ECONOMICS WORKING PAPERS

Job Loss Fears in the First Years of Generative Artificial Intelligence

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

Economics Working Paper 26125

HOOVER INSTITUTION
434 GALVEZ MALL
STANFORD UNIVERSITY
STANFORD, CA 94305-6010
August 2026

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 large-scale, multi-wave survey measuring workers' perceptions of AI-related job loss risk and AI adoption in the workplace, which we estimate at 30–40% of U.S. workers through the first half of 2026. 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, which is comparable to workers' realized risk of job loss from any cause over a similar horizon. 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. Instead, 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 independent of any realized job loss.

---

* Hartley: UT-Austin School of Civic Leadership and Visiting Fellow, Hoover Institution, Stanford University, jonathan.hartley@austin.utexas.edu, Jolevski: World Bank, fjolevski@worldbank.org, Melo: RAND Corporation, vitormelo95@gmail.com, Moore: University of Michigan, moorebd@umich.edu. This paper previously circulated under the title "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 Daron Acemoglu, Susan Athey, David Autor, Tania Babina, Jose Maria Barrero, Nick Bloom, Erik Brynjolfsson, John Cochrane, Mihai Codreanu, Tyler Cowen, David Deming, Josh Gans, Josh Gottlieb, Erik Hurst, Melissa Kearney, Brian Kovak, Liya Palagashvili, Lindsey Raymond, Eric Schmidt, Gonzalo Schwarz, Jesse Shapiro, and Isaac Sorkin. 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). In recent years, technologists, media commentators, and academics have increasingly warned that advances in AI may lead to significant job displacement. (Amodei, 2026; Steinberger, 2026; Brynjolfsson et al., 2025a) Large technology and telecommunications firms such as Amazon, Microsoft, Intel, Verizon, IBM, and Block have cited AI as a reason for large layoffs since 2025. Firms report that they expect AI to reduce headcount (Davies et al., 2026; Yotzov et al., 2026) and are increasingly evaluating employees on AI adoption in 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 purpose and future employability (McNamara and Thornton, 2025).

Despite many reasons for concern, an emerging empirical literature tends to show that the labor market effects of LLMs are fairly limited in terms of employment and wages (Humlum and Vestergaard, 2026). Even papers documenting negative effects of LLMs do so only for a small group of younger and entry-level workers in highly exposed occupations (Brynjolfsson et al., 2025a; 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 technology use shapes those beliefs, and how sharply they diverge from what the data show.

This paper pursues three questions about the labor market consequences of generative AI. First, has the diffusion of GenAI produced detectable effects on labor demand and job losses in the first several years since ChatGPT's public release? Second, how large is workers' fear of losing their jobs due to AI? Third, if realized disruption remains limited, are these fears instead shaped by how intensively and for which tasks workers use the technology?

To address these questions, we conduct a large-scale survey of U.S. workers, fielded quarterly across six waves from December 2024 through June 2026, with roughly 25,100 total respondents. We document that GenAI adoption at work reached approximately 35% among U.S. workers by the end of 2025 (a level similar to contemporaneous survey evidence from Bick et al. 2026). Consistent with previous work, we show adoption is sharply stratified by education and income, with graduate-degree holders and workers earning above $200,000 annually adopting at higher rates.

We then investigate whether AI exposure and adoption have led to the much-feared labor market disruption through 2025. To do so, we use a continuous-treatment difference-in-differences design, treating occupations as exposed to varying doses of AI, applied to Current Population Survey transitions into unemployment or non-employment and Lightcast vacancy postings. We do not detect evidence of AI-related layoffs or a decline in vacancy postings that has animated public discussion. While point estimates directionally indicate disruption (positive for layoffs, negative for vacancies), they are quantitatively small and not statistically significant. These nulls are robust to alternative sample restrictions, to a binary top-versus-bottom-quartile design, and to substituting an alternative occupational exposure measure. This lack of evidence for substantial GenAI labor market disruption is consistent with much of the existing literature (Humlum and Vestergaard, 2026; Chandar, 2025; Hellsten et al., 2025; Aldasoro et al., 2026; Hyman et al., 2026).

Despite this lack of detectable disruption, U.S. workers report strikingly high fear of AI-driven displacement: the average worker believes there is a one-in-five chance of losing their job due to AI over the next two years according to our survey. Among workers actively using GenAI on the job, this figure rises to 25%. These fears are particularly noteworthy given the average annual *realized* job loss probability (for any reason) of around 16–18%. In other words, workers fear AI-induced displacement at a rate that rivals their total contemporaneous job loss risk from all sources. We show that more AI-exposed jobs (regardless of usage) tend to have higher fears of job loss due to the technology.

Why are workers' fears of job loss so elevated given the lack of realized disruptions? Our survey evidence suggests a learning-by-doing story, in which workers form beliefs about substitutability from direct experience with what the technology can do, rather than from observation of actual labor market outcomes. We arrive at this conclusion for two reasons. First, we show that the more intensively workers use GenAI on the job, the more they fear AI-induced job loss. This relationship holds even controlling for occupational group and demographic cells, suggesting that the more workers see the technology's capabilities, the more they worry about its capacity to replace their work.

Second, the specific tasks for which fearful and non-fearful adopters deploy GenAI correlate intuitively with core job functions. 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. In such cases, workers watch the technology perform substantive or technical tasks central to their occupation. Meanwhile, those with below-median AI job loss fear rely on it mainly for writing and learning strategies, where the technology is more likely to complement or accelerates workers' own output.

It is important to document the stark gap between workers' elevated anxiety about AI and the technology's lack of early disruption for several reasons. First, it is useful to have a historical record of workers' job loss expectations regarding a potentially transformational technology at the onset of its widespread adoption. Prior waves of technological change have generated similar anxieties, but fears were documented only after the realized disruption, rather than measured contemporaneously as workers form their expectations (Mokyr et al., 2015). Second, heightened fear of job loss can have real economic consequences, independent of actual realization. For example, workers alter consumption in anticipation of layoffs even before actual displacement (Stephens Jr, 2004; Karahan et al., 2019). Such fears can also depress occupational entry, as Salman (2026) documents for commercial truckers who are more exposed to the capabilities of autonomous vehicles. Displacement risk also affects financial investing by boosting the equity premium (Gârleanu et al., 2012) and steering workers to direct search towards less risky firms (Lehner et al., 2025).

In addition to the literature on job loss expectations, we also contribute to a growing literature on AI's adoption and labor market impacts (Chen and Wang, 2024; Brynjolfsson et al., 2025b; Engberg et al., 2025; Hampole et al., 2025). We complement recent survey evidence from Bick et al. (2026), extending their analysis to the question of subjective displacement risk of AI. Our labor market analysis builds on the event-study designs of Brynjolfsson et al. (2025a) and Liu et al. (2025), who document negative employment effects for entry-level workers.

The remainder of the paper proceeds as follows. Section 2 describes our survey design, data, and documents AI adoption patterns. Section 3 presents our empirical analysis of labor demand and job separations using Lightcast job postings data and the Current Population Survey. Section 4 uses our survey to characterize workers' beliefs about displacement risk due to AI and the utilization and task-level correlates of those beliefs. Section 5 concludes.

2 Survey and Labor Market Data

2.1 Generative AI Adoption and Job Loss Risk Survey

This paper's first contribution is to conduct a survey of workers (similar to the Pew and Bick et al. (2026) GenAI surveys) with IncQuery asking about the use of LLMs in specific tasks and occupations. For the purpose of this paper, the survey's two main goals are (1) identifying which occupations are using LLMs and in what ways and (2) documenting AI-related job displacement risk and how it varies across workers according to their exposure and interaction with the technology.[1]

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 9,000 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 is workers in the U.S. The survey was self-administered 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. However, we differ in that our survey primarily focuses on the intensive margin of AI usage, whereas many of their questions focus on the extensive margin.[2] We further distinguish our survey from others that

[1] In this way, we build on Bloom and Makridis (2026) who survey 10,360 workers via Gallup about AI displacement risks. However, the authors only focus on political partisan splits and do not relate job loss fears to AI adoption or utilization.

[2] We follow the authors by defining GenAI: "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

have asked about workers' fear of job loss, such as Pew, which elicit categorical response options (e.g., how "worried" versus "hopeful" workers feel) (Lin and Parker, 2025) rather than a continuous, numeric subjective probability of job loss due to AI, as we describe below.

Specifically, we ask "What do you think is the probability (0-100%) that you will lose your job due to generative AI in the next two years?" The elicitation of job loss expectations as a numeric probability has a longstanding tradition in the literature on labor market expectations (Manski and Straub, 2000; Mueller and Spinnewijn, 2023). We similarly ask workers about the probability of their co-workers losing their jobs due to AI.

The survey also asks "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. 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. In wave 6 (Mar. 2026) it was asked to all respondents. It was not asked in wave 1 (Dec. 2024).

For the purpose of maximizing our sample size of respondents adopting GenAI, we ask demographic, occupation, and industry questions before screening out respondents who do not use GenAI at work. We maintain the demographic data of those screened out in order to produce estimates of AI use by occupation and demographic characteristics. 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 implement use several attention checks, one before the screener and one after.

Among those screened for AI use on the job, we ask how intensively they use the technology. 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?"

To understand what tasks for which workers use LLMs, we 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 the full list of options).

See Appendix C for our complete survey instrument.

2.2 Exposure and Adoption Measures

Our preferred measure of occupational exposure is Eloundou et al.'s (2024) β-score, which describes how exposed occupations are to LLMs with the assistance of an additional technology. Specifically, this score measures the capacity of LLMs 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. We prefer this exposure measure because our paper focuses on subjective job loss risk, and the Eloundou et al. (2024) score is the metric most closely related to

Model is a type of artificial intelligence model that uses machine learning to understand and generate human language."

task displacement. Measures that instead capture LLM's capacity for *augmentation* are less appropriate since they may not track workers' displacement fears which are likely borne of substitution.

We can relate the occupational exposure score to an occupation's GenAI adoption at work in our survey. AI exposure strongly predicts adoption among occupation groups, with an estimated coefficient of 0.55 (similar to 0.67 documented by Bick et al. (2026) with their survey adoption measure). Appendix Figure A.11 shows this relationship. For example, architecture/engineering and business/finance occupations have both high predicted AI exposure and high 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.

The "predicted adoption" measure will be used in section 3 for the reduced form analysis of the effects of AI on the labor market. This predicted adoption measure allows us to interpret regression coefficients in more transparent units of adoption rather than units of exposure at the occupational level. In each case, the underlying variation (and thus the precision of the estimates) is the same.

We estimate an average GenAI adoption rate of roughly 40% throughout 2025[3], closely in line with estimates from Gallup, Bick et al. (2026), and the New York Fed Survey of Consumer Expectations (Hashim et al., 2026). Appendix Figure A.1a shows the time series of adoption through 2026Q2. As shown in prior work (Pulito et al., 2026), adoption of GenAI at work exhibits strong positive gradients in education and income. The highest GenAI adoption rates are in information services, management, real estate, construction, and education, with substantially lower rates in sectors such as agriculture, mining, and government (Appendix Figure A.5). Conditional on adoption, one-third of users report using GenAI every weekday (Appendix Figure A.8), and the median user works 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.3–A.10.

2.3 Labor Market 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), pro-

[3]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 that estimates how many total Americans a given respondent should represent for a particular age/sex/education/race cell.

viding nationally representative data on employment status, occupation, and industry for the civilian population aged 16 and older.

