

Working Paper Series
LEWP No. 14

**Searching Without Moving: Generative AI
and Worker Mobility**

Casquet, M.A.
Tercero-Lucas, D.
Victoria-Lanzón, C.

August 26

Available online at: <https://www.aeet.eu/labour-economics-wp-2/>

Searching Without Moving: Generative AI and Worker Mobility

Miguel Ángel Casquet¹, David Tercero-Lucas^{2,3}
and Carlos Victoria-Lanzón ^{*1}

¹Universidad Complutense de Madrid

²Comillas Pontifical University – ICADE, Madrid

³Institute for Research in Technology, ICAI School of Engineering, Comillas Pontifical University, Madrid

August 2026

Abstract

Generative AI is increasingly used at work, but whether it changes workers' opportunities outside their current jobs remains unclear. We use the 2025 Survey of Household Economics and Decisionmaking to study workplace AI use alongside randomized hypothetical wage changes and recent labor-market outcomes. Workers respond strongly to larger wage cuts, but this response does not differ between AI users and non-users. The main difference we find relates to job search behavior: AI users are about 5 percentage points more likely to have applied for another job during the previous year. Furthermore, application rates increase with the frequency of AI use. These results suggest that workers who use generative AI at work search more actively for other jobs, although we do not find evidence that they are more likely to actually change jobs.

JEL classification: J24, J31, J62, J64, O33.

Keywords: Generative artificial intelligence; job search; worker mobility; outside options; wage responsiveness.

1 Introduction

One in four U.S. workers used generative artificial intelligence (AI) at work in 2025. Among users, 81 percent reported that it saved time, and more expected the technology to improve their careers than feared that it would replace their jobs ([Board of Governors](#)

^{*}*Corresponding author.* Email: cvictori@ucm.es. Postal address: Universidad Complutense de Madrid, Campus de Somosaguas, 28223 Pozuelo de Alarcón, Madrid, Spain.

of the Federal Reserve System, 2026a). Such optimism raises a question that existing evidence has left open: how is workplace AI use related to workers’ opportunities outside their current firm?

Research on generative AI has concentrated mainly on analyzing its effects on productivity, task performance, and occupational exposure (Brynjolfsson et al., 2025; Eloundou et al., 2024; Noy and Zhang, 2023). However, workers may also use this technology to identify vacancies, prepare applications, or perform tasks beyond their previous experience. Greater search activity need not imply better outside opportunities: generative AI may make it easier to explore alternative jobs without increasing the probability of receiving or accepting a better offer. Distinguishing job search from realized mobility is therefore important for assessing the labor-market consequences of workplace AI use.

A useful distinction is between changes in productivity in the current job and workers’ opportunities outside the firm. Generative AI may improve performance in the current job without increasing the value of outside options, especially when its benefits depend on firm inputs or organizational practices. At the same time, AI may make job search easier by reducing the cost of finding vacancies, comparing alternatives, or preparing applications. Job search may therefore increase even if realized mobility does not.

We study that distinction using the 2025 Survey of Household Economics and Decisionmaking (SHED). The survey contains a new module on generative AI use at work and a randomized experiment assigning employees to a one-year wage freeze or to a wage reduction of 1, 5, or 10 percent. Respondents then report their likelihood of actively searching for another job or leaving their current position. Random assignment identifies workers’ responses to increasingly adverse wage changes. Comparisons between AI users and non-users remain descriptive because workplace AI use is not randomly assigned. We complement the experiment with information on confidence in finding work, recent job applications, voluntary exits, new job starts, and employer changes.

Our results show that larger hypothetical wage cuts sharply increase stated search

or exit intentions. However, AI users do not react more strongly than non-users. The main difference appears in reported job-search behavior: workers who use generative AI are about 5 percentage points more likely to have applied for another job during the previous year. This result survives a multiple-testing adjustment. Furthermore, application rates increase monotonically with the frequency of AI use. We find much weaker evidence for other outcomes further along the mobility process. For instance, AI users are not systematically more likely to start a new job or change employer. Generative AI use is therefore associated with more job search, but not with a stronger response to wage cuts or greater realized mobility.

The paper contributes in three ways. First, it shifts attention from the effects of generative AI within the current job to workers' behavior and opportunities outside the current employment relationship. Second, it separates active job search from perceived employability and realized mobility, allowing us to locate where differences associated with workplace AI use emerge along the mobility process. Third, it combines observational evidence on workplace AI use with randomized variation in hypothetical wage changes, providing a separate test of whether AI users respond differently when the value of the current match deteriorates.

The distinction between search and realized mobility also has implications for labor-market and skills policy. Generative AI use is diffusing much more broadly across the workforce than explicit employer demand for specialized AI skills. Across 14 OECD countries, vacancies requiring AI-related skills accounted for less than one percent of online job postings between 2019 and 2022 ([Borgonovi et al., 2023](#)), while recent evidence suggests that fewer than one percent of workers will require advanced AI-specific capabilities such as model development or programming. For most workers, AI instead increases the importance of broader digital, data-analytical, managerial, and problem-solving skills ([OECD, 2026](#)). This distinction matters for worker mobility. If generative AI primarily lowers the cost of identifying vacancies, evaluating alternatives, or preparing applications, wider access to the technology may increase job-search activity without necessarily improving the arrival rate or quality of outside offers, or creating

a credential or capability that outside employers explicitly reward. In search-theoretic terms, AI may therefore lower the cost of search before materially changing offer arrival rates or the distribution of outside offers. Policies aimed at translating AI adoption into better labor-market opportunities may consequently need to go beyond promoting AI use itself, by improving information about vacancies and skill requirements, reducing matching frictions, and strengthening the recognition and portability of AI-enabled and complementary skills so that greater search activity can more readily translate into productive job-to-job transitions.

A related policy question concerns the extent to which the capabilities associated with workplace AI use belong to workers rather than to their current firms. AI use at work is already widespread—more than one quarter of European workers report experimenting with AI ([Cedefop, 2025](#))—even though demand for specialized AI skills remains concentrated in a relatively narrow set of occupations and sectors ([Borgonovi et al., 2023](#)). At the same time, workplace adoption may reflect employer-provided tools, training, and organizational deployment rather than the accumulation of fully portable human capital. Employer provision appears to be an important source of variation in worker AI use ([Kim et al., 2026](#)), while theoretical work distinguishes personal AI tools, which may enhance worker portability, from enterprise systems whose value is co-specialized with proprietary data, organizational routines, and firm-specific infrastructure ([Call et al., 2026](#)). Consistent with this distinction, employer-funded training is already an important channel through which workers acquire AI-related capabilities ([OECD, 2026](#)), and international organizations increasingly emphasize training and social dialogue as central to the governance of workplace AI ([OECD, 2026](#); [Doellgast et al., 2025](#)). The policy challenge is therefore not simply to increase AI adoption, but to ensure that the capabilities workers acquire through AI remain visible and transferable across employers rather than raising productivity only within the current match. This may require complementing AI training with recognition of prior learning, recognized credentials, career guidance, and better information on emerging skill requirements ([International Labour Organization, 2026](#)).

