AI and Labor-Market Reallocation*

Hie Joo Ahn† Nicholas A. Carollo‡

July 10, 2026

Abstract

This paper documents how artificial intelligence has affected stocks and flows in the U.S. labor market. Combining CPS and JOLTS data with occupation-level AI exposure and adoption measures from OpenAI and Lightcast, respectively, we construct labor force stocks, worker flows, and market tightness by AI exposure and adoption. Since the introduction of LLMs, workers with high AI exposure and adoption have experienced larger declines in job-finding and job-switching rates than other groups, while their within-job activity switching has increased noticeably relative to others. The positive effects of AI on nominal wage growth and hours worked attenuate markedly in the post-LLM period, suggesting that LLM-driven reallocation has operated through within-firm task reorganization and a reconfiguration of labor inputs, alongside weaker demand for workers most exposed to AI. To quantify LLM-driven reallocation pressure, we construct individual-level hirability and separability indices and show that dispersion in job-finding prospects has risen markedly since 2023, largely driven by the AI factors. The natural rate of unemployment—recovered from the trend components of unemployment inflows and outflows by AI exposure and adoption, as well as from a theoretical model of reallocation—is estimated to have risen by about 0.2 percentage point since the introduction of LLMs, albeit with considerable uncertainty.

JEL classification: E24, J21, J62, J64, O33.

Keywords: artificial intelligence, labor market flows, unemployment, natural rate of unemployment, reallocation, wages

*The views expressed in this paper are those of the authors and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff. We thank Travis Berge, Leland Crane, Andrew Figura, Chris Nekarda, John Stevens and Paul Soto for helpful comments; any remaining errors are our own.

†Corresponding author. Federal Reserve Board of Governors, 20th Street and Constitution Avenue NW, Washington, DC 20551, U.S.A. Email: econ.hjahn@gmail.com

‡Federal Reserve Board of Governors, 20th Street and Constitution Avenue NW, Washington, DC 20551, U.S.A. Email: carollo.nicholas@gmail.com

We are suffering, not from the rheumatics of old age, but from the growing pains of over-rapid changes, from the painfulness of readjustment between one economic period and another. The increase of technical efficiency has been taking place faster than we can deal with the problem of labour absorption. We are being afflicted with technological unemployment. — Keynes (1930)

1 Introduction

The rapid diffusion of artificial intelligence—and large language models (LLMs) in particular—has prompted significant debate about its consequences for the labor market. Whether AI will predominantly complement human workers, substitute for them, or trigger broad reallocation across occupations and industries is a central theme of these discussions. A growing literature has examined the productivity and employment effects of AI for specific worker groups and occupations using micro-level data. However, much less is known about how LLMs reshape the reallocation and functioning of the labor market. In this paper, we examine the effects of LLMs on aggregate labor market dynamics, focusing on labor market flows and stocks—particularly unemployment inflows and outflows—and the natural rate of unemployment. To this end, we integrate traditional labor market data—such as CPS, JOLTS, and O*NET—with nontraditional data sources from OpenAI, Anthropic, and Lightcast.

Our approach of measuring AI's labor market effects has a few key aspects. First, we match individuals in the CPS microdata to measures of AI exposure and AI adoption for their detailed occupation. Specifically, we draw on occupation-level exposure scores from Eloundou et al. [2023] and Handa et al. [2025], which characterize the extent to which generative AI could potentially perform the tasks associated with each occupation. For AI adoption, we compute the share of positions that include AI and machine-learning related skills from Lightcast job postings, capturing the extent to which firms have actually integrated AI into production. This adoption measure is strongly correlated with the Census Business Trends and Outlook Survey at the industry level, but allows us to measure the adoption intensity for detailed occupations, which BTOS lacks. In other words, AI exposure captures the supply of skills that can be theoretically

replaced by AI, and AI adoption measures actual demand for (or integration of) AI-related skills. We consider both dimensions of AI intensity separately and jointly to comprehensively assess their independent and interaction effects. By assigning AI exposure and adoption scores to each individual in the CPS, we directly assess how AI has affected the unemployment rate, job-finding and job-loss rates, and various other labor market outcomes.