To study job loss, we use labor market transition rates which are constructed 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.[4]

To define layoffs, use the CPS to observe 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 minor group occupation (3-digit) 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. Lightcast collects job postings by scraping listings from a variety of employer websites and online job posting websites. Hershbein and Kahn (2018) validate the Lightcast (formally Burning Glass) job postings data using Job Openings and Labor Turnover Survey, showing that aggregate and industry-level vacancy patterns are broadly-similar between the two sources.

3 Effects of LLMs on the Labor Market

3.1 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 following 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-3.5 rollout:

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

[4]We avoid the COVID-19 period because of issues with estimating occupation fixed effects for March and April 2020, which included historically large layoffs due to the onset of the pandemic. In particular, including these months in the pre-period estimation may bias the estimation of a given occupation's separation or vacancy rate because the COVID-19 shock may be systematically correlated with AI exposure. Details about our estimation strategy are in Section 3.

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 this reason, we do not regress labor market outcomes on our *adoption* measure from the survey, as adoption may be endogenous to labor market conditions. 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 translation obviates this concern by isolating only the variation in adoption attributable to theoretical exposure.

We harmonize all data sources in the difference-in-differences analysis to the 3-digit SOC code level (94 occupation groups)[5], including aggregating the occupation-level exposure β-scores using 2022 Occupational Employment Statistics employment weights. 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 because it is a theoretical measure of exposure, whereas Anthropic's score is constructed largely from observed usage. 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).

3.2 Results

Table 1 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. The pattern of point estimates is consistent in that vacancy effects are negative and separation effects are positive. Nevertheless, in the two and a half years following the release of ChatGPT, we can rule out anything

[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.

close to the widespread displacement (consistent with Brynjolfsson et al. (2025a) and Humlum and Vestergaard (2026)).

Table 1: 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.

Vacancies

We find no evidence that GenAI exposure has reduced job openings. The continuous-treatment estimates in the first column of Table 1 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 $N \geq 20$ 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).

Only when we use the Massenkoff and McCrory (2026) (Anthropic) exposure score in place of the Eloundou et al. (2024) β-score does the vacancy posting coefficient become negative and statistically significant at the 1% level (Appendix Table A.7). This point estimate implies 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 as a sign of a labor market disruption because, as discussed in Section 3.1, the Anthropic score is constructed from observed patterns of AI usage rather than from task content. Therefore, 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 (see columns 2 and 3 of Table 1). The estimated effect on the monthly EU rate is 0.47 percentage points per unit of the Eloundou et al. (2024) β-score (standard error of 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.

Appendix Figure A.13 serves as a robustness check by aggregating job transitions to a quarterly frequency and running an above- and below-median employment-weighted AI exposure event study where we consider below-median occupations to be untreated and mark above-median occupations with a binary treatment indicator. Across both panels, coefficients are small and statistically indistinguishable from zero throughout the sample period. This also validates our empirical approach by showing no evidence of differential pre-trends prior to the ChatGPT release, supporting the parallel trends assumption.

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.

Thus, we show that workers in AI-exposed occupations are not actually losing their jobs at higher rates or facing dimmer job prospects than their counterparts in less exposed occupations. However, as we will document in the next section, workers have substantial fears about job insecurity due to AI. Given these fears are not formed from seeing realized disruption, we investigate this gap between perceived and realized displacement risk and consider its sources.

4 Perceived Job Loss Risk due to AI

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 1 shows the distribution of these job loss probabilities for GenAI adopters at work (red) and non-adopters (blue). The main difference between adopters and non-adopters is that the latter are more certain of their job security, with 53% reporting less than a 5% chance of AI-induced job loss. This compares to just 40% for AI adopters. A non-negligible share of adopters (10%) assess 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 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 that workers change their behaviors of consumption (Stephens Jr, 2004; Hendren, 2017) and labor market decisions (Lehner et al., 2025; Salman, 2026) in response to layoff risk.

We validate the quality of this subjective job loss measure in several ways. First, we compare the distribution of our AI-job loss expectations response to those from two other 2026 surveys from Verasight and Ipsos posing similar questions to U.S. workers. Both surveys asked, without a time horizon, "How much do you think you are at risk losing your job due to AI?". Verasight allowed answers of quantifiable bins (e.g., No risk, 1–25%, 26–50%, More than 50%) while Ipsos allowed qualitative answers (e.g., "Not at all," "Only a little," "A fair amount," "A great deal"). We map the qualitative answer options from Ipsos to the numeric job loss risk bins from Verasight and compare both measures to our survey measure (described in more detail in section 2.1). Appendix Figure B.1 compares our March 2026 survey measure (conducted among all workers) to the external sources and documents a very similar distribution. Thus, we can feel confident that our survey measure is capturing true AI job loss fear among U.S. workers, as reported by other outlets. More details about the survey question comparisons are provided in Appendix B.1.

Second, given the substantial fears of job loss reported in our survey, we test whether anchoring workers on realistic job loss probabilities affects the answers provided. As described in further detail in Appendix B.2, the seventh wave of our survey conducted in June 2026 (which is not used in the main analysis) caveats the question soliciting job loss expectations with the statement "about 15% of U.S. workers have involuntarily lost their job for any reason over a two-year period in recent years."[6] While this addition was not randomized across survey participants within wave, we compare the answers to this question with and without the anchoring. We show that contextualizing typical job loss rates does not noticeably affect workers' answers about their own probability of layoff due to AI.

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.

[6]This calculation is from the authors' investigation in the Survey of Income and Program Participation (SIPP).

Figure 1: Job Loss Perceptions Due to GenAI During the Next Two Years by Adoption

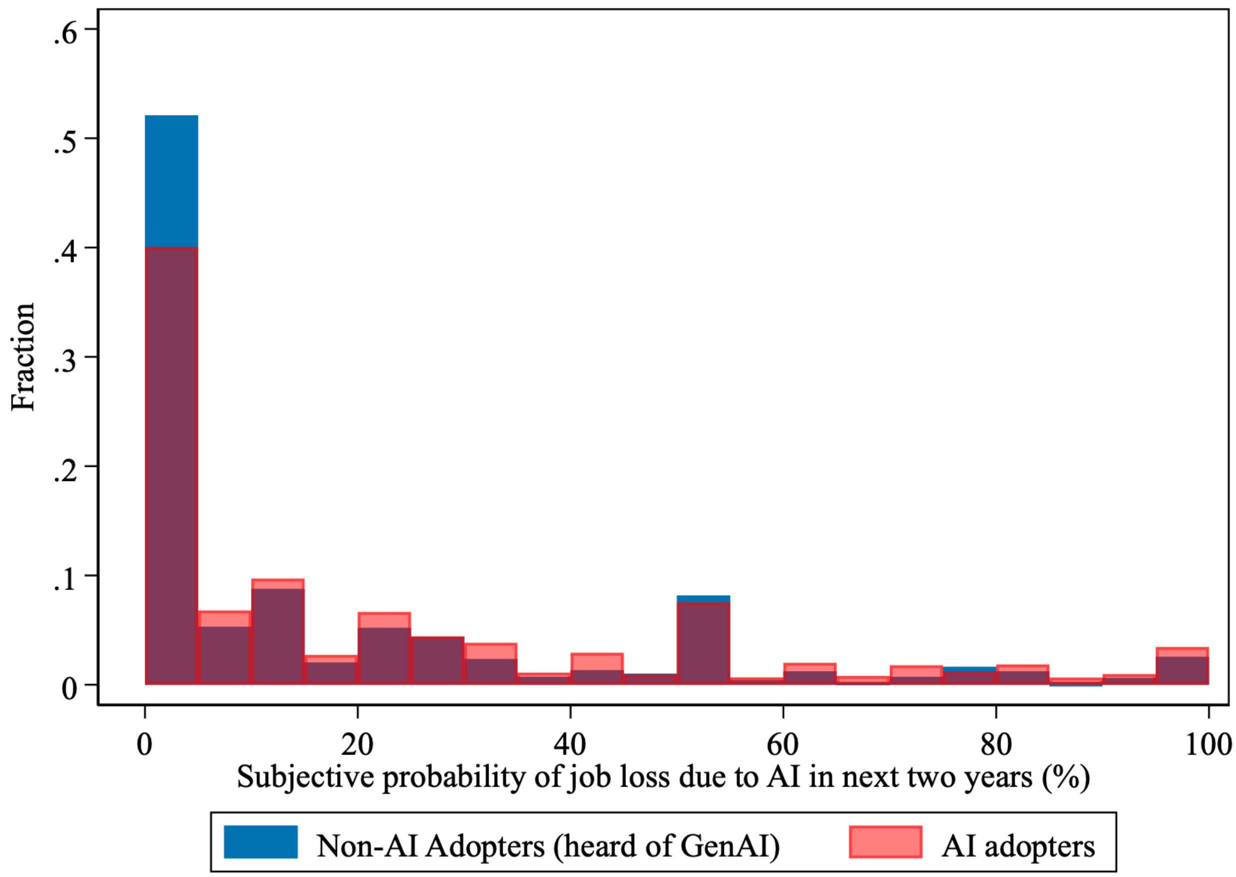

Source: Authors' survey. *Note*: Histogram of subjective job loss perceptions due to generative AI in the next two years among non-AI adopters at work but who have heard of generative AI (blue) and AI users at work (red). Survey sample is weighted population weights by age-sex-education-race from the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

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. Figure 2 illustrates this positive relationship between exposure and job loss fear at the three-digit SOC code level. For example, large occupation groups that are more exposed, such as "computer occupations" and "information and record clerks," exhibit heightened fear of job loss due to AI in our survey, exceeding 25%. Conversely, much less-exposed occupations such as K-12 teachers and health technicians report average fears of AI-induced displacement of less than 15%. Nevertheless, there is dispersion across similarly AI-exposed occupations: computer occupations report double the job loss fears (34%) compared to financial clerks (17%). Appendix Figure A.12 shows the histogram of job loss beliefs by exposure according to Eloundou et al. (2024), where above-median exposure occupations report a 21% expected job loss rate compared to 16% for below-median exposed occupations.