The remainder of the paper is organized as follows. Section 2 discusses the related literature and presents the main hypotheses. Section 3 describes the SHED data, the generative AI module, and the randomized wage scenarios. Section 4 presents the empirical strategy. Section 5 reports the main results, examines heterogeneity in job-search activity by AI use intensity, and presents the robustness exercises. Section 6 concludes.

2 Literature and hypotheses

This paper contributes to the emerging literature on the labor-market consequences of generative AI. A first strand of this literature documents the rapid diffusion of generative AI at work and emphasizes the distinction between occupational exposure and actual use. Adoption is already widespread, but varies substantially across demographic groups, occupations, and workers facing similar technological exposure (Bick et al., 2026; Humlum and Vestergaard, 2025b). This distinction is central for our analysis, which relies on workers' reported use of generative AI rather than on occupation-level measures of potential exposure.

A growing body of experimental and quasi-experimental evidence shows that access to generative AI can raise productivity, improve task performance, and accelerate learning. These gains have been documented in professional writing, software development, and customer support, and in some settings are particularly large among less experienced workers (Brynjolfsson et al., 2025; Gambacorta et al., 2026; Noy and Zhang, 2023). This evidence establishes that generative AI can alter productivity within the current job. It is less clear, however, whether such gains improve workers' opportunities outside the current employment relationship.

Labor-market adjustment to AI may also occur through changes in tasks, occupations, and mobility rather than immediately through earnings or hours. Recent worker-level evidence shows that AI exposure reshapes task content primarily within occupations and induces worker reallocation across jobs, often toward less AI-exposed

industries, while displacement effects remain limited (Gathmann et al., 2025). College-educated workers exposed to AI appear to reallocate from occupations with high AI exposure and low complementarity toward high-exposure occupations in which AI is more complementary to labor (Cazzaniga et al., 2025). More directly, evidence linking actual AI adoption to administrative labor-market outcomes suggests substantial task reorganization and some movement by adopters toward higher-paying occupations in which AI chatbots are more relevant, despite precise null effects on average earnings and recorded hours (Humlum and Vestergaard, 2025a). These findings make worker mobility a natural margin through which the consequences of generative AI may emerge, while also suggesting that different stages of labor-market adjustment need not occur simultaneously.

Job search provides an especially relevant earlier stage of this adjustment process. Generative AI is already used to identify vacancies and prepare application materials, indicating that the technology has entered the production function of job search itself (Hartley et al., 2026). Our analysis moves from the use of AI for job search to a different question: whether workers who use generative AI at work are systematically more likely to explore outside employment opportunities, and whether greater exploration is accompanied by stronger perceived outside options or completed job transitions. This distinction allows us to separate search activity from the realization of outside opportunities through actual mobility.

Adoption may also matter at both the extensive and intensive margins. Worker responses to generative AI vary with the frequency of use, with heavier users displaying stronger changes in career-related learning (Kim et al., 2026). We therefore complement the user–non-user comparison with an intensive-margin analysis, asking whether the association between workplace AI use and job search becomes stronger as generative AI is used more frequently.

A related literature points to the potential portability of AI-related human capital across employers. Employer demand for AI skills has grown rapidly, and vacancies requiring these skills command substantial wage premia (Alekseeva et al., 2021). More

generally, AI-related skills appear particularly valuable when combined with complementary capabilities ([Stephany and Teutloff, 2024](#)). These findings concern specialized AI skills rather than the use of broadly accessible generative AI tools, but they highlight the possibility that AI capabilities may have value beyond the current firm. Whether routine use of generative AI creates portable capabilities is ultimately an empirical question.

Our second point of departure is the literature on outside options, job search, and employer wage-setting power. In frictional labor markets, workers' outside options depend on both the opportunities actually available to them and their beliefs about those opportunities. The set and value of alternative jobs affect wages and job-to-job mobility, while imperfect information about outside wages can shape search and bargaining behavior; separations also respond causally to changes in relative pay ([Bagga, 2023](#); [Dube et al., 2019](#); [Jäger et al., 2024](#)). Outside options can therefore manifest themselves at several distinct stages: perceived employability, active job search, voluntary separation, realized mobility, and responsiveness to deterioration in the current match.

Generative AI can, in principle, affect these margins through different channels. Recent work distinguishes portable, worker-controlled AI capabilities from firm-specific complementarities embedded in enterprise systems ([Call et al., 2026](#)). When workers retain control over the tools and capabilities they develop, AI use may reduce dependence on firm resources and strengthen opportunities elsewhere. When productivity gains instead rely on proprietary data, organizational routines, or firm systems, AI may raise productivity in the current match without increasing the value of the worker to outside employers. This distinction motivates our separation between the technology of job search, the strength of realized outside options, and workers' responsiveness to adverse wage changes.

These mechanisms generate distinct predictions across different stages of the mobility process. First, generative AI may affect the technology of job search without necessarily improving the opportunities available to workers. By lowering the cost of identifying vacancies, evaluating alternatives, or preparing application materials, AI

use may increase external job-search activity even if the arrival rate or value of outside offers remains unchanged. Second, if the capabilities associated with generative AI are portable across firms and improve workers' external employability, AI users should have stronger outside options. This should be reflected in greater confidence in obtaining alternative employment and, ultimately, in a higher probability of voluntary separation, starting a new job, or changing employer. By contrast, if the productivity gains from AI are largely firm-specific, or if adoption has not yet translated into more valuable external opportunities, differences in realized mobility need not emerge. Finally, stronger outside options should affect workers' responses to changes in the value of their current match. If workplace AI use identifies workers with more valuable external opportunities, these workers should respond more strongly when current compensation deteriorates. The randomized wage scenarios allow us to test this prediction separately from the observational associations between AI use and labor-market behavior.

We test these predictions using the 2025 Survey of Household Economics and Decisionmaking (SHED), which introduced new questions on the use of generative AI. The SHED is particularly well suited to our research question because it measures workers' actual generative AI usage, rather than assigning exposure based on occupation, and independently randomizes the wage scenario presented to each employed respondent. Our design therefore does not identify the causal effect of adopting generative AI. Instead, it tests whether workers who use generative AI at work respond differently to adverse wage changes than non-users. To the best of our knowledge, existing research has not directly examined this link between individual generative AI adoption and wage responsiveness.

These arguments lead to three empirical hypotheses:

H1: Job search. Workers who use generative AI at work should engage more actively in external job search if the technology lowers the cost of exploring alternative employment opportunities.

H2: External employment opportunities. If generative AI improves workers' ex-

ternal employment opportunities, users should report greater perceived employability and exhibit greater realized labor-market mobility.

H3: Wage responsiveness. If generative AI users have stronger outside options, they should respond more strongly to adverse wage scenarios by becoming more likely to search for another job or leave their current position.