Next, we construct market tightness by AI exposure and adoption. Because the JOLTS data are reported by industry but not occupation, we allocate vacancies to occupations based on the occupational composition of Lightcast job postings within each industry and month. This measure allows us to analyze how AI has shaped the interaction between labor supply and demand and provides information about the labor market slack and functioning.

Our analysis produces four main findings. *First*, workers in high–AI-exposure and high–AI-adoption jobs have lower unemployment rates than those in low-exposure, low-adoption jobs, primarily reflecting their lower separation (EU) rates. Although AI exposure and adoption are associated with lower job-finding (UE) rates—consistent with longer search durations for higher-skilled positions—this effect is more than offset by correspondingly lower separation rates.

Second, however, the unemployment rate of the high-exposure/high-adoption group rose more than that of other groups, particularly those with lower exposure and adoption who are typically more cyclically sensitive. This differential weakness is driven by the hiring margin rather than the separation or layoff margin: since the introduction of LLMs, job-finding rates have declined more for high-exposure/high-adoption workers than for other groups, while job-loss rates show little differential change across the groups. A similar pattern is observed in job-to-job transition rates, another margin of hiring. Consistently, market tightness for workers with high AI exposure and adoption has remained below pre-COVID levels since 2023, while tightness for low-exposure workers has largely returned to its pre-COVID level. These results suggest that LLMs have disproportionately reduced labor demand for high-exposure, high-adoption workers, with adjustments occurring primarily along the hiring margin rather than the separation margin, consistent with prior microeconomic evidence [Brynjolfsson et al., 2025, Hosseini and Lichtinger, 2025].

Third, LLMs have reshaped labor allocation within firms. The rate of within-job activity switching rises disproportionately for workers with high–AI-exposure, pointing to increased task

reorganization. We also find that the positive effects of AI on nominal wage growth and hours worked attenuate markedly in the post-LLM period. In the pre-LLM period, both AI exposure and adoption are associated with higher nominal wage growth, and AI exposure is linked to longer hours. These effects weaken following the introduction of LLMs, suggesting softer labor demand for highly exposed workers and a reconfiguration of labor input.

Fourth, these uneven effects of LLMs signals a potential buildup of reallocation pressure, wherein workers are unable to move to sectors with better reemployment prospects due to mismatch or other frictions. Reallocation pressure is traditionally measured with the cross-sectional dispersion of market tightness, and employment growth (e.g., Lilien, 1982, Jackman and Roper, 1987, Şahin et al., 2014). To assess the reallocation pressure driven by the AI factors, we construct individual-level *hirability* and *separability* indices that measure an individual's propensity of job finding from unemployment and of job loss, respectively. In the estimation, we include market tightness and industry and time fixed effects as controls to distinguish LLMs' effects from those of macroeconomic as well as industry-specific conditions. We show that the cross-sectional dispersion in the hirability index increased noticeably from the end of 2022, with the AI factors accounting for nearly all of the rise. By contrast, the dispersion of the separability index shows no comparable increase, consistent with LLMs primarily affecting reallocation through the hiring margin (Hosseini Maasoum and Lichtinger, 2026).

Fifth, the overall empirical evidence points to upward pressure on the natural rate of unemployment in the post-LLM period.[1] We employ two approaches to quantify LLMs' effects on the natural rate. To begin with, we extend the theoretical model of structural unemployment from Jackman and Roper [1987] and establish that greater cross-sectional dispersion in job-finding probabilities reflects increased frictions in reallocating workers to sectors with better reemployment prospects. Calibrating this theoretical condition with the estimated dispersion in individual-level hirability and LLMs' contributions, we demonstrate that LLMs likely raised the structural unemployment rate by between one-tenth and two-tenths of a percentage point, depending on parameter calibration.