Second, we show that workers who use AI regularly for more hours per week at their job also

Figure 2: Job Loss Fears by GenAI Exposure Across Occupations

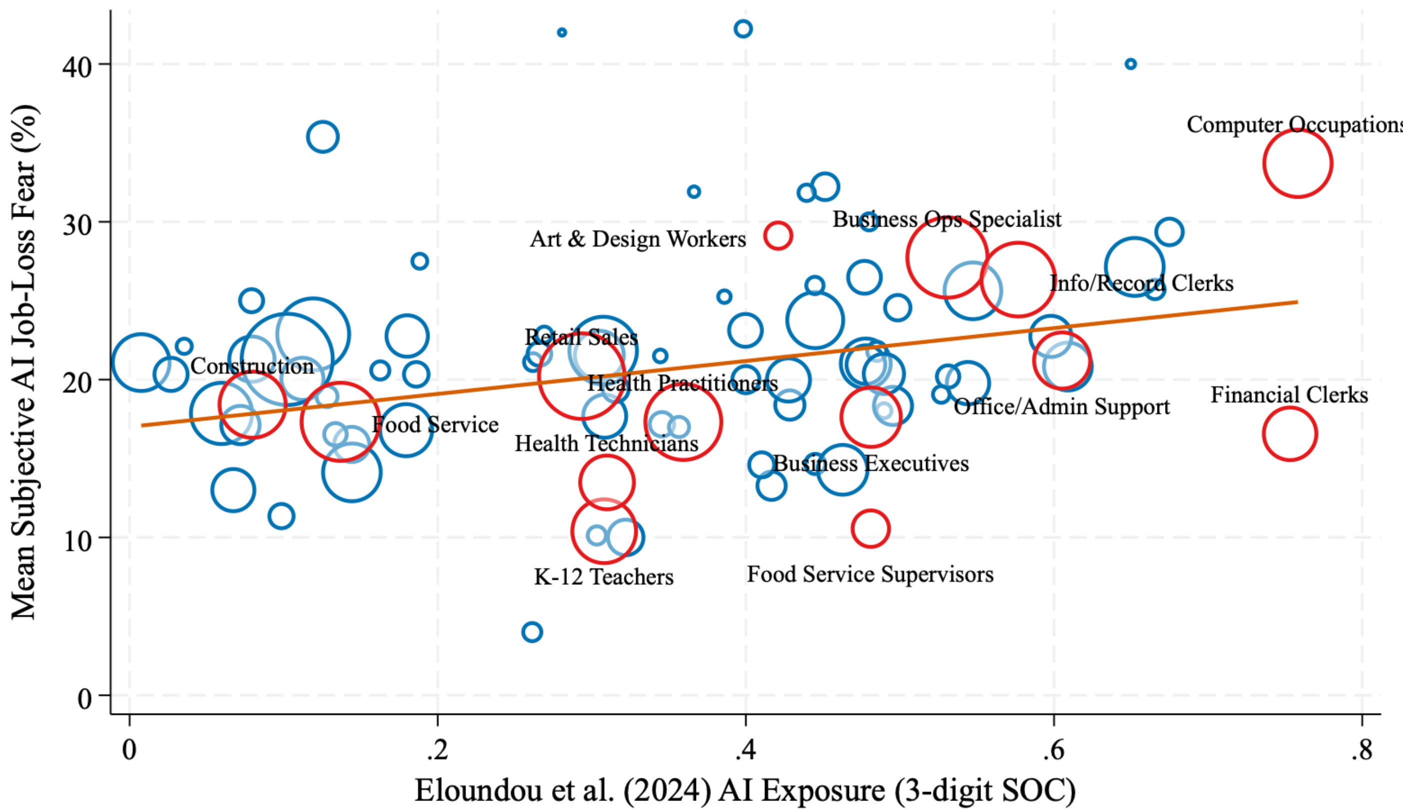

Note: Scatterplot at the 3-digit SOC occupation level for all cells with five or more respondents in 2026Q1 wave (84 occupations). Employment weighted slope of 10.43, *t*-value of 4.08. Marker sizes are proportional to 2022 OES employment share. Labeled markers for specific occupations are denoted with red circles.

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 2 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 one's 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 likely 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 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

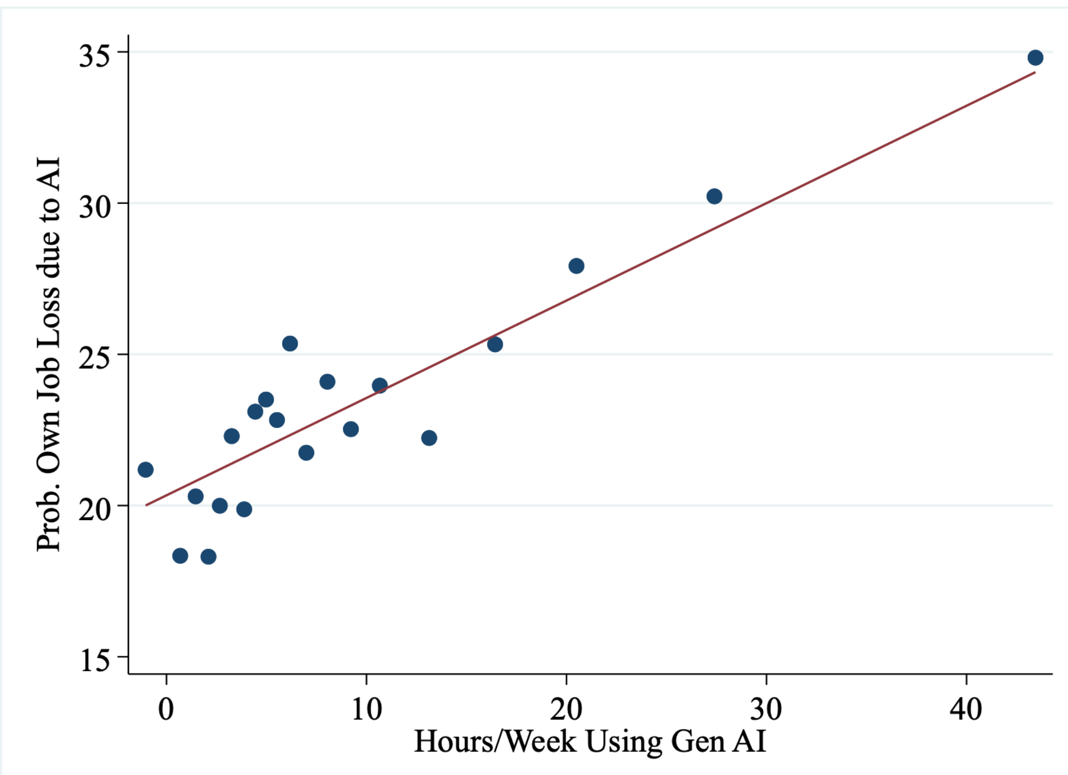

Note: Binscatter shows relationship between two-year subjective job loss rate due to AI with respect to the reported number of hours per week the respondent uses AI at work, among the sample that uses AI at work at all. Regression includes occupation group (2-digit SOC) fixed effects and controls for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

Table 2: GenAI Tasks at Work and AI Displacement Fears

Task	Share Using AI for Task Among: High Job-Loss Fear	Share Using AI for Task Among: Low Job-Loss Fear	High-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.

5 Conclusion

In the several years after ChatGPT's release, significant labor market disruption has not manifested. Using a large-scale, multi-wave survey of over 25,000 U.S. workers fielded between December 2024 and March 2026, we find that roughly 40% use GenAI on the job. Yet more AI-exposed jobs have not been differentially slow to hire or quick to fire from 2022 to 2025.

Against the backdrop of limited labor market disruption, workers express substantial fear about their job security due to GenAI. The average worker believes there is a one-in-five chance of being displaced by AI within two years – a figure that rises to one-in-four among those who use the technology at work. Workers' fear of AI-induced displacement alone is on par with recent realized job loss rates for any reason. Our survey provides a valuable, contemporaneous record of workers' fear of job loss as this revolutionary technology diffuses in real-time.

We provide suggestive evidence that these fears are amplified by workers watching GenAI perform a task central to one's own job. Workers in more AI-exposed occupations and those who use AI more intensively on the job report the highest displacement fears, a relationship that holds even controlling for occupation and demographics. Workers who deploy GenAI for core technical functions such as programming and technology design report markedly higher fears than those who use it primarily for writing. Together these patterns point to a learning-by-doing mechanism: workers form beliefs about the technology's capacity to replace them not from observing labor market outcomes, but from direct, hands-on experience with what the technology itself can do.

These perceptions have real economic consequences independent of whether displacement ever materializes, as workers adjust consumption and job search in anticipation of job loss (Stephens Jr, 2004; Lehner et al., 2025). Even without realized mass unemployment, heightened subjective displacement risk may distort job-to-job transitions by making workers excessively cautious (Clymo et al., 2026). We are still in the early years of generative AI's diffusion into the workplace, and both the technology's realized labor market impacts and workers' beliefs about it should continue to be monitored to understanding AI's ultimate consequences for the American workforce.

References

ALDASORO, I., L. GAMBACORTA, R. PAL, D. REVOLTELLA, C. WEISS, AND M. WOLSKI (2026): "AI adoption, productivity and employment: evidence from European firms," Tech. rep., EIB Working Papers.

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*.

BLOOM, N. AND C. MAKRIDIS (2026): "The Politics of AI," Tech. rep., National Bureau of Economic Research.

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 (2025a): "Canaries in the Coal Mine?: Six Facts about the Recent Employment Effects of Artificial Intelligence," .

BRYNJOLFSSON, E., D. LI, AND L. RAYMOND (2025b): "Generative AI at work," *The Quarterly Journal of Economics*, 140, 889–942.

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.

CHANDAR, B. (2025): "Tracking employment changes in AI-exposed jobs," *Available at SSRN 5384519*.

CHEN, M. A. AND J. WANG (2024): "Displacement or augmentation? The effects of AI innovation on workforce dynamics and firm value," *The Effects of AI Innovation on Workforce Dynamics and Firm Value (April 8, 2024)*.

CLYMO, A., P. DENDERSKI, Y. MERCAN, AND B. SCHOEFER (2026): "Cautious Careers: Job Mobility under Incomplete Markets," Tech. rep., National Bureau of Economic Research.

DAVIES, E., D. FRANCIS, F. JOLEVSKI, G. NAYYAR, K. STAPLETON, AND S. YU (2026): "Adopting Fast and Slow: Cross-Country Evidence on Business Adoption of Artificial Intelligence," .

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

ENGBERG, E., M. KOCH, M. LODEFALK, AND S. SCHROEDER (2025): "Artificial intelligence, tasks, skills, and wages: Worker-level evidence from Germany," *Research Policy*, 54, 105285.

GÂRLEANU, N., L. KOGAN, AND S. PANAGEAS (2012): "Displacement risk and asset returns," *Journal of Financial Economics*, 105, 491–510.

HAMPOLE, M., D. PAPANIKOLAOU, L. D. SCHMIDT, AND B. SEEGMILLER (2025): *Artificial intelligence and the labor market*, w33509, National Bureau of Economic Research.

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.

HELLSTEN, M., S. KHANNA, M. LODEFALK, AND Y. YAKYMOVYCH (2025): "The Effects of Artificial Intelligence on Jobs: Evidence from an AI Subsidy Program," .

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

HERSHBEIN, B. AND L. B. KAHN (2018): "Do recessions accelerate routine-biased technological change? Evidence from vacancy postings," *American Economic Review*, 108, 1737–1772.

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 (2026): "Still waters, rapid currents: Early labor market transformation under generative AI," *NBER Working Paper*, 33777.

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," .

LIN, L. AND K. PARKER (2025): "US workers are more worried than hopeful about future AI use in the workplace," *Pew Research Center*.

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*.

MANSKI, C. F. AND J. D. STRAUB (2000): "Worker Perceptions of Job Insecurity in the Mid-1990s: Evidence from the Survey of Economic Expectations," *The Journal of Human Resources*, 35, 447–479.

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.

MOKYR, J., C. VICKERS, AND N. L. ZIEBARTH (2015): "The history of technological anxiety and the future of economic growth: Is this time different?" *Journal of economic perspectives*, 29, 31–50.

MUELLER, A. I. AND J. SPINNEWIJN (2023): "Expectations data, labor market, and job search," *Handbook of Economic Expectations*, 677–713.

PULITO, G., M. PYTLIKOVA, S. SCHROEDER, AND M. LODEFALK (2026): "Who Adopts AI? Evidence on Firms, Technologies and Workers," Tech. rep., GLO Discussion Paper.

SALMAN, D. (2026): "End of the Road? Autonomous Vehicles and Displacement Risk," *Autonomous Vehicles and Displacement Risk (May 06, 2026). Kilts Center at Chicago Booth Marketing Data Center Paper*.

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.