3 Data and descriptive statistics

3.1 Sample

We use the public-use microdata from the 2025 Survey of Household Economics and Decisionmaking (SHED), administered by the Board of Governors of the Federal Reserve System. The survey was fielded between October 17 and October 28, 2025, through Ipsos’s KnowledgePanel, a probability-based online panel recruited using address-based sampling. The public-use file contains 12,934 U.S. adults. Post-stratification weights account for differential non-response and coverage and align the full survey sample with the U.S. adult population along demographic, geographic, educational, and income dimensions ([Board of Governors of the Federal Reserve System, 2026a](#)).¹

The 2025 SHED introduced a module on generative AI use at work. Respondents were asked whether they had used generative AI at their job during the previous month. Our main measure, AI_i , equals one for workers reporting such use and zero otherwise. The module was administered to 7,488 respondents, of whom 6,285 reported working for someone else. A small number of respondents are classified as employees in the relevant survey question but as self-employed or working in a family business in the employer-type classification. We exclude those cases from the main analytical sample, leaving 6,025 workers.² The survey also records frequency of use among AI users, ranging from less than once per week to multiple times per day ([Board of Governors of the Federal](#)

¹We use the main cross-sectional survey weights for descriptive statistics and weighted regressions.

²Retaining these observations produces similar estimates. We report the corresponding results as a robustness check.

[Reserve System, 2026b](#)).

Our main experimental outcome comes from a question administered to respondents working for an employer. Before answering the question, each employee is randomly assigned to one of four scenarios: the employer keeps the worker’s nominal pay unchanged for one year or reduces it by 1, 5, or 10 percent. Respondents then report whether they would be *very likely*, *somewhat likely*, *not that likely*, or *not at all likely* to actively look for another job or leave their current position. We compare the wage-response profiles of generative AI users and non-users, defining our primary outcome as reporting either *very likely* or *somewhat likely* and using *very likely* alone as a robustness check. We complement the experimental outcome with a measure of perceived employability. Respondents are asked how confident they are that, if they wanted to, they could obtain a job within the next three months; respondents who are currently working are explicitly asked about obtaining a *new* job. The available responses are *not confident*, *somewhat confident*, and *very confident*. In the full SHED sample, 26 percent of respondents report being not confident, 35 percent somewhat confident, and 24 percent very confident; the remaining respondents are not working and report no current interest in working ([Board of Governors of the Federal Reserve System, 2026b](#)). Because our analysis conditions on current employment, the question provides a direct subjective measure of perceived employability outside the current job.

The survey also provides complementary measures of labor-market behavior over the previous 12 months. We use three indicators reporting whether the respondent applied for a new job, started a new job, or voluntarily left a job during that period. The SHED reports that these behaviors are relatively common: 23 percent of respondents applied for another job, 13 percent started a new job, and 8 percent voluntarily left a job ([Board of Governors of the Federal Reserve System, 2026b](#)). We code each item as a separate binary outcome. We also construct an indicator for changing employer, equal to one for respondents who started a new job and report that their current main job is with a different employer. Applications and voluntary separations capture the search and separation margins, whereas new-job starts and employer changes capture

realized mobility.

These measures play complementary roles in the empirical analysis. The experimental outcome captures how workers respond to an exogenously assigned deterioration in pay, allowing us to compare wage responses across AI users and non-users. Confidence in finding a job provides a measure of perceived employability, while applications, voluntary exits, new-job starts, and employer changes describe recent labor-market behavior. Generative AI use is not randomly assigned and refers to the previous month,³ whereas these outcomes cover the previous 12 months. We therefore treat the corresponding estimates as descriptive associations rather than causal effects.

The main analytical sample contains 6,025 workers with valid information on recent generative AI use and the main experimental outcome. Analyses of the additional labor-market outcomes use the same sample whenever the corresponding outcome is available; the confidence measure is available for 5,965 respondents. The SHED also provides demographic, socioeconomic, occupational, and job-level characteristics. Base-line specifications control for age, gender, race and ethnicity, education, household income, occupation, and industry. Employment characteristics are added in extended specifications. Appendix Tables B1 and B2 define the variables used in the empirical analysis. The former covers generative AI use, the randomized wage scenario, and the labor-market outcomes. The latter reports the demographic, human-capital, and employment controls.

3.2 Descriptive statistics

Table 1 reports descriptive statistics for the analytical sample and compares AI users with non-users. Generative AI users account for 26.0 percent of the weighted sample. The two groups are similar in gender composition and differ only modestly in age,

³Because generative AI use is measured over the month preceding the survey, the indicator may partly reflect conditions specific to that period. A common aggregate shock would not by itself generate differences between AI users and non-users, although period-specific shocks that affected the two groups differently cannot be ruled out. Our specifications condition on occupation, industry, and a broad set of worker characteristics, and the main association with job applications is also robust to entropy balancing.

although AI users are somewhat less likely to be Black or Hispanic. Differences are considerably larger in education and household income. Among AI users, 69.5 percent have a bachelor’s degree or higher, compared with 39.5 percent of non-users. The corresponding shares living in households with annual income of at least \$100,000 are 78.4 and 55.9 percent. We account for these differences in the adjusted specifications reported below.

Table 1: Descriptive statistics and mean differences by generative AI use

Variable	Full sample			By AI use		Difference
	N	Mean	SD	Non-users	AI users	
<i>Panel A. Main variables</i>						
AI user	6,025	0.260	0.439			
Wage freeze	6,025	0.249	0.433	0.247	0.257	0.010
1% wage cut	6,025	0.254	0.436	0.255	0.252	-0.003
5% wage cut	6,025	0.249	0.432	0.255	0.230	-0.025*
10% wage cut	6,025	0.248	0.432	0.243	0.261	0.018
Likely to search or leave	6,025	0.548	0.498	0.545	0.555	0.010
Very likely to search or leave	6,025	0.275	0.446	0.275	0.272	-0.003
Confident of finding a job	5,965	0.720	0.449	0.723	0.712	-0.011
Applied for another job	6,025	0.273	0.445	0.260	0.309	0.050***
Left a job voluntarily	6,025	0.093	0.290	0.090	0.101	0.011
Started a new job	6,025	0.178	0.383	0.178	0.180	0.002
Changed employer	6,025	0.102	0.303	0.101	0.106	0.005
<i>Panel B. Demographic characteristics</i>						
Age	6,025	42.951	14.262	43.174	42.317	-0.858**
Female	6,025	0.478	0.500	0.480	0.474	-0.006
White, non-Hispanic	6,025	0.606	0.489	0.603	0.615	0.013
Black, non-Hispanic	6,025	0.113	0.317	0.121	0.091	-0.031***
Hispanic	6,025	0.175	0.380	0.183	0.151	-0.032***
<i>Panel C. Human-capital characteristics</i>						
Bachelor's degree or higher	6,025	0.474	0.499	0.395	0.695	0.300***
Household income $\geq$ \$100,000	6,025	0.618	0.486	0.559	0.784	0.225***

Notes: The table reports survey-weighted descriptive statistics and mean differences by generative AI use. Differences are defined as AI users minus non-users. Statistical significance is based on two-sided tests of equality of means obtained from survey-weighted bivariate regressions with heteroskedasticity-robust standard errors. Female, the race/ethnicity categories, bachelor’s degree or higher, and household income of at least \$100,000 are binary indicators. Occupation, industry, and employment characteristics are omitted for brevity. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Source: Authors’ elaboration.