Moreover, we statistically estimate the natural rate of unemployment for workers groups

[1]We treat the structural unemployment rate as a component of the natural rate of unemployment, which includes both frictional and structural unemployment (Daly et al., 2012).

with different AI exposure and adoption by feeding each group's trend job-loss and job-finding rates into Shimer [2012]'s two-state model.[2] We employ a Bayesian trend–cycle decomposition to extract group-specific trend components of unemployment flows rates. One merit of this approach is to allow us to control pre-LLM trends and business cycles for the different groups. Reflecting the persistent weakening of job-finding rates for high-exposure workers relative to low-exposure workers, the natural rate is estimated to have risen by about two-tenths of a percentage point since the introduction of LLMs through mid-2026, albeit with considerable uncertainty. This statistical estimate aligns with the prediction of the theoretical model.

Our empirical findings have policy implications. The weaker job-finding rates of workers with high AI exposure and adoption than others and their steady decline since the LLMs' introduction could be the joint outcomes of uncertainty about both the macroeconomic environment and the capabilities and limitations of LLMs. To the extent that the low-hiring, low-firing environment may itself be an endogenous response to uncertainty, policies that reduce uncertainty may help alleviate reallocation pressure and facilitate more efficient labor market adjustment. Recent signs of recovery in job-finding and job-switching rates for high-exposure/high-adoption workers, along with a improved market tightness among them, could signal somewhat reduced uncertainty associated with the hiring of workers in high-AI-intensity jobs.

This paper contributes to the literature on the labor market effects of artificial intelligence along several dimensions. First, it provides the first comprehensive analysis of how AI shapes labor market stocks and flows along with other labor market outcomes, with a particular focus on unemployment dynamics and the natural rate of unemployment. Our focus on aggregate labor market distinguishes this paper from prior work, which has primarily examined AI's effects on specific occupations, demographic groups, or selected labor market outcomes such as employment and productivity [Brynjolfsson et al., 2018, Webb, 2020, Felten et al., 2023, Eloundou et al., 2023, Acemoglu et al., 2022, Hampole et al., 2025, Brynjolfsson et al., 2025, Hosseini and Lichtinger, 2025, Crane and Soto, 2026]. Second, our approach considers both the supply and demand sides of LLMs by jointly incorporating AI exposure and adoption in analyzing their effects on labor market reallocation and structural unemployment, whereas existing studies typically consider these factors in isolation. This joint framework is suitable for the measurement

[2]This approach of estimating the natural rate is considered by Tasci [2012], Crump et al. [2019], Ahn [2023].

of market tightness, and reveals differential effects of AI adoption: the job-finding rate was weaker among workers in occupations with high AI exposure and high adoption than among those with high AI exposure but low adoption.

Last, this paper is the first to develop both a conceptual and empirical framework for assessing the effects of LLMs on the natural rate of unemployment. To this end, we introduce three novel components. First, to shed light on the forces driving changes in the natural rate, we construct individual-level hirability and separability indices that capture the effects of AI on job-finding and job-loss propensities, as well as their contribution to reallocation pressure. Second, we extend the theoretical framework of Jackman and Roper [1987] by incorporating the dispersion of the hirability index, allowing us to quantify changes in structural unemployment due to LLMs. Third, we adopt a Bayesian trend–cycle model of unemployment flows, building on Shimer [2012], to estimate the natural rate of unemployment for disaggregated groups defined by AI exposure and adoption, as well as in the aggregate. A more detailed discussion of the related literature is provided in the literature review section below.

Section 2 discusses related literature. Section 3 describes the datasets used for the empirical analyses. Section 4 reports demographics and labor market outcomes by AI exposure and adoption, focusing on stocks and flows of workers and jobs. Section 5 introduces the individual-level separability and hirability indices and analyzes the extent to which AI contributes to the dispersion in the likelihood of job findings and job separations. Section 6 discusses implications for the natural rate of unemployment and the trend labor force participation rate.

