generative AI in the next two years?
Number | Required | Min: 0 | Max: 100
__________ %

85. What do you think is the probability (0-100%) that at least one of your coworkers will lose their job due to generative AI in the next two years?
Number | Required | Min: 0 | Max: 100
__________ %
Allow skipping with checkbox labeled: "I do not have any coworkers."

13. Have you used generative AI tools before?
Multiple choice | Required | Single-select

Yes

No **[TERMINATE]**

14. Do you use generative AI for your job?
Multiple choice | Required | Single-select

Yes

No **[TERMINATE]**

Generative AI Usage at Work

15. Check off which generative AI tools you use.
Multiple choice | Required | Multi-select

ChatGPT

Midjourney

Claude

Gemini

Github Copilot

Scribe

Embedded products

DeepSeek

Apple Intelligence

Meta AI (Llama)

Adobe Firefly

Other (please specify) __________

16. How many days each week do you use generative AI tools at work?
Number | Required | Min: 0 | Max: 7 | Decimals: 0

17. How many hours per week do you use generative AI tools at work?
Number | Required | Min: 0 | Max: 168 | Decimals: 1

18. When did you most recently use generative AI tools at work?
Multiple choice | Required | Single-select

During the most recent day at work

During the most recent week at work

During the most recent month at work

During the most recent year at work

More than a year ago

19. In what tasks are you using generative AI/Large Language models at work?
Multiple choice | Required | Multi-select | Randomize
Tasks listed in Appendix Table A.2

____

20. In reference to your most recent task in which you used generative AI at work, how long did it take you to complete the task using generative AI? (in number of minutes)
Number | Required | Min: 0
___________ minutes

21. How did you use generative AI for this task?
Multiple choice | Required | Single-select | Randomize

GenAI did the task for me.

GenAI helped me to do the task quicker.

GenAI provided information/advice/strategies on how to do the task better.

GenAI verified checked the task for me.

Other ___________

22. In reference to your most recent task in which you used generative AI at work, how long would it have taken you to complete the same task <u>without</u> generative AI?
Combination (hours + minutes) | Required
Ensure minutes row < 60. Flag qc_time_to_complete_without_AI if time to complete without AI is shorter than time to complete with AI.

95. In the last three months, has the number of tasks increased, decreased, or remained the same?
Multiple choice | Required | Single-select

Increased

Decreased

Remained the same

96. In the last three months, have your tasks:
Multiple choice | Required | Single-select

Changed significantly (doing mostly different tasks from before)

Changed slightly (doing mostly similar tasks but with different complexities)

Remained the same

Employment Details

Show the following block if Q9 is "Employed full-time" or "Employed part-time":

23. How many days do you usually work **from home** in a given week?
Number | Required | Min: 0 | Max: 7 | Decimals: 2
__________ days

24. How many hours per week do you usually work for your job? (0-168)
Number | Required | Min: 0 | Max: 168
__________ hours

25. How many days do you usually work for your job? (0-7)
Number | Required | Min: 0 | Max: 7 | Decimals: 1
__________ days
Flag qc_num_days_worked if worked fewer days than days worked from home.

26. How many days do you usually commute to your job?
Number | Required | Min: 0 | Max: 7 | Decimals: 1
__________ days
Flag qc_num_days_commuted if worked fewer days than days commuted. Flag qc_commute_plus_WFH if days worked from home + days commuted does not equal total days worked.

Show Q27 if Q26 answer > 0:

27. Which of the following best explains why you commuted to work last week?
Multiple choice | Required | Single-select

Some aspects of my job could not be done from home

Some or all of my job could have been done from home, but my employer required me to commute

Some or all of my job could have been done from home, but I preferred to commute

28. For how long have you been working for this employer?
Multiple choice | Required | Single-select

0-2 years

3-5 years

6-9 years

10 or more years

Show the following block if Q9 is "Unemployed":

29. When was the last time you worked for pay or profit?
Multiple choice | Required | Single-select

1 to 2 months ago

2 to 6 months ago

6 to 12 months ago

One to two years ago

More than two years ago

30. Have you searched for a job in the past 3 years, since November 2022?
Multiple choice | Required | Single-select

Yes

No

Show Q31 if Q30 is “Yes”:

31. Did you use generative AI to search for a job? For example, to help write cover letters or resumés, prepare for interviews, or conduct research about job openings.
Multiple choice | Required | Single-select

Yes

No

32. Do you directly use computer at home?
Multiple choice | Required | Single-select

Yes

No

33. Do you use computer for your job?
Multiple choice | Required | Single-select

Yes

No

34. Do you use internet at any location?
Multiple choice | Required | Single-select

Yes

No

Final Demographics

35. How would you describe your marital status?
Multiple choice | Required | Single-select

Married, spouse currently lives in this household

Married, spouse currently lives outside this household

Widowed

Divorced

Separated

Never Married

Show Q36 if Q35 is “Widowed,” “Divorced,” “Separated,” or “Never Married”:

36. Do you have a partner/boyfriend/girlfriend who currently lives in this household?
Multiple choice | Required | Single-select

Yes

No

37. Number of children under age 18 in the household
Multiple choice | Required | Single-select

No children

1 child

2 children

3 or more children

Show Q38 if Q37 indicates one or more children:

38. What is the age of the youngest child under age 18 that currently lives with you?
Number | Required | Min: 0 | Max: 18

54. What year were you born?
Number | Required | Min: 1900 | Max: 2010
Flag qc_age_birthyear if implied age is more than a year off from age reported in Q2.

39. Attention Check: The color of grass is green. This actually is an attention check question. Make sure that you select purple as an answer so we know you are paying attention.
Multiple choice | Required | Single-select

Green

Purple

Blue

Black

White

Brown Flag qc_attention_check1 if selected choice is not "Purple."