4 Empirical Strategy

Our empirical analysis has two parts. The first examines whether AI users respond differently to an exogenous deterioration in current pay. The second compares AI users and non-users across perceived employability, job search, and realized mobility. The randomized wage scenarios identify the effect of hypothetical changes in compensation on stated search or exit intentions. Comparisons across AI-use status remain conditional associations because workplace generative AI use is not randomly assigned.

4.1 Randomized wage scenarios

Respondents are randomly assigned to one of four hypothetical scenarios: a one-year wage freeze or a wage reduction of 1, 5, or 10 percent. After seeing the assigned scenario, respondents report how likely they would be to actively search for another job or leave their current position. Our main dependent variable equals one when the respondent reports being very or somewhat likely to search or leave, and zero otherwise.

We estimate the following linear probability model:

$$Y_i = \alpha + \sum_{s=1}^3 \beta_s W_{is} + \gamma AI_i + \sum_{s=1}^3 \delta_s (W_{is} \times AI_i) + X_i' \theta + \varepsilon_i, \quad (1)$$

where Y_i denotes the binary search-or-leave outcome, W_{is} identifies the three wage-cut scenarios, and the wage-freeze group is the omitted category. AI_i equals one when respondent i reports using generative AI at work. The vector X_i contains age, age squared, gender, race and ethnicity, education, household income, occupation, and industry. The full specification additionally controls for working arrangements, job autonomy, union status, and the desire to work additional hours.

The coefficients β_s measure the effect of each wage cut among non-users of generative AI. The coefficient γ captures the difference between AI users and non-users under the wage-freeze scenario. The interaction coefficients δ_s test whether AI users respond differently to each wage reduction. Random assignment identifies the effects of the wage scenarios and their heterogeneity across AI-use groups. It does not identify the

causal effect of adopting generative AI.

4.2 Associational analysis of labor-market outcomes

The experiment measures stated responses to changes in current pay, but it does not show whether AI users differ along other stages of the mobility process. We therefore estimate associations between workplace AI use and five labor-market outcomes: confidence in finding a job, applying for another job, leaving a job voluntarily, starting a new job, and changing employer.

For each outcome k , we estimate:

$$L_{ik} = \alpha_k + \rho_k AI_i + X_i' \theta_k + u_{ik}, \quad (2)$$

where L_{ik} denotes the corresponding binary outcome. The coefficient ρ_k measures the conditional difference between AI users and non-users. The control vector includes age, age squared, gender, race and ethnicity, education, household income, occupation, and industry.⁴ All baseline specifications are estimated as linear probability models using the survey weights provided by the SHED. Heteroskedasticity-robust standard errors are reported throughout. Robustness exercises replace the linear probability model with a logit model, remove the survey weights, impose a common sample across outcomes, and add employment controls. We also apply the Holm procedure to the five baseline coefficients on AI use to account for multiple hypothesis testing.


5 Results

5.1 Responses to randomized wage changes

We first exploit the randomized wage scenarios to examine whether generative AI use is associated with a stronger response to adverse changes in pay. Respondents were

⁴Employment characteristics are added only as a robustness check because they may be jointly determined with recent job-search and mobility outcomes.

assigned to a one-year wage freeze or to a wage cut of 1, 5, or 10 percent. They then reported their likelihood of actively searching for another job or leaving their current position. Figure 1 presents the survey-weighted probabilities separately for AI users and non-users. Random assignment identifies the effect of the wage scenario on stated search or exit intentions.⁵



Notes: The figure reports survey-weighted percentages of respondents who state that they would be very or somewhat likely to actively search for another job or leave their current position. Error bars represent 95% confidence intervals. The estimation sample contains 6,025 observations.

Search or exit intentions increase monotonically with the size of the wage cut. The response to a 1 percent cut relative to a wage freeze is modest, whereas the 5 and 10 percent cuts produce substantially larger changes. The point estimates for AI users and non-users are very similar across the four scenarios, with small gaps that show no consistent pattern.

Table 2 quantifies these patterns. Relative to the wage-freeze scenario, a 1 percent wage cut increases the probability of being very or somewhat likely to search for another

⁵Because AI use is not randomly assigned, comparisons between AI users and non-users remain descriptive.

Table 2: Responses to randomized wage changes by generative AI use

	(1) No controls	(2) Demographics	(3) Human capital and sector	(4) Full controls
1% wage cut	0.065*** (0.022)	0.065*** (0.022)	0.067*** (0.021)	0.066*** (0.021)
5% wage cut	0.262*** (0.022)	0.262*** (0.021)	0.265*** (0.021)	0.269*** (0.021)
10% wage cut	0.346*** (0.021)	0.348*** (0.020)	0.350*** (0.020)	0.350*** (0.020)
AI user	0.021 (0.030)	0.004 (0.030)	0.019 (0.030)	0.022 (0.030)
AI user $\times$ 1% cut	-0.024 (0.043)	-0.017 (0.042)	-0.015 (0.042)	-0.014 (0.041)
AI user $\times$ 5% cut	-0.012 (0.043)	-0.012 (0.042)	-0.009 (0.042)	-0.008 (0.042)
AI user $\times$ 10% cut	-0.007 (0.040)	0.001 (0.039)	0.008 (0.039)	0.011 (0.039)
Observations	6,025	6,025	6,025	6,025
R-squared	0.080	0.126	0.157	0.176

Notes: The dependent variable equals one when the respondent reports being very or somewhat likely to actively search for another job or leave their current position. The omitted wage scenario is a one-year wage freeze. All specifications include a constant and use survey weights. Heteroskedasticity-robust standard errors are reported in parentheses. Demographic controls include age, age squared, gender, and race or ethnicity. Human capital and sector controls include education, household income, occupation, and industry. Employment controls include working arrangements, job autonomy, union status, and the desire to work additional hours. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

job or leave the current position by about 6 percentage points among non-users, while the corresponding increases for 5 and 10 percent cuts are approximately 27 and 35 percentage points. The coefficient on AI use under the wage-freeze scenario is small, and none of the interactions between AI use and the wage cut indicators is statistically significant.⁶

A plausible interpretation separates productivity in the current match from the value of an outside option. Standard models of on-the-job search link mobility decisions to offer arrival rates, expected wage offers, and search costs rather than to productivity

⁶A joint test fails to reject the null hypothesis that the three interaction coefficients are jointly equal to zero ($F(3, 5934) = 0.16$, $p = 0.924$), providing no evidence of differential responses to the randomized wage reductions by generative AI use.

alone (Faberman et al., 2022; Mortensen, 1986). Evidence on generative AI points mainly to productivity gains within existing jobs, with larger effects among less experienced workers (Brynjolfsson et al., 2025). Such gains need not improve the value of outside employment opportunities. Beliefs and mobility frictions provide a second margin. Workers often anchor expectations about outside pay on their current wage, and employer changes involve substantial switching costs (Caldwell et al., 2025; Jäger et al., 2024). AI use may therefore be associated with productivity gains in the current job without implying stronger outside employment opportunities. The absence of differential responses to the randomized wage cuts is consistent with that distinction.