2 Literature Review

Our work contributes, first, to a rapidly growing empirical literature on the labor market effects of artificial intelligence. One branch of this literature focuses on identifying which occupations and industries are potentially exposed to AI-related disruptions based on their task structures [Brynjolfsson et al., 2018, Webb, 2020, Felten et al., 2023, Eloundou et al., 2023]. Recent work in this vein leverages large-scale data on user interactions with LLMs to disentangle which tasks AI may automate and which it may augment [Handa et al., 2025, Tomlinson et al., 2025].

In addition, the recent supplement of Business Trends and Outlook Survey by Census asks

whether firms are using AI to substitute for tasks, supplement existing tasks, or introducing new tasks. One takeaway from this survey is that more firms report using AI in an augmentative way than report using AI to fully automate tasks.[3] This finding is consistent with our observation of an increase in within-job activity changes among workers with high AI exposure and high AI adoption.

Yin et al. [2026] raises concerns about the stability of AI exposure measures, underscoring the need to consider multiple measures for robustness. We draw on AI exposure measures from OpenAI and Anthropic alongside adoption measures from Lightcast and Business Trends and Outlook Survey to ensure reliable measurement.

Moving beyond these measures, [Manning and Aguirre, 2026] develop an occupation-level adaptive capacity index that measures how capable workers are of adjusting to job displacement, and find that many workers highly exposed to AI also have high adaptive capacity perhaps suggesting that they might be well-placed to manage a labor market transition. The hirability and separability indices we construct in this paper extend this idea to individual workers and allow us to characterize the extent to which AI has already contributed to reallocation pressure in the labor market.

Another branch of literature examines the employment effects of AI. [4] Studies focusing on traditional AI and machine learning applications have generally found small net employment effects and evidence of within-firm task reallocation [Acemoglu et al., 2022, Hampole et al., 2025]. Autor and Thompson [2025] argue that, in an extension of the task framework, the key determinant is whether AI substitutes for expert or inexpert tasks within an occupation. Autor and Thompson further find that substituting for expert tasks tends to lower wages but increase employment, whereas substituting for inexpert tasks raises wages but reduces employment. Recent work focusing on the adoption of generative AI, however, has documented adverse employment effects for early-career workers in highly-exposed occupations [Brynjolfsson et al., 2025, Hosseini and Lichtinger, 2025, Tucker, 2026] and reduced employment growth for coders

[3] See `https://www.census.gov/programs-surveys/btos.html` for more details.

[4] When assessing AI's effects on employment, Acemoglu and Restrepo [2019] provide a useful framework for evaluating how new technologies shape labor demand. They highlight three key channels: (1) the displacement effect, (2) the reinstatement effect (job creation), and (3) the productivity effect. Depending on which force dominates, AI may either increase or decrease demand for human workers—both in aggregate and across skill groups. Importantly, the timing of these effects may differ, with existing jobs potentially displaced more rapidly than new ones are created.

in particular [Crane and Soto, 2026].[5]

Meanwhile, The Budget Lab at Yale [2025a] found no relationship between measures of AI exposure and recent changes in employment or unemployment. The Economic Innovation Group (Eckhardt and Goldschlag, 2025) also reached this conclusion. Meanwhile, a September 2025 piece from the St. Louis Fed finds some evidence that unemployment rose more between 2022 and 2025 for more AI-exposed occupations.

Recent studies have examined AI's effects on wages and inequality. For example, Freund and Mann [2026] develop a general-equilibrium model of job transformation and find that wage effects are non-monotonic in exposure: moderate exposure raises average wages while the most exposed workers lose substantially, with large dispersion within occupations. Althoff and Reichardt [2026] find that AI reduces wage inequality by enabling lower-skill workers to compete for previously inaccessible jobs. Our findings are broadly consistent: job-finding rates are lower and declined faster among workers with higher AI intensity than among those with lower AI intensity in the post-LLM era; AI exposure and adoption are associated with higher wage levels and growth, though these premia narrowed after the introduction of LLMs, reflecting weaker demand for high-exposure/high-adoption workers. Moreover, LLMs' effects differ across workers—job switchers with high AI exposure enjoy higher wage growth, consistent with Freund and Mann [2026], while those with high AI adoption experience lower wage growth, pointing to wage disparities along the exposure-adoption dimension.