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

Appendix Figure A.1 shows the time series of AI adoption at work in the U.S. (panel A.1a) and measured adoption in 2025 internationally for over a dozen countries (panel A.1b). U.S. adoption in 2025 is around 40%, 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 A.1a) and the New York Fed Survey of Consumer Expectations (Hashim et al., 2026). The late-2025 decline in measured adoption mirrors plateauing interest in GenAI tools, demonstrated by declining Google search volume for "ChatGPT" and "Gemini", which we document in Appendix Figure A.14. Internationally, the United States has one of the highest rates of AI adoption, along with Australia, the U.K., and Ireland. Italy, Canada, France and Japan all have lower rates of AI adoption.

With respect to education, we find in our survey (consistent with Pew and Bick et al. (2026)) that more educated workers are more likely to use GenAI (Appendix Figure A.3). 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.

Industries most likely to use AI at work include "Information Service" 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), shown in Appendix Figure A.5. Workers in industries such as "Agriculture, Forestry, Fishing," "Mining, Quarrying, Oil & Gas," and "Government & Military" use it far less.

AI adoption is increasing with incomes for the vast majority of workers. Those earning $75,000-$99,999 use AI at work twice the rate of those earning $25,000–$34,999 (Appendix Figure A.4). Over a third of those earning $100,000–$149,999 use AI at work, and the share is nearly half for those making above $200,000 annually.

With respect to gender, a greater proportion of men claim to use GenAI at work than women in the survey (Appendix Figure A.6). 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 (Appendix Figure A.7).

We also document important facts about the intensive margin of using GenAI tools. Appendix Figure A.8 shows that conditional on using AI at work, about 33% of workers use the technology five days per week (every weekday). Roughly 12% of adopter users use such tools at work only one day a week. 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). Appendix Figure A.9 shows that 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 that there is intermittent use of AI throughout the workday.

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

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

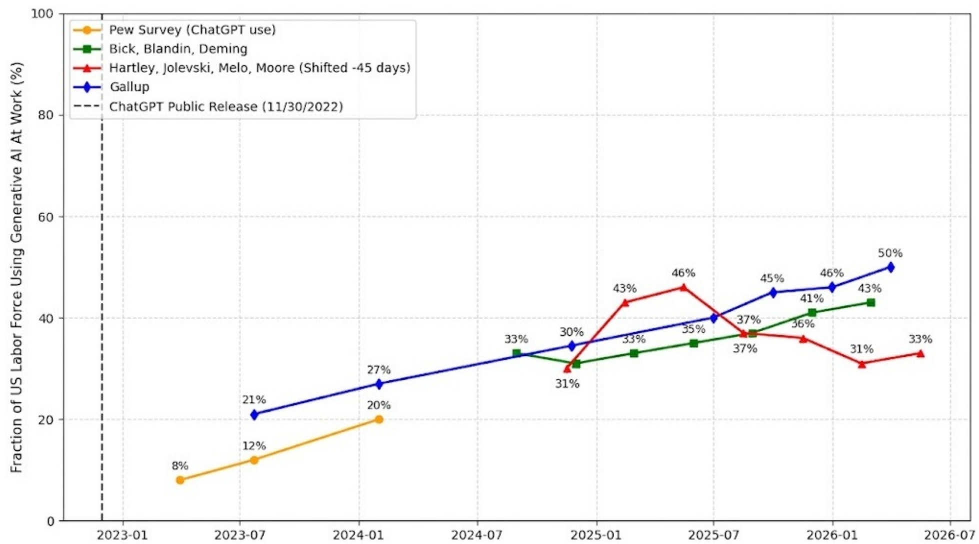

(b) Adoption at Work Across Countries in 2025

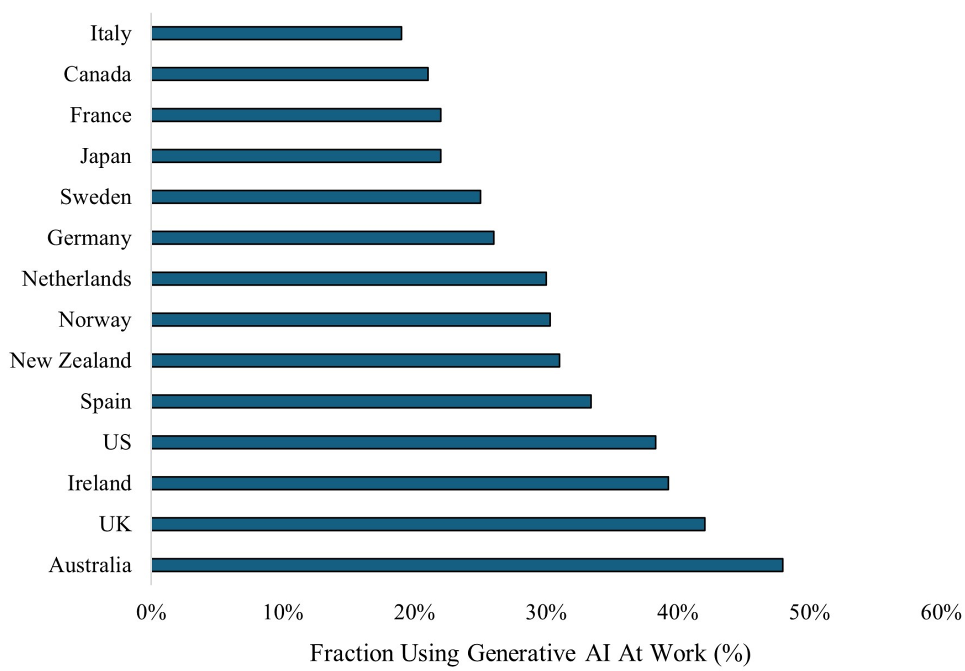

Note: Top panel shows adoption over time. First survey (orange) from Pew Charitable Trusts (2023-2024). Second survey (green) from Bick, Blandin, and Deming (2026) Third survey (red) from Hartley, Jolevski, Melo, and Moore (2026) Bottom panel shows international adoption in 2025 from Hartley, Jolevski,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

Figure A.2: Average Number of Hours of Generative AI Use At Work Over Time

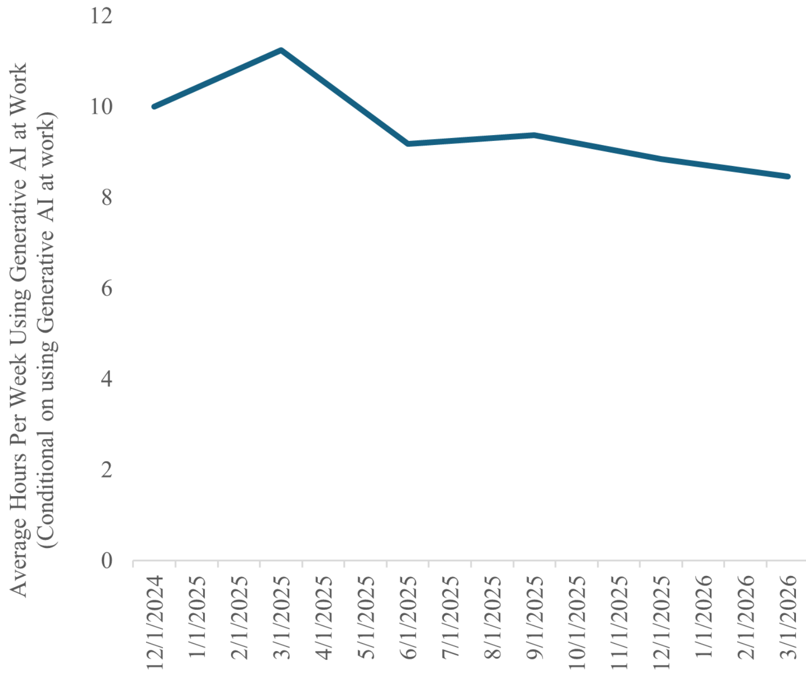

Source: Authors' survey.

Figure A.3: Generative AI Use By Level of Education

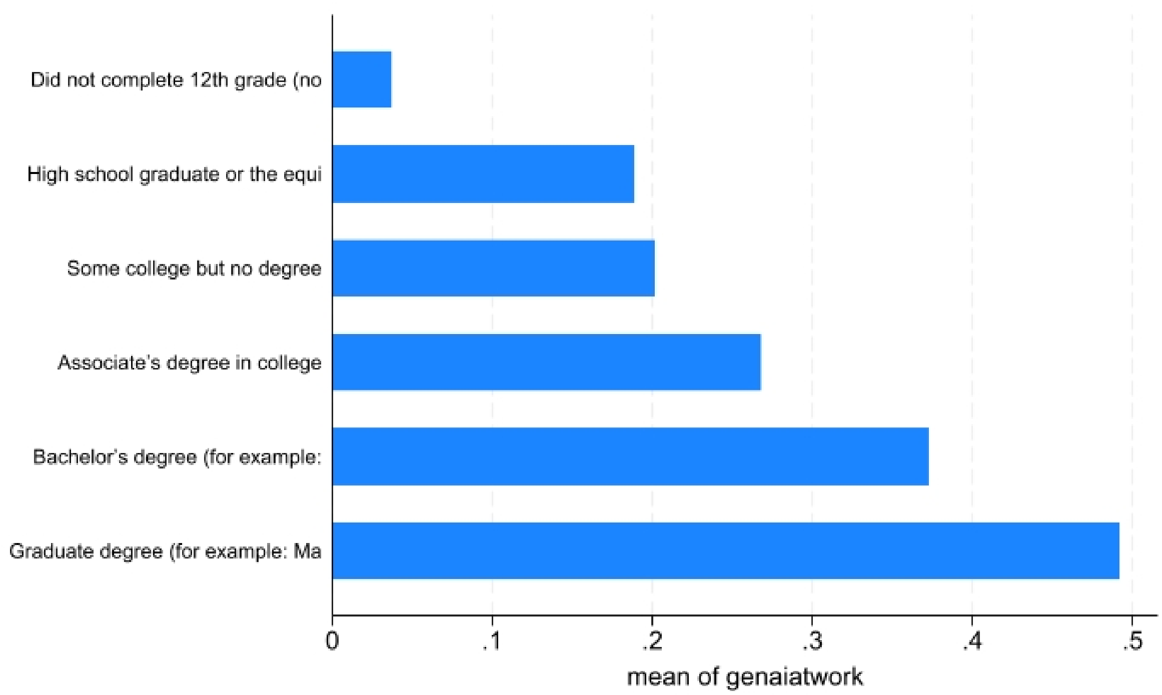

Source: Authors' survey.