5.2 Job search and labor market transitions

The randomized wage scenarios test whether workers’ responses to adverse wage scenarios differ by AI-use status. A separate implication concerns the search process itself. Generative AI may reduce the cost of gathering information about vacancies, preparing applications, and evaluating alternative jobs without increasing offer arrival rates or improving the quality of available matches. We therefore examine distinct margins of the mobility process: perceived employability, active search, voluntary separation, and realized employment transitions. The analysis allows us to assess whether any association with AI use is concentrated at the search margin or extends to completed mobility.⁷

Table 3 reports weighted linear probability models with demographic, human-capital, and sector controls. The sequence of outcomes, from column (1) to column (5), allows us to distinguish perceived employability and active search from realized mobility. The estimates remain associational because generative AI use is not randomly assigned.

AI users are 5 percentage points more likely to have applied for another job during the previous 12 months. The difference is sizeable relative to the 26 percent application rate among non-users. AI use is also significantly associated with a 2.2 percentage point

⁷We report the full set of available outcomes covering perceived employability, job search, separation, and realized mobility, including outcomes for which the estimated association is null.

Table 3: Generative AI use and labor market outcomes

	(1) Confident	(2) Applied	(3) Left voluntarily	(4) Started job	(5) Changed employer
AI user	0.027* (0.015)	0.050*** (0.015)	0.022** (0.010)	0.018 (0.013)	0.014 (0.010)
Mean among non-users	0.723	0.260	0.090	0.178	0.101
Observations	5,965	6,025	6,025	6,025	6,025
R-squared	0.059	0.115	0.061	0.118	0.058

Notes: The table reports weighted linear probability models. The dependent variables indicate whether the respondent is confident of finding a job within the next three months, applied for a new job, voluntarily left a job, started a new job, or changed employer during the previous 12 months. The confidence outcome excludes respondents who are not working and are not interested in working. Changed employer equals one when the respondent started a new job and reports that the current main job is with a different employer. All specifications control for age, age squared, gender, race and ethnicity, education, household income, occupation, and industry. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

increase in voluntary exits. Confidence in finding a job is 2.7 percentage points higher among AI users, although the estimate is significant only at the 10 percent level. AI use is not significantly associated with starting a new job or changing employer. The baseline estimates therefore point to larger differences at the search margin than in realized mobility.

Because the analysis considers five related outcomes, Table C1 in the appendix reports p -values adjusted using the Holm method (Holm, 1979). The association with job applications remains significant at the 1 percent level after the adjustment. The adjusted p -value for voluntary exits is 0.107, while confidence, new job starts, and employer changes are not statistically significant. The multiple-testing adjustment reinforces the conclusion that the strongest evidence concerns job-search activity rather than completed mobility. Across the five outcomes, the clearest difference between AI users and non-users appears at the job-search margin. Although the baseline estimates provide some evidence of differences in perceived employability and voluntary separations, only the association with job applications remains statistically significant after adjusting for multiple testing. We find no robust difference in new-job starts or employer changes. The evidence therefore points to greater exploration of outside employment opportunities rather than systematically greater realized mobility.

The pattern is consistent with generative AI being more closely associated with the search margin than with the value or availability of outside offers. Workers can use generative AI tools to identify vacancies, prepare application materials, and reduce the time cost of exploring alternative employment, uses that are already documented among job seekers by [Hartley et al. \(2026\)](#). Completing a job transition, however, additionally requires an outside opportunity to arrive and to be sufficiently attractive to accept. Our results therefore suggest that the clearest labor-market correlate of workplace generative AI use is concentrated at the search margin of the mobility process.

5.3 Job-search activity and the intensity of generative AI use

We next examine whether the association between generative AI use and job applications also appears along the intensive margin of adoption. Because applications are the only labor-market outcome that remains statistically significant after the multiple-testing adjustment, we treat this exercise as complementary evidence rather than as an independent test of the baseline result. Frequency of use provides a natural dimension along which to examine this association, as recent evidence also documents stronger career-related responses among heavier users of generative AI ([Kim et al., 2026](#)). Appendix Table [D1](#) reports the distribution of generative AI use frequency among AI users.

To study heterogeneity along the intensive margin, we estimate

$$Y_i = \alpha + \sum_{k=2}^5 \beta_k \mathbf{1}\{F_i = k\} + X_i' \theta + \varepsilon_i, \quad (3)$$


where Y_i is an indicator for having applied for another job during the previous 12 months and F_i denotes the frequency of generative AI use at work.⁸ The five frequency categories are ordered from lowest to highest intensity. Figure [2](#) reports adjusted application probabilities for all five categories. This normalization does not restrict the subsequent predictions: adjusted probabilities are recovered for all five frequency cate-

⁸The estimation sample is restricted to respondents who report using generative AI at work.

gories. The vector X_i contains the same demographic, human-capital, occupation, and industry controls as in the baseline specification. All estimates use the SHED survey weights and heteroskedasticity-robust standard errors.

Figure 2 reports adjusted application probabilities for each of the five frequency categories, based on estimates from equation (3). Because frequency enters the model categorically, the reported probabilities do not impose a linear relationship between AI-use intensity and job-search activity.

Figure 2: Adjusted probability of job application by frequency of generative AI use



Notes: The figure reports adjusted probabilities of having applied for another job during the previous 12 months for each of the five frequency-of-use categories among generative AI users. Predictions are based on equation (3), use the SHED survey weights, and condition on the baseline covariates. Error bars report 95% confidence intervals. The estimation sample contains 1,579 observations.

The adjusted application probabilities exhibit a monotonic pattern across the five frequency categories.⁹ An auxiliary ordinal specification implies an average increase of about 3.8 percentage points in the probability of having applied for another job when moving to a higher frequency category. The estimate should be interpreted as a summary of the ordered gradient rather than as implying equal distances between adjacent categories. The pattern extends the extensive-margin result: the association with job-

⁹A polynomial contrast based on the unrestricted categorical specification identifies a significant linear component ($F(1, 1508) = 15.79$, $p < 0.001$), while the quadratic, cubic, and quartic components are not statistically significant. See Figure D1 for the estimated differences relative to workers using generative AI less than one day per week.

search activity is stronger among workers who use generative AI more frequently. The gradient should be interpreted as descriptive. Frequency of use is not randomly assigned and may capture differences in workers, occupations, or workplaces that are not fully absorbed by the observed controls. The estimates do not imply that more frequent AI use causes additional job search. They show instead that the user–non-user difference in applications is accompanied by a systematic intensive-margin pattern.