Distinguished from previous studies, our paper provides the first comprehensive analysis of how AI shapes labor market stocks, flows, and outcomes, with a particular focus on unemployment dynamics and the natural rate of unemployment.

Importantly, this paper contributes to a broader macro-labor literature on labor market reallocation. To begin with, the literature on effects of reallocation on unemployment, Lilien [1982] interprets increases in the dispersion of sectoral employment growth during downturns as evidence of sectoral reallocation, implying that the associated rise in unemployment is largely structural. In contrast, Abraham and Katz [1986] argue that countercyclical dispersion can arise from heterogeneous sensitivities of sectors to aggregate shocks, so that the comovement

[5]Brynjolfsson et al. [2025] find that since late 2022, early-career workers in occupations highly exposed to AI automation have experienced a decline in employment relative to more experienced workers in the same occupations and workers of all ages in occupations less exposed to AI.

between employment dispersion and the unemployment rate does not necessarily indicate an increase in structural unemployment during recessions. Different from these views, Neumann and Topel [1991] claim that the low-frequency evolution of sectoral employment shares captures permanent demand shift across industries, creating the persistent reallocation pressure that drives up structural unemployment.[6] Jackman and Roper [1987] argue that dispersion in market tightness signals changes in structural unemployment. Building on this framework, Şahin et al. [2014] construct a mismatch index. We extend the framework of Jackman and Roper [1987] and show that LLMs are the primary source of dispersion in individual-level hirability and increased reallocation pressure.

Recent studies also focus on LLMs' reallocation effects. Hampole et al. [2025] show that, although highly exposed occupations experience lower demand, productivity increases lead firms to expand employment elsewhere (albeit mostly for "traditional" AI applications that pre-date generative AI). Baslandze et al. [2026] document compositional labor reallocation both within and across firms, with routine clerical roles declining and relative demand for skilled technical roles increasing. The Budget Lab at Yale [2025b] showed that the occupational mix of the US labor market is changing somewhat faster than it has historically. A survey conducted by the Federal Reserve Bank of New York (Abel et al. [2025]) shows that firms are more likely to retrain workers in AI-exposed occupations than to lay them off. While some firms report reducing hiring in response to AI, others report increasing demand for workers with AI-related skills.

Our results are broadly consistent with these findings, as we also find that small negative employment effects appear to stem from reduced hiring rather than increased separations. In addition, we find job-to-job transition rates declined, aligning with reduced job-finding rates, but the rate of activity changes within the same job picked up among workers with high AI exposure. This observation suggests LLM-driven reallocation has occurred through task reorganization within firms and reconfiguration of labor input alongside a decline in demand for workers with high AI exposure and high AI-adoption jobs.

All told, our paper complements the microeconomic evidence of previous literature by

[6] Rissman [1993] adopts the measure of Neumann and Topel [1991] for the empirical investigation of wage Phillips curve.

providing the first systematic analysis of AI's influence on unemployment flows and stocks, with particular emphasis on its implications for structural unemployment.