Figure A.4: Generative AI Use By Income

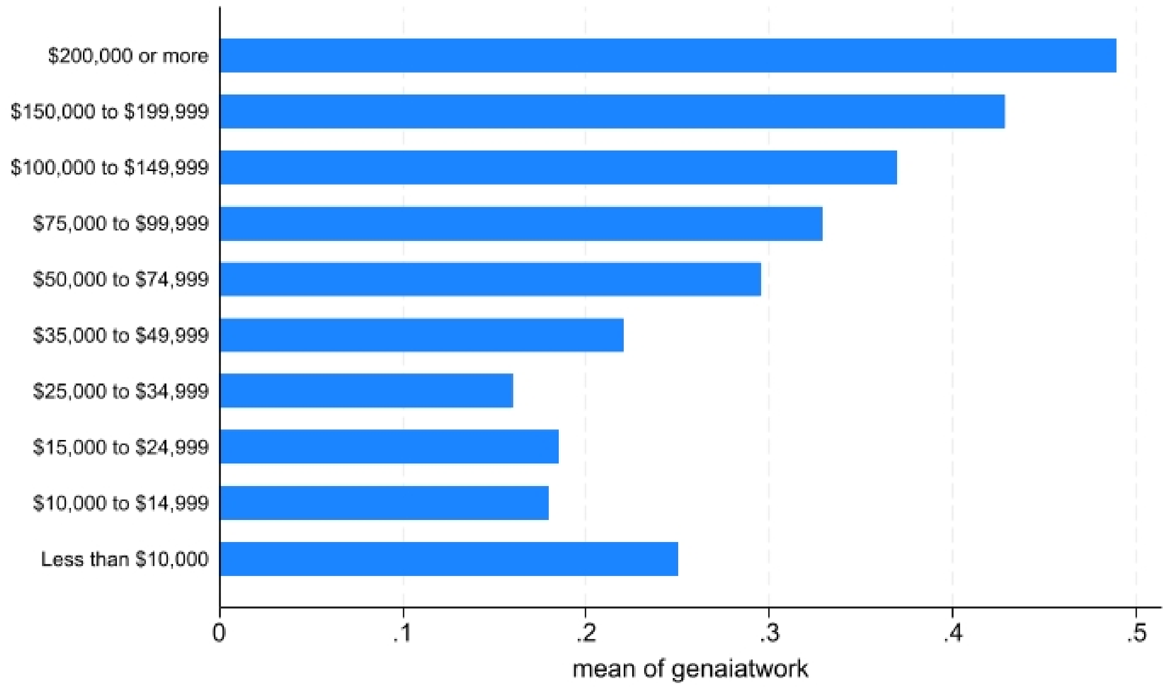

Source: Authors' survey.

Figure A.5: Generative AI Use By Industry

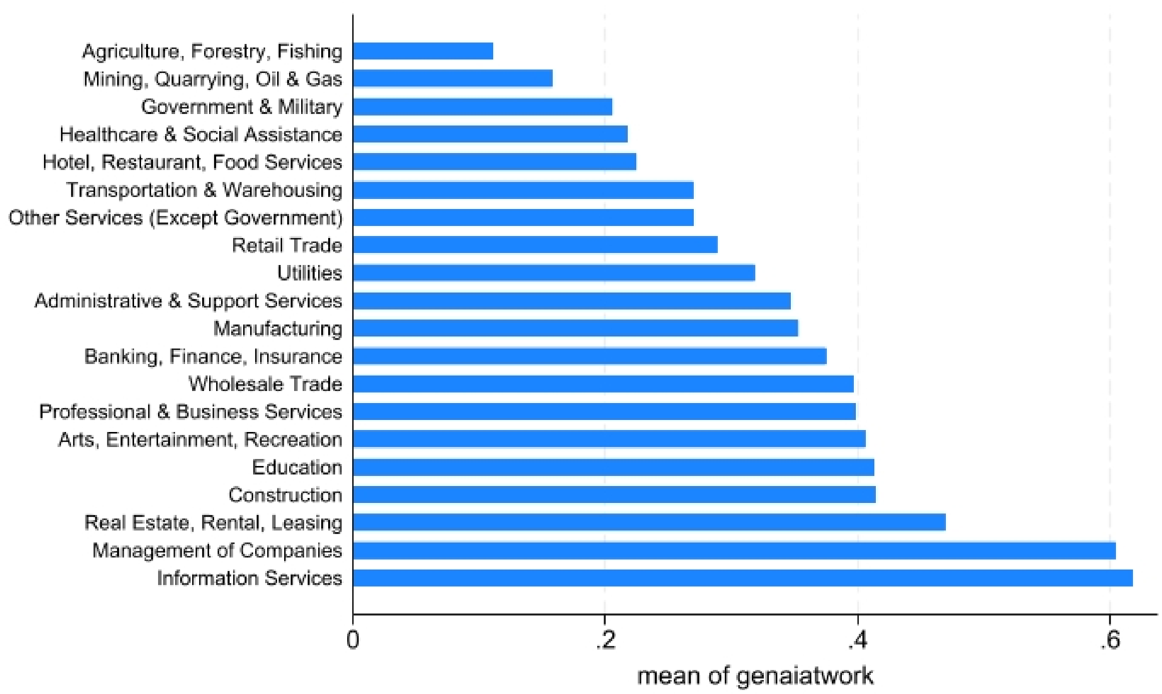

Source: Authors' survey.

Figure A.6: Generative AI Use By Gender

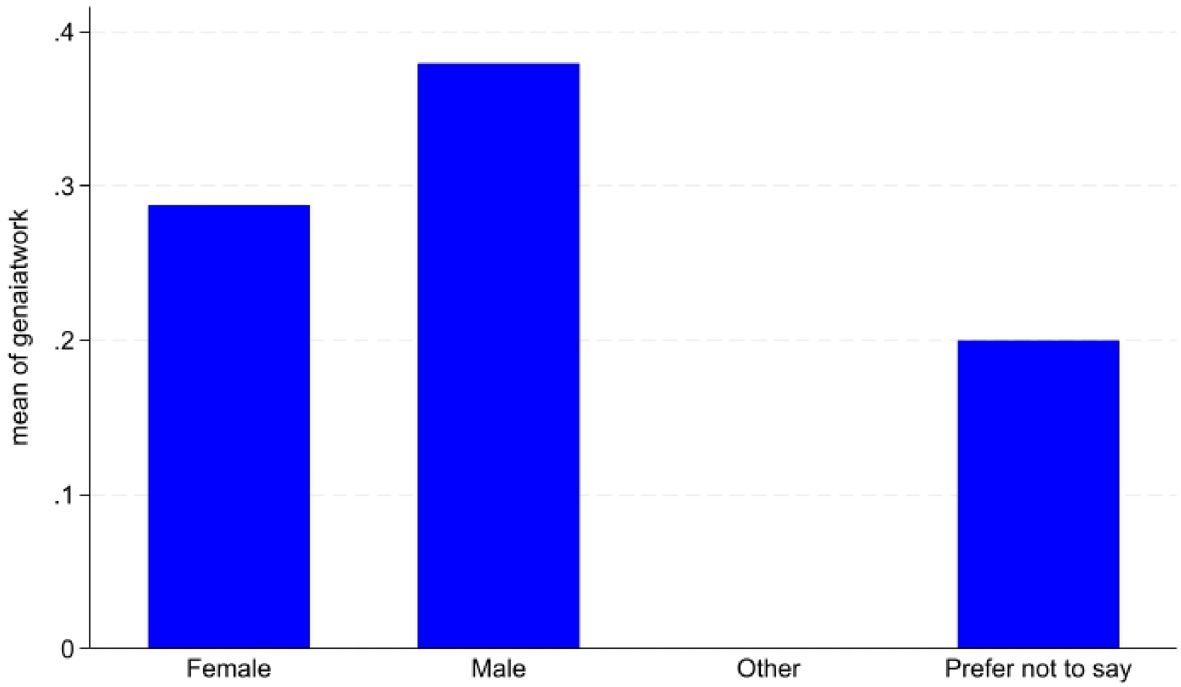

Source: Authors' survey.

Figure A.7: Generative AI Use By Race

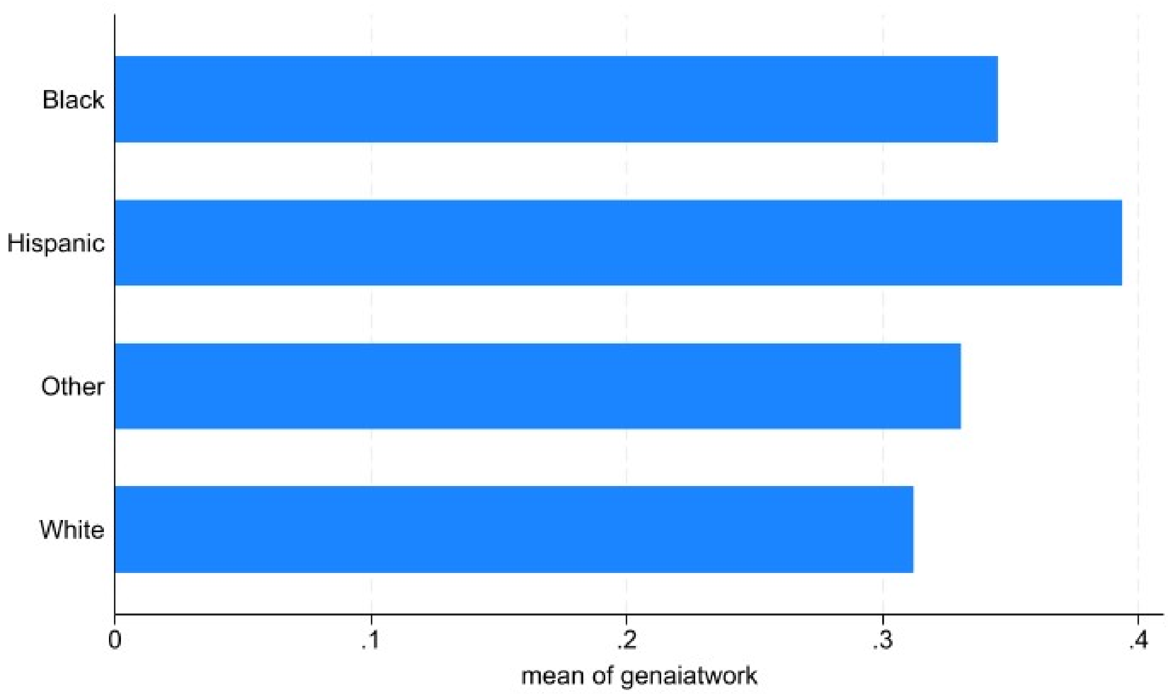

Source: Authors' survey.

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

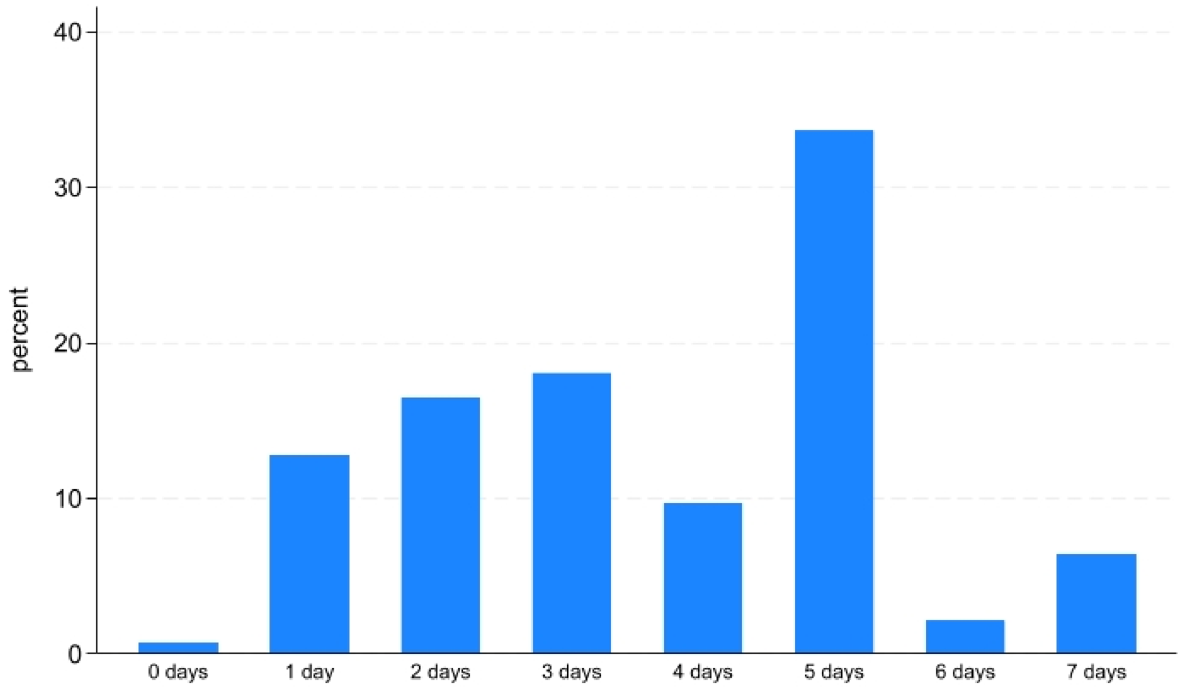

Source: Authors' survey.