5.4 Robustness checks

5.4.1 Robustness: Responses to randomized wage changes

Table 4: Robustness of responses to randomized wage changes

	(1) Logit	(2) Very likely	(3) Unweighted	(4) Fully classified sample	(5) No employer-type excl.
1% wage cut	0.066*** (0.021)	0.062*** (0.016)	0.074*** (0.020)	0.075*** (0.022)	0.061*** (0.021)
5% wage cut	0.273*** (0.021)	0.231*** (0.018)	0.278*** (0.019)	0.286*** (0.021)	0.260*** (0.020)
10% wage cut	0.354*** (0.020)	0.287*** (0.019)	0.339*** (0.019)	0.353*** (0.021)	0.338*** (0.020)
AI user	0.029 (0.029)	0.027 (0.021)	0.031 (0.028)	0.029 (0.031)	0.013 (0.030)
AI user $\times$ 1% cut	-0.016 (0.040)	0.001 (0.031)	-0.031 (0.039)	-0.009 (0.043)	-0.009 (0.041)
AI user $\times$ 5% cut	-0.020 (0.041)	-0.010 (0.036)	-0.024 (0.040)	-0.035 (0.043)	-0.003 (0.041)
AI user $\times$ 10% cut	-0.007 (0.038)	-0.008 (0.036)	0.013 (0.037)	0.015 (0.040)	0.021 (0.038)
Observations	6,025	6,025	6,025	5,635	6,285

Notes: The table reports robustness checks for the preferred specification reported in the main text. Column (1) estimates a logit model and reports contrasts between adjusted predicted probabilities. Column (2) uses a stricter definition of the dependent variable, equal to one only when the respondent reports being very likely to actively search for another job or leave the current position. Column (3) estimates the preferred specification without survey weights. Column (4) excludes observations assigned to unclassified, unavailable, or analogous residual categories in the control variables. Column (5) removes the employer-type restriction used in the preferred sample. Those observations are excluded from the preferred sample because their employer-type classification is inconsistent with the survey classification identifying respondents as working for someone else. The omitted wage scenario is a one-year wage freeze. Unless otherwise indicated, specifications use survey weights. All specifications include age, age squared, gender, race or ethnicity, education, household income, occupation, industry, working arrangements, job autonomy, union status, and the desire to work additional hours. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

We perform a series of robustness checks to assess the sensitivity of the baseline results. Table 4 reports the estimates. The results remain stable across specifications. Wage cuts continue to produce a positive and monotonic increase in the probability of searching for another job or leaving the current position. The stricter outcome yields smaller estimates for the 5 and 10 percent wage cuts, but the pattern remains economically large and precisely estimated. The interactions between AI use and the wage scenarios remain small and statistically insignificant in every column. The main conclusion therefore does not depend on the functional form, the definition of the outcome, the survey weights, the treatment of unavailable control information, or the composition of the sample.

5.4.2 Robustness: Job search and labor market transitions

Table 5 examines the sensitivity of these results to alternative specifications. Panel A reports average marginal effects from weighted logit models. Panel B removes the survey weights. Panel C estimates all five outcomes using a common sample. Panel D adds employment controls, while Panel E removes the employer-type restriction used in the preferred sample.

The association between AI use and job applications is highly stable. The estimated difference ranges from 4.8 to 6.2 percentage points and remains statistically significant at the 1 percent level in every specification.¹⁰ The association with voluntary exits is also positive across all panels. However, results for confidence and completed transitions are more sensitive to the specification: the coefficient on confidence becomes statistically insignificant after removing the survey weights or adding employment controls. Starting a new job and changing employer are significant only after adding employment controls. The latter estimates should be interpreted cautiously because the additional employment controls may themselves be jointly determined with recent job-search and

¹⁰The association with job applications also persists when we restrict the sample to respondents who neither started a new job nor changed employer during the previous 12 months ($\hat{\beta} = 0.056$, $p < 0.01$). The baseline association is therefore not driven by respondents who also report a completed employment transition during the same period.

Table 5: Robustness checks: Generative AI use and labor market outcomes

	(1) Confident	(2) Applied	(3) Left voluntarily	(4) Started job	(5) Changed employer
<i>Panel A. Weighted logit marginal effects</i>					
AI user	0.026* (0.014)	0.048*** (0.015)	0.026** (0.010)	0.019 (0.013)	0.016 (0.010)
Observations	5,965	6,025	5,980	6,025	6,025
<i>Panel B. Unweighted LPM</i>					
AI user	0.023 (0.014)	0.056*** (0.014)	0.021** (0.009)	0.018 (0.012)	0.014 (0.010)
Observations	5,965	6,025	6,025	6,025	6,025
<i>Panel C. Common outcome sample</i>					
AI user	0.027* (0.015)	0.053*** (0.015)	0.022** (0.010)	0.020 (0.013)	0.016 (0.010)
Observations	5,965	5,965	5,965	5,965	5,965
<i>Panel D. Employment controls</i>					
AI user	0.022 (0.015)	0.062*** (0.015)	0.028*** (0.010)	0.024* (0.013)	0.021** (0.010)
Observations	5,965	6,025	6,025	6,025	6,025
<i>Panel E. No employer-type exclusion</i>					
AI user	0.026* (0.014)	0.052*** (0.015)	0.023** (0.010)	0.020 (0.013)	0.013 (0.010)
Observations	6,222	6,285	6,285	6,285	6,285

Notes: The table reports robustness checks for the association between generative AI use and labor market outcomes. Panel A reports average marginal effects from weighted logit models. For the voluntary-exit outcome, 45 observations are dropped because one category perfectly predicts non-exit. Panel B estimates the preferred specification without survey weights. Panel C estimates all outcomes using the common sample of respondents with available information for the five dependent variables. Panel D adds controls for working arrangements, job autonomy, union status, and the desire to work additional hours. Panel E removes the employer-type restriction imposed in the preferred sample. All specifications control for age, age squared, gender, race and ethnicity, education, household income, occupation, and industry. Unless otherwise indicated, specifications use survey weights. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

mobility outcomes. The robustness exercises therefore identify active search as the most consistent margin associated with generative AI use.

As an additional check, Appendix Table E1 reweights non-users using entropy balancing (Hainmueller, 2012) so that the weighted distribution of their observed demographic, human-capital, occupation, and industry characteristics matches that of generative AI users. The estimated difference in job applications remains close to the baseline

result, at 4.5 percentage points, while the estimates for starting a new job and changing employer remain statistically insignificant.

6 Conclusion

The rapid diffusion of generative AI at work raises questions that go beyond its effects on productivity within the firm. Workers may also use the technology to gather information about alternative jobs, prepare applications, or develop capabilities that are valuable to other employers. However, it is less clear whether these changes translate into better opportunities outside the current job. We study that distinction using the 2025 SHED, which combines information on workplace generative AI use with labor-market behavior and a randomized experiment in which employees face hypothetical wage freezes or wage cuts of different magnitudes. The design lets us compare how workers respond to lower pay with the relationship between AI use and different stages of job mobility.

The results show differences between job search and actual job transitions. Workers who use generative AI at work are about 5 percentage points more likely to report having applied for another job during the previous 12 months. Among AI users, application probabilities also increase monotonically with the frequency of use, and the association with applications is the only one that remains statistically significant after adjusting for multiple testing. The evidence is much weaker at later stages of the mobility process. AI users are not systematically more likely to start a new job or change employer, and the associations with perceived employability and voluntary exits are not robust across specifications and after multiple-testing adjustment. The randomized wage scenarios point in the same direction. Larger wage cuts substantially increase stated search or exit intentions, but we find no evidence that the response differs between AI users and non-users. Workplace generative AI use is therefore associated with greater job-search activity, without corresponding evidence of greater realized mobility or stronger responsiveness to adverse wage changes.