3 Data

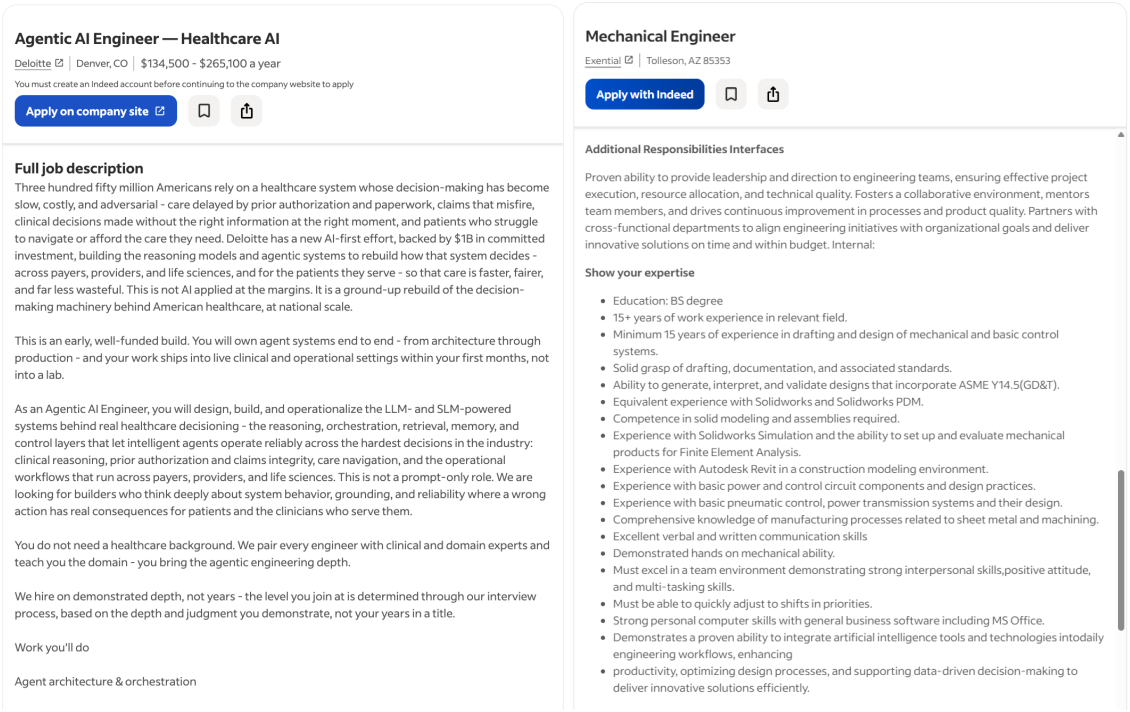
3.1 AI Exposure and Adoption

AI exposure We obtain our preferred measure of an occupation's exposure to AI from Eloundou et al. [2023], which we refer to as OpenAI exposure scores throughout this paper. This ranking was constructed by labeling direct work activities and task descriptions from the Occupational Information Network (O*NET) based on an assessment of whether access to an LLM or LLM-powered system would reduce the time required to complete the activity or task by at least 50% without loss of quality. We use the authors' GPT-4 assigned β exposure measure, which considers the use of complementary tools or software on top of LLMs. Occupation-level exposure scores are averages over task-level exposure, where supplemental O*NET tasks are assigned half the weight of core tasks.

In robustness checks, we use an alternative measure of AI exposure based on data from the Anthropic Economic Index [Handa et al., 2025]. This dataset provides descriptions of hundreds of distinct AI applications derived from millions of prompts submitted to Claude. For each AI application (e.g. "develop time series forecasting models"), we convert the text to a numerical word embedding and do same for each O*NET task description. Following [Hampole et al., 2025], we consider a task exposed to AI if the cosine similarity between the task embedding and *any* AI application embedding is above the 95th percentile of similarity across all task-application pairs. We then aggregate task-level exposure to occupations as described above.

AI adoption Both the OpenAI and Anthropic exposure scores capture the extent to which generative AI applications might *feasibly* perform the same tasks as a human worker. However, the degree to which AI is actually being used to accomplish exposed tasks varies considerably across occupations and industries and has risen over time as the underlying technology has improved. We therefore supplement our AI exposure scores with a measure of AI utilization constructed from online job postings. Using data from Lightcast, we compute the share of job

Figure 1: EXAMPLE OF JOB POSTINGS REQUIRING AI AND ML SKILLS