Figure A.9: Number of Hours At Work Each Week Using Generative AI Tools

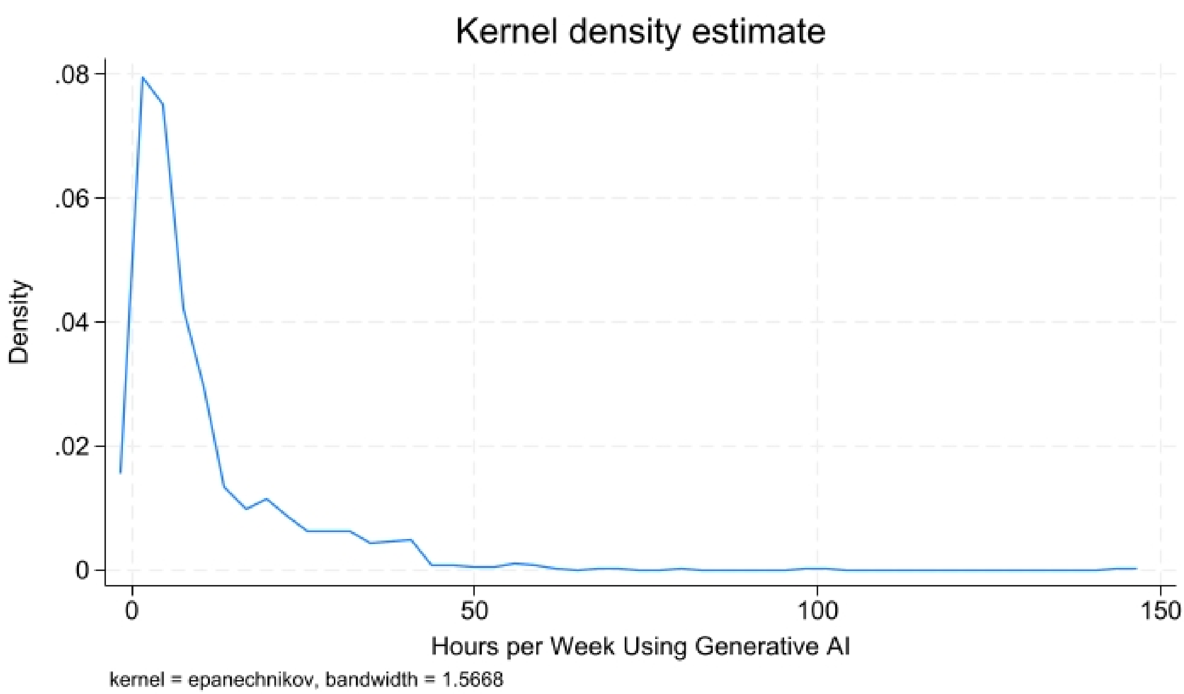

Source: Authors' survey.

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

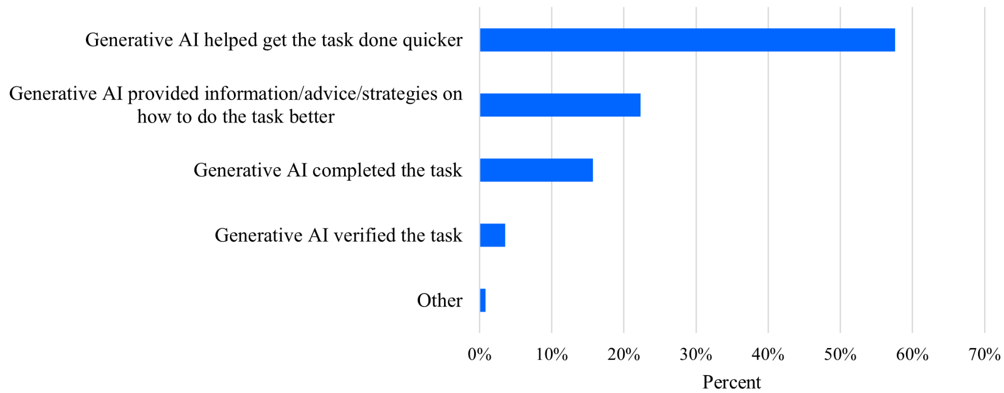

Source: Authors' survey.

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

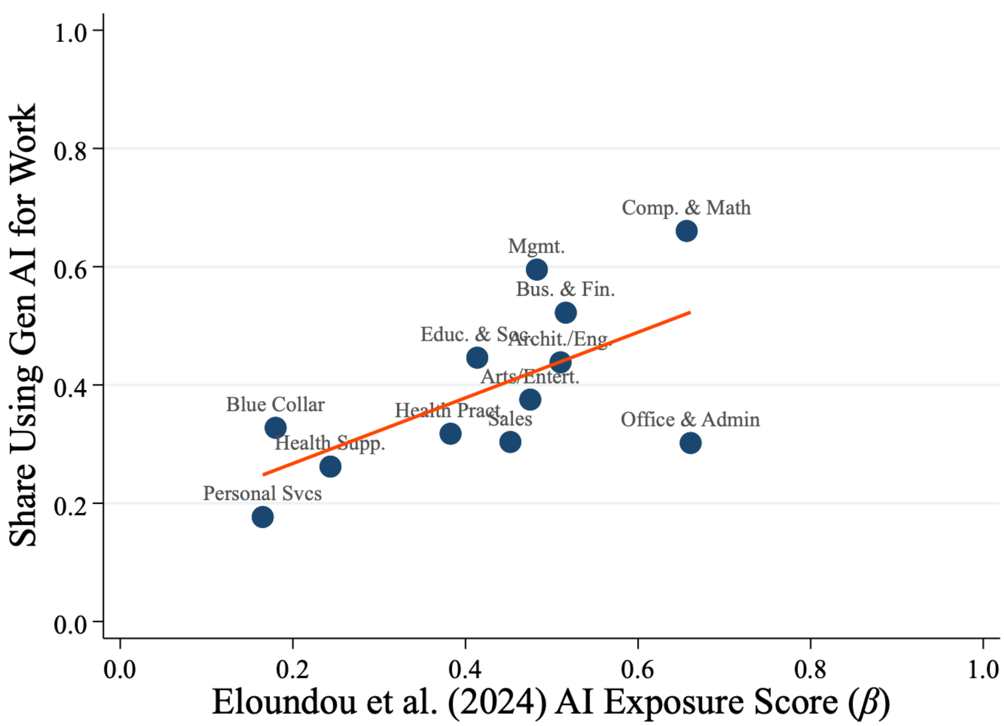

Source: Authors' survey & Eloundou, Manning, Mishkin, and Rock (2024) LLM β-exposure measure. Coefficient: 0.55; Intercept: 0.16.

Figure A.12: Job Loss Perception Due to GenAI During the Next Two Years by Exposure

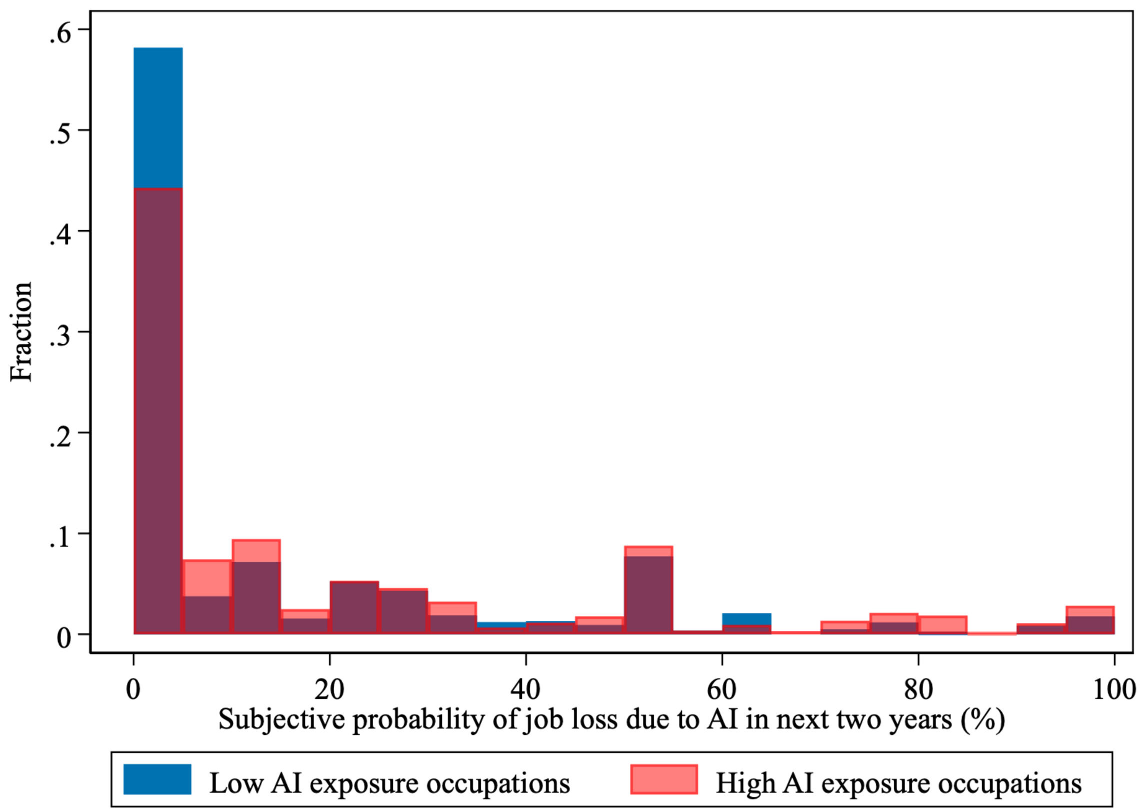

Sources: Authors' survey and β exposure score by Eloundou et al. (2024).