The difference between search and mobility is important. Job search is only one step in the process of changing jobs. A technology can lower the cost of identifying vacancies, comparing jobs, or preparing applications without changing the rate at which attractive offers arrive or the gains workers expect from accepting them. Our results suggest

that generative AI use matters more at this earlier stage of the job-mobility process. They provide less support for the idea that workplace AI use already gives workers better outside options. If that were the case, we might also expect workers to feel more confident about finding another job, change jobs more often, or respond more when their current job becomes less attractive. This distinction also matters for how we interpret the labor-market effects of generative AI. Productivity gains inside the current firm need not immediately increase worker mobility across firms, and greater exploration of outside opportunities need not imply that those opportunities have become more valuable.

Several features of the data limit how far the results can be taken. Generative AI adoption is not randomly assigned, and differences between users and non-users may reflect worker, occupation, or workplace characteristics that remain unobserved after conditioning on a rich set of covariates. Frequency of use is subject to the same limitation. Timing also prevents a causal interpretation of the observed labor-market associations: generative AI use refers to the previous month, whereas applications and employment transitions are measured over the previous 12 months. The randomized wage experiment solves a different identification problem. It identifies responses to changes in hypothetical compensation, but not the causal effect of AI use and not necessarily workers' behavior following realized wage changes. Future longitudinal data could link the adoption and intensity of workplace AI use to the successive stages of job search, offers, accepted jobs, and subsequent earnings. This would help determine whether the higher search activity documented here is a temporary feature of early AI adoption, a persistent reduction in search frictions, or the first stage of a change in worker mobility as generative AI becomes more integrated into work.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used ChatGPT (OpenAI) to assist with language editing, improve the clarity and organization of the manuscript, identify minor typographical and formatting inconsistencies, and improve the readability and organization of the code. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

- Alekseeva, L., Azar, J., Giné, M., Samila, S., and Taska, B. (2021). The Demand for AI Skills in the Labor Market. *Labour Economics*, 71:102002.
- Bagga, S. (2023). Firm Market Power, Worker Mobility, and Wages in the US Labor Market. *Journal of Labor Economics*, 41(S1):S205–S256.
- Bick, A., Blandin, A., and Deming, D. J. (2026). The Rapid Adoption of Generative AI. *Management Science*, 0(0).
- Board of Governors of the Federal Reserve System (2026a). Economic Well-Being of U.S. Households in 2025. Report, Board of Governors of the Federal Reserve System, Washington, DC.
- Board of Governors of the Federal Reserve System (2026b). Supplemental Appendixes to the Report on the Economic Well-Being of U.S. Households in 2025. Report, Board of Governors of the Federal Reserve System, Washington, DC. Appendix B: Consumer Responses to Survey Questions.
- Borgonovi, F., Calvino, F., Criscuolo, C., Samek, L., Seitz, H., Nania, J., Nitschke, J., and O’Kane, L. (2023). Emerging Trends in AI Skill Demand across 14 OECD Countries. OECD Artificial Intelligence Papers 2, OECD Publishing, Paris.
- Brynjolfsson, E., Li, D., and Raymond, L. R. (2025). Generative AI at Work. *The Quarterly Journal of Economics*, 140(2):889–942.
- Caldwell, S., Haegele, I., and Heining, J. (2025). Why Workers Stay: Pay, Beliefs, and Attachment. NBER Working Paper 33445, National Bureau of Economic Research.
- Call, M. L., Jiang, K., and Idso, C. (2026). Star Advantage: Employee Value Creation and Capture in the Age of Artificial Intelligence. *Human Resource Management*, 65(1):151–167.

- Cazzaniga, M., Pizzinelli, C., Rockall, E., and Tavares, M. M. (2025). Exposure to Artificial Intelligence and Occupational Mobility: A Cross-Country Analysis. *Economía*, 24(1):314–339.
- Cedefop (2025). Skills Empower Workers in the AI Revolution: First Findings from Cedefop’s AI Skills Survey. Policy Brief 9201, Publications Office of the European Union, Luxembourg.
- Doellgast, V., Appalla, S., Ginzburg, D., Kim, J., and Thian, W. L. (2025). Global Case Studies of Social Dialogue on AI and Algorithmic Management. ILO Working Paper 144, International Labour Organization, Geneva.
- Dube, A., Giuliano, L., and Leonard, J. (2019). Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior. *American Economic Review*, 109(2):620–663.
- Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs Are GPTs: Labor Market Impact Potential of LLMs. *Science*, 384(6702):1306–1308.
- Faberman, R. J., Mueller, A. I., Şahin, A., and Topa, G. (2022). Job Search Behavior Among the Employed and Non-Employed. *Econometrica*, 90(4):1743–1779.
- Gambacorta, L., Qiu, H., Shan, S., and Rees, D. M. (2026). Generative AI and Labour Productivity: A Quasi Experiment on Coding. *Journal of Financial Stability*, 84:101543.
- Gathmann, C., Grimm, F., and Winkler, E. (2025). AI, Task Changes in Jobs, and Worker Reallocation. CEPR Discussion Paper 19864, Centre for Economic Policy Research.
- Hainmueller, J. (2012). Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies. *Political Analysis*, 20(1):25–46.
- Hartley, J. S., Jolevski, F., Melo, V., and Moore, B. (2026). The Labor Market Effects of Generative Artificial Intelligence. SSRN Working Paper.

- Holm, S. (1979). A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, 6(2):65–70.
- Humlum, A. and Vestergaard, E. (2025a). Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI. NBER Working Paper 33777, National Bureau of Economic Research. Revised March 2026.
- Humlum, A. and Vestergaard, E. (2025b). The Unequal Adoption of ChatGPT Exacerbates Existing Inequalities among Workers. *Proceedings of the National Academy of Sciences of the United States of America*, 122(1):e2414972121.
- International Labour Organization (2026). Lifelong Learning and Skills for the Future. World of work report, International Labour Organization, Geneva.
- Jäger, S., Roth, C., Roussille, N., and Schoefer, B. (2024). Worker Beliefs About Outside Options. *The Quarterly Journal of Economics*, 139(3):1505–1556.
- Kim, J. H., Mester, N., and Sun, G. (2026). AI and the Labor Market: A Worker’s-Eye View. SSRN Working Paper.
- Mortensen, D. T. (1986). Job Search and Labor Market Analysis. In Ashenfelter, O. and Layard, R., editors, *Handbook of Labor Economics*, volume 2, pages 849–919. Elsevier.
- Noy, S. and Zhang, W. (2023). Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Science*, 381(6654):187–192.
- OECD (2026). AI and Skills: What We Know So Far. Policy brief, OECD Publishing, Paris.
- Stephany, F. and Teutloff, O. (2024). What Is the Price of a Skill? The Value of Complementarity. *Research Policy*, 53(1):104898.

A Additional descriptive evidence

A.1 Job-search intentions by generative AI use

Table A1 reports the unweighted distribution of stated job-search or exit intentions by generative AI use, pooling the four randomized wage scenarios. The distributions are similar across groups. Among AI users, 55.4 percent report being very or somewhat likely to search for another job or leave their current position, compared with 54.1 percent of non-users. AI users are also less likely to select the lowest response category, with shares of 16.2 and 18.8 percent, respectively.

Table A1: Generative AI use and stated job-search intentions

	Very likely	Somewhat likely	Not that likely	Not at all likely	Total
Non-user	1,205 (27.1)	1,201 (27.0)	1,205 (27.1)	835 (18.8)	4,446 (100.0)
AI user	432 (27.4)	442 (28.0)	450 (28.5)	255 (16.2)	1,579 (100.0)
Total	1,637 (27.2)	1,643 (27.3)	1,655 (27.5)	1,090 (18.1)	6,025 (100.0)

Notes: The table cross-tabulates generative AI use with respondents' reported likelihood of actively searching for another job or leaving their current position. Entries report unweighted frequencies. Row percentages are reported in parentheses.

Source: Authors' elaboration.

B Variables

Tables B1 and B2 describe the construction of the variables used in the analysis. The first table covers the main explanatory variables and outcomes. The second reports the demographic, human-capital, and employment controls.

Table B1: Variable definitions: main variables

Variable	Question number	Definition
AI user	D50	Indicator equal to one when the respondent reports using generative AI at work during the previous month.
AI usage intensity	D51	Frequency of reported generative AI use at work during the previous month.
Wage scenario	DOV_D36_AMOUNT	Randomly assigned hypothetical wage scenario. The four categories are a one-year wage freeze and wage cuts of 1, 5, and 10 percent.
Likely to search or leave	D36B	Indicator equal to one when the respondent reports being very likely or somewhat likely to actively search for another job or leave the current position under the assigned wage scenario.
Changed employer	D37A	Indicator equal to one when the respondent started a new job during the previous 12 months and reports that the current main job is with a different employer.
Applied for a new job	D44	Indicator equal to one when the respondent reports applying for a new job during the previous 12 months.
Started a new job	D44	Indicator equal to one when the respondent reports starting a new job during the previous 12 months.
Voluntarily left a job	D44	Indicator equal to one when the respondent reports voluntarily leaving a job during the previous 12 months.
Confident of finding a job	D1B	Indicator equal to one when the respondent is somewhat or very confident of finding a job within the next three months. Respondents who are not working and are not interested in working are excluded.

Source: Authors' elaboration.

Table B2: Variable definitions: control variables

Variable	Source	Definition
<i>Demographic controls</i>		
Age	ppage	Respondent's age in years. Age enters the regressions in linear and quadratic form.
Gender	ppgender	Categorical variable distinguishing male and female respondents.
Race and ethnicity	ppethm	Categorical variable distinguishing White non-Hispanic, Black non-Hispanic, other non-Hispanic, Hispanic, and multiracial non-Hispanic respondents.
<i>Human-capital controls</i>		
Education	ppeducat	Highest educational attainment: less than high school, high school, some college or associate degree, and bachelor's degree or higher.
Household income	ppinc7	Annual household income, grouped into seven categories ranging from less than \$10,000 to \$150,000 or more.
Occupation	occupation	Categorical classification of the respondent's occupation in the current job.
Industry	industry	Categorical classification of the industry associated with the respondent's current job.
<i>Employment controls</i>		
Usual working hours	D48	Categorical variable capturing the respondent's usual working hours.
Multiple-job holding	D4	Categorical variable capturing multiple-job holding.
Work schedule	D30	Categorical variable capturing the respondent's usual work schedule.
Remote work	D34A	Categorical variable capturing the extent to which the respondent works remotely.
Job autonomy: task choice	D28_a	Frequency with which the respondent can choose which tasks to work on.
Job autonomy: task execution	D28_b	Frequency with which the respondent can choose how to perform tasks at work.
Union status	UNION100	Categorical variable indicating whether the respondent belongs to a labor union or a similar employee association.
Desire to work additional hours	D1E	Indicator for whether the respondent wanted to work additional hours during the previous month.

Notes: The source column reports the variable in the SHED data used to construct each control. Employment controls are included in the corresponding extended specifications.

Source: Authors' elaboration.

C Multiple-testing adjustment

Table C1 reports the multiple-testing adjustment for the five labor-market outcomes examined in Table 3. We apply the Holm procedure to the five baseline coefficients on generative AI use. The association with job applications remains statistically significant after the adjustment, whereas none of the remaining outcomes does.

Table C1: Multiple-testing adjustment

Outcome	p -value	
	Unadjusted	Holm-adjusted
Confident	0.067	0.202
Applied	0.001	0.004
Left voluntarily	0.027	0.107
Started job	0.164	0.325
Changed employer	0.163	0.325

Notes: The table reports unadjusted and Holm-adjusted p -values for the coefficient on generative AI use across the five outcomes reported in Table 3. The adjustment is applied jointly to the five baseline weighted linear probability models.

D Generative AI use frequency

Table D1 reports the distribution of generative AI use frequency among workers who reported using generative AI at work during the previous month.


Table D1: Frequency of generative AI use at work

Frequency	N	Percent
Multiple times per day	224	14.19
Daily	224	14.19
Several days per week	373	23.62
Weekly	290	18.37
Less than one day per week	468	29.64
Total	1,579	100.00

Notes: The table reports the unweighted distribution of the frequency of generative AI use at work among respondents who reported using generative AI at their job during the previous month. Percentages are calculated over the 1,579 respondents with available information on frequency of use.

Source: Authors' elaboration.

Figure D1: Differences in job-application probability by frequency of generative AI use



Notes: The figure reports estimated differences in the probability of having applied for another job during the previous 12 months across generative AI use-frequency categories. The estimates are based on equation (3) and are expressed relative to workers using generative AI less than one day per week. Estimates use the SHED survey weights and condition on the baseline covariates. Error bars report 95% confidence intervals. N = 1,579.

E Entropy-balancing robustness

The baseline regressions condition on a broad set of worker characteristics, but generative AI users differ substantially from non-users in several observable dimensions, particularly education and household income. We assess whether these differences account for the associations reported in Table 3 using entropy balancing (Hainmueller, 2012). The procedure reweights non-users so that their observed characteristics match those of generative AI users. We balance age and age squared, gender, race and ethnicity, education, household income, occupation, and industry, using the SHED survey weights as base weights. The balancing constraints are satisfied for the included covariates, substantially reducing the observable differences between users and non-users documented in Table 1.

Table E1 reports the resulting estimates for the five labor-market outcomes. The association with job applications remains close to the baseline estimate. After balancing, generative AI users are 4.5 percentage points more likely to report having applied for another job during the previous 12 months.

Table E1: Entropy-balancing robustness

	(1) Confident	(2) Applied	(3) Left voluntarily	(4) Started job	(5) Changed employer
AI user	0.017 (0.016)	0.045*** (0.016)	0.022** (0.010)	0.014 (0.013)	0.014 (0.010)
Observations	5,965	6,025	6,025	6,025	6,025

Notes: The table reports linear probability models after entropy balancing. Non-users of generative AI are reweighted to match AI users on age, age squared, gender, race and ethnicity, education, household income, occupation, and industry. The SHED survey weights are used as base weights in the balancing procedure. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

The entropy-balancing exercise therefore preserves the main pattern in the baseline results. The strongest difference between generative AI users and reweighted non-users remains concentrated at the job-search margin, with no corresponding difference in realized mobility. Entropy balancing addresses differences in observed worker composition but does not account for unobserved differences between AI users and non-users. The

estimates should therefore continue to be interpreted as descriptive associations rather than causal effects.