Figure A.13: Effect of Predicted AI Exposure on Job Separations

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

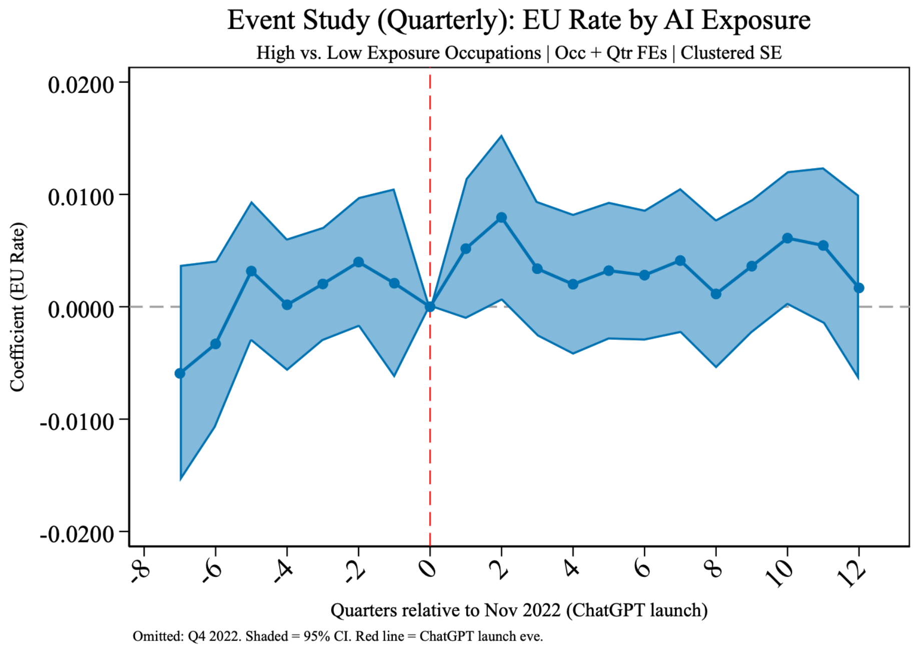

(b) Employment-to-Nonparticipation (*EN*) Transition Rate

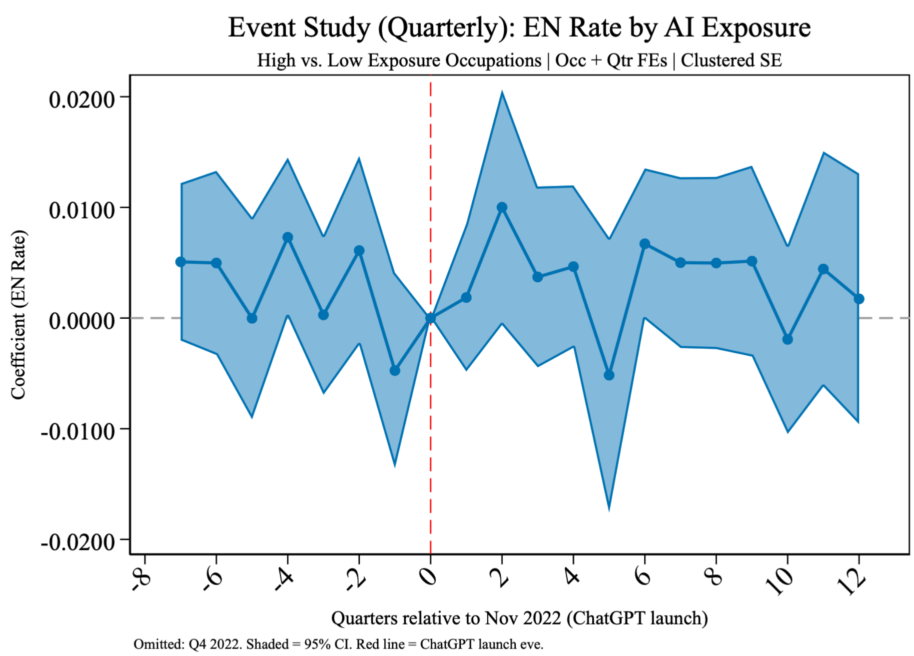

Note: Dependent variable on top panel is *EU* transition rate and on bottom panel is *EN* (NILF) transition rate. Control group are occupations with below median level of employment-weighted predicted AI adoption (0.33). Treatment group includes occupations with above median level of employment-weighted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

Figure A.14: ChatGPT and Gemini Google Search Volume in the United States

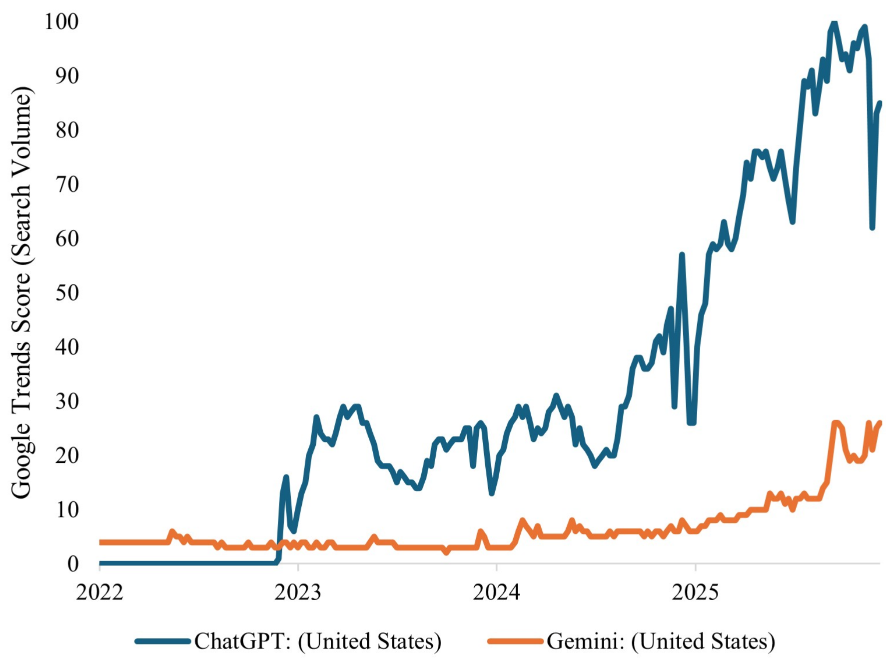

Source: Google Trends.

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)
$\text{Exposure}_o \times Post_t$	-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 × 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 Validating Job Loss Expectations

B.1 Validation Against Outside Measures

We use two other external surveys, implemented by Verasight and Ipsos, to validate our subjective job loss risk measure and assess whether the relationship between AI usage and subjective job loss risk is robust. These surveys elicited job loss probability expectations due to AI with slightly different wording but at the same time (spring 2026) among a similar population (U.S. workers), making them sensible benchmarks for validation.

WE compare the three surveys along several dimensions. First, there are differences in sample frames and recruitment. Our survey uses a Dynata probabilistic sample, whereas Verasight recruits its own verified panel via address-based sampling and dynamic targeting, and the Ipsos KP Omni draws from Ipsos's KnowledgePanel, a long-standing probability-based panel. All three are weighted to benchmarks from the Current Population Survey (see Appendix Table B.1 for comparisons of the survey methodologies).

Second, our survey's sixth wave (where the job loss fear question was asked to all workers) was fielded in March 2026, while Verasight and Ipsos were fielded in May and June, respectively. In general, we don't expect attitudes about AI's job and task substitution capabilities to have shifted dramatically (if at all) among American workers in this narrow window.

Third, and most importantly for this validation exercise, the usage-intensity measures are not constructed identically: our own survey measures weekly hours of GenAI use at work, while Verasight measures general AI use in daily life, and Ipsos measures a narrower and qualitatively different construct (use of an AI agent for companionship or intimate conversation) that is at best a weak proxy for work-related GenAI engagement. Naturally, the Ipsos usage measure is the weakest comparison, but we report it primarily because its job loss risk question is well-matched to our own measure. Exact wording of the questions are presented below:

HJMM Survey.

- **Job-loss fear:** "What do you think is the probability (0-100%) that you will lose your job due to generative AI in the next two years?" (numeric, 0–100; variable `ProbOwnJobLoss`)
- **AI adoption screener:** "Do you use generative AI for your job?"
(Yes/No; variable `UseGenAIForJob`)
- **Usage intensity:** "How many hours per week do you use generative AI tools at work?"
(numeric; variable `HrsPerWeekGenAI`)

Verasight.

- **Job-loss fear (q18):** "How much do you think you are at risk of losing your job due to Artificial Intelligence?"
Response options: *No risk*; *0%–25%*; *25%–50%*; *More than 50%*; *Not sure / I don't know*.
- **Usage intensity (q29):** "How frequently do you use AI in your day to day life?"
Response options: *Never*; *Once a month*; *A few times a month*; *Once a week*; *Everyday / multiple times a day*.

Verasight's usage question asks about general AI use in daily life, not specifically GenAI use at work.

Ipsos KP Omni.

- **Job-loss fear (AI2):** "How much do you think you are at risk of losing your job due to Artificial Intelligence (AI)?"
  Response options: *A great deal*; *A fair amount*; *Only a little*; *Not at all*; *Not currently employed*.
- **AI engagement (AI4):** "There are many uses for online Artificial Intelligence (AI) platforms such as ChatGPT, Claude, and Gemini. Some people use AI to be more productive at work or in school. Others use it to create artwork or music. There are some who create AI agents for companionship or intimate conversation. Have you ever used an online AI agents for companionship or intimate conversation?"
  Response options: *Yes, and I still talk to my AI agent*; *Yes, but I no longer use it for this*; *No, but I use AI for other purposes*; *No, I do not use AI*; *Other*.

Despite slight differences in the wording, all three measures agree qualitatively about the degree of AI job loss fears among U.S. workers in 2026. A comparison is shown in Figure B.1. We map all three surveys onto a common four-category AI job loss risk scale: *No risk*, *1–25%*, *26–50%*, and *>50%*, similar to the Verasight scale for comparability. For our own survey, we bin the continuous job-loss risk variable at the same thresholds. For Ipsos, we map the four qualitative labels onto the same scale by rank order: *Not at all* → No risk (%), *Only a little* → 1–25%, *A fair amount* → 26–50%, *A great deal* → >50%. We report this comparison using Wave 6's full sample, as it is an appropriate comparison to Verasight and Ipsos, given that neither screens on AI adoption. The four categories across the three surveys follow very similar patterns, with similar magnitudes in each bin, giving us confidence in our estimates.

Beyond the assessment that Americans have substantial fears of AI job loss, we also illustrate that all three surveys agree directionally that these fears are positively correlated with greater AI engagement. Figure B.2 depicts this relationship for each of the three surveys, with their own unique scales of fear and measures of AI usage. All three show fear rising with AI usage or engagement. Panel A (our measure) shows the clearest relationship, as both usage and fear are measured on a continuous numerical basis. The Verasight survey (Panel B) uses a five-category of usage frequency that doesn't measure the intensive margin, while the Ipsos Omnibus survey (Panel C) measures the sophistication of AI use, showing that more complex uses of AI, such as agents, are associated with elevated fear of job loss. For our own survey, we restrict to GenAI adopters at work and bin hours of weekly AI use into six categories (0, 1–2, 3–5, 6–10, 11–20, 21+); Panel A plots the mean of job-loss risk by bin with 95% confidence intervals. For Verasight, Panel B plots the weighted mean of the four categories by the five usage-frequency categories. For Ipsos, we restrict to employed respondents (excluding "Not currently employed"). Panel C combines two agent-use categories ("still talk to my AI agent" and "no longer use it for this,") into a single *Ever used AI agent* category, yielding three ordered groups: *Do not use AI*, *Uses AI (no agent)*, and *Ever used AI agent*.

Below we provide further infomration on survey implementation:

Verasight Omnibus Survey. Fielded May 21–26, 2026, to a sample of 1,000 U.S. adults recruited from the Verasight Community (address-based sampling, person-to-person text recruitment, and dynamic online targeting). Respondents are verified via multi-step authentication (SMS from a registered mobile carrier, reCAPTCHA v3) and screened for low-quality response behavior (straight-lining, speeding). Data are weighted to the April 2026 Current Population Survey (CPS) on age, race/ethnicity, sex, income, education, region, and metropolitan status. The margin of sampling error is $\pm$3.1 percentage points at the 95% confidence level.

Ipsos KnowledgePanel (KP) Omnibus. Fielded June 26–28, 2026, to a sample of 1,024 U.S. adults (ages 18+) drawn from KnowledgePanel, a probability-based panel recruited via address-based sampling from the USPS Delivery Sequence File. Of 2,082 panel members fielded, 1,024 completed the survey (a 49% completion rate; all completes qualified). Design weights are adjusted via iterative proportional fitting (raking) to 2025 March CPS supplement benchmarks on gender $\times$ age, race/ethnicity, education, census region, metropolitan status, and household income. The margin of sampling error is $\pm$3.1 percentage points at the 95% confidence level (design effect 1.02).

HJMM Survey (Wave 6). As described in Section 2.1, our survey is fielded quarterly to a probabilistic sample drawn from Dynata, self-administered via CAWI. Wave 6 was fielded in March 2026 to 4,593 respondents. All completed respondents passed a preliminary attention check. AI adopters who answered the longer version of the survey underwent a second attention check. Unlike Waves 2–5, in which the job-loss-risk question was asked only to respondents who screened in as GenAI adopters at work, Wave 6 asked this question of all respondents, which is what permits a general-population comparison to the two external surveys below. Survey weights are constructed from the November 2024 CPS on education (two groups), age (four groups), gender (two groups), and race (four groups), as described in Section 2.1.

Table B.1: Comparison of Survey Methodologies

	Own Survey (Wave 6)	Verasight	IPSOS KP Omni
Field dates	March 2026	May 21–26, 2026	June 26–28, 2026
Sample frame	Dynata	Verasight Community	KnowledgePanel (ABS)
Mode	Online, self-admin.	Online	Online
N (completed)	4,593	1,000	1,024
Weighting benchmark	Nov. 2024 CPS	Apr. 2026 CPS	Mar. 2025 CPS

Note: Table compares methodologies of the HJMM survey from wave 6 against external benchmarks.

Figure B.1: Distribution of Job-Loss Risk Perceptions across Surveys

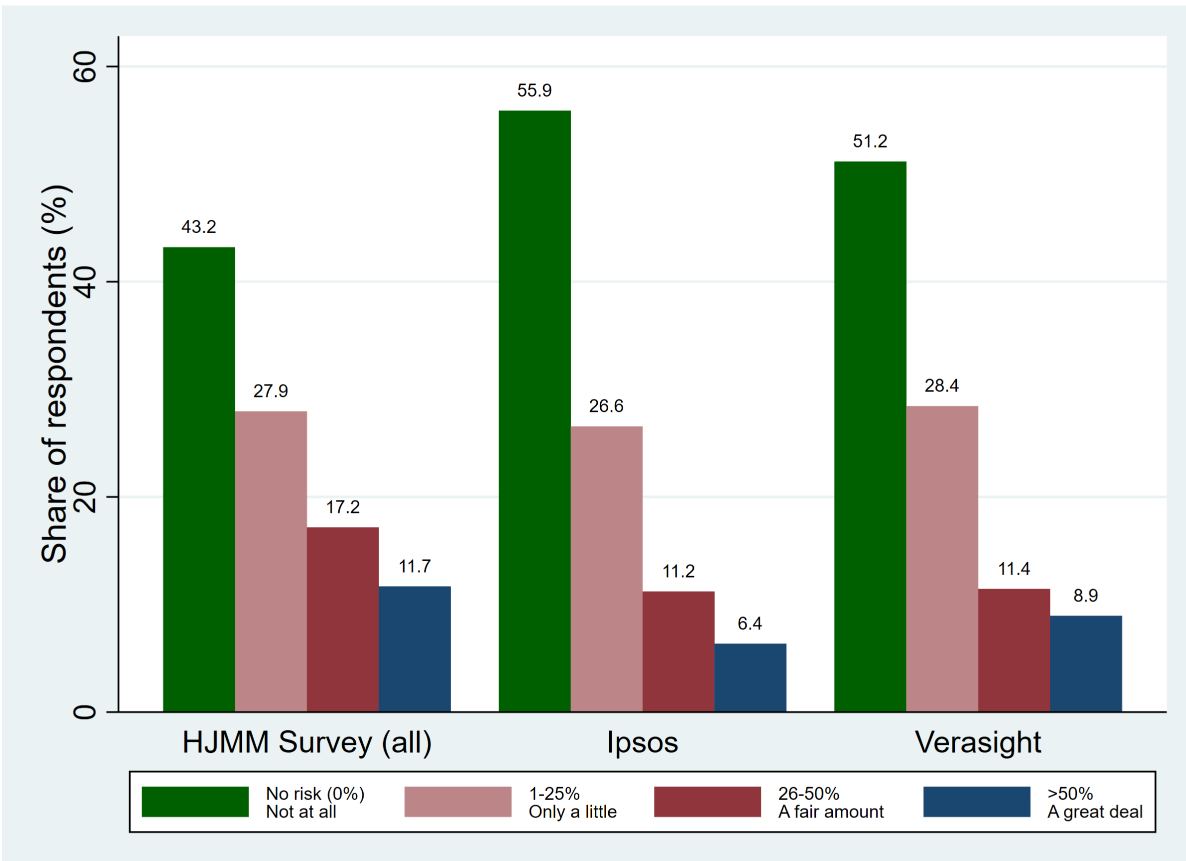

Sources: Authors' own calculations. The categories follow the Verasight scaling. The responses from the HJMM survey are binned in those categories. The categories for the Ipsos survey are presented in the 2nd row of the labels.

Figure B.2: AI Usage and Job-Loss Fear across Surveys

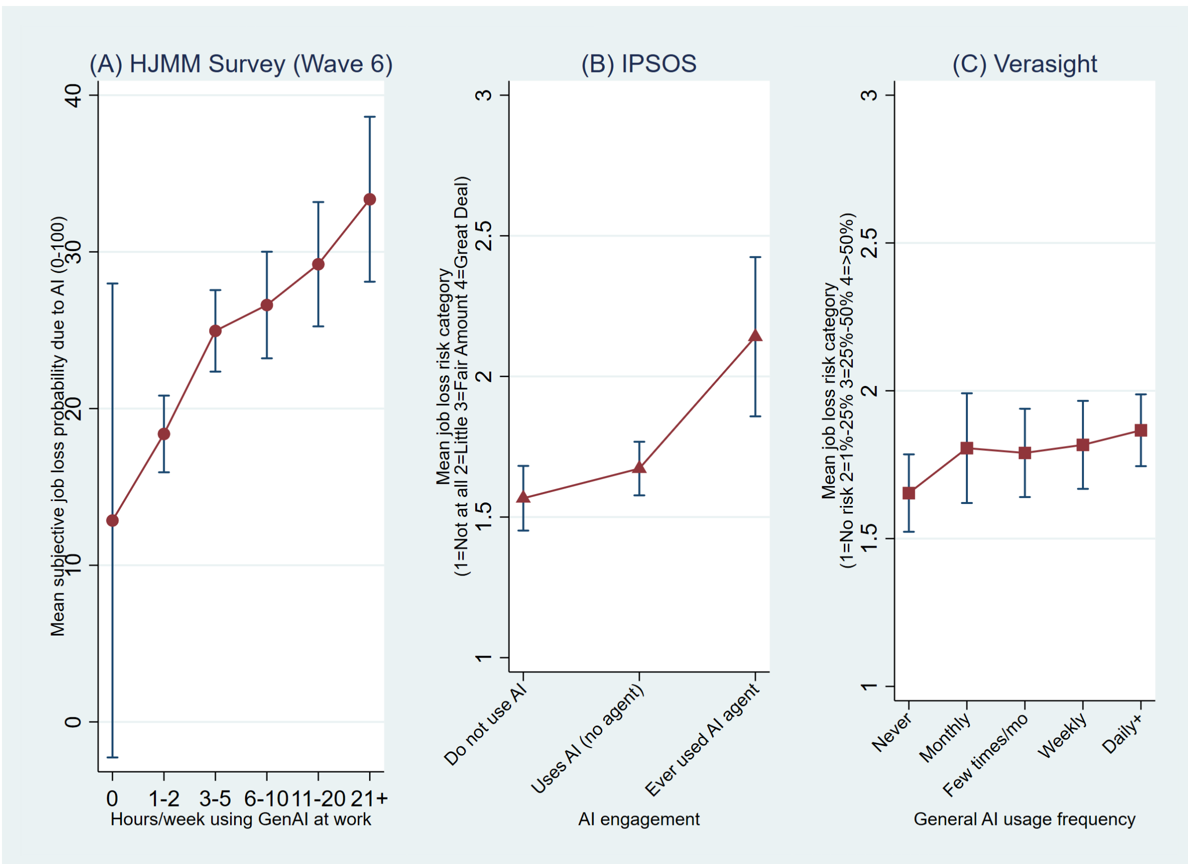

Sources: Authors' own calculations. Error bars show 95% confidence intervals. Panel A shows the HJMM (Wave 6) with the intensive margin of use of AI at work. Panels B and C show coarser, less work-specific use of AI. Job-loss risk scales are not comparable across panels.

B.2 Robustness to Anchoring Job Loss Expectations

The seventh wave of the survey, administers in June 2026, modifies the questions eliciting job loss expectations due to AI by stating the following to survey respondents:

What do you think is the probability (0-100%) that you will lose your job due to generative AI in the next two years? For reference, about 15% of U.S. workers have involuntarily lost their job for any reason over a two-year period in recent years.

Prior to this modification, the survey simply elicited expectations without an anchor for total job loss rates over two years.

Table B.2 shows a t-test that there is no statistically significant difference in the self-reported subjective job loss probabilities due to AI between wave 6 (without anchoring) and wave 7 (with anchoring). The p-value for the difference between the means of only 0.6 percentage points is $p = 0.33$, and the 95% confidence intervals of the means almost entirely overlap.

Lastly, we show that our key relationship – the positive association between AI adopters' job loss fears due to the technology and the intensity of their usage – is robust to anchoring recipients about job loss probabilities. Specifically, Table B.3 conducts statistical test of the *differences* between the positive slopes between job loss fears due to AI and the weekly measure of hours using AI among adopters. All four specifications control for two-digit SOC code. In all four cases (regressions with and without demographic covariates, including and excluding those who failed either attention check), there is no statistically significant difference between the positive slopes for wave 6 and wave 7, with p-values ranging from 0.68 to 0.94.

Table B.2: Two-Sample t-Test with Unequal Variances

Survey Wave	N	Mean	Std. Error	Std. Dev.	95% CI
6	3,934	20.5419	0.4451	27.9172	[19.6693, 21.4146]
7	3,754	19.9387	0.4259	26.0953	[19.1037, 20.7738]
Combined	7,688	20.2474	0.3084	27.0428	[19.6428, 20.8520]
Difference $(6 - 7)$		0.6032	0.6160		[-0.6044, 1.8108]

Notes: Welch's two-sample t-test allowing for unequal variances. The estimated mean difference is 0.6032, with $t = 0.9792$ and Satterthwaite degrees of freedom $= 7682.73$. The two-sided p-value is 0.3275.

Table B.3: Robustness of Job Loss Fear-AI Utilization Gradient to Anchoring

Sample	Specification	Difference in Slopes	Robust SE	F	p-value	N
All	No demog covariates	0.0089	0.1150	0.0060	0.9385	2,911
All	Demog covariates	0.0214	0.1149	0.0348	0.8521	2,911
Attentive	No demog covariates	-0.0506	0.1231	0.1689	0.6811	2,690
Attentive	Demog covariates	-0.0247	0.1230	0.0403	0.8409	2,690

Notes: This table reports the difference in the estimated coefficient on hours-per-week using generative AI between Waves 6 and 7 from regressions with perceived probability of own job loss due to AI on the left hand side. Specifications vary by attentiveness restriction and inclusion of demographic controls. All models include occupation fixed effects and use robust standard errors. Demographic controls are female, non-white demographics, age, and college education. The p-value tests (H_0: $\beta_6 = \beta_7$).

C 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 without 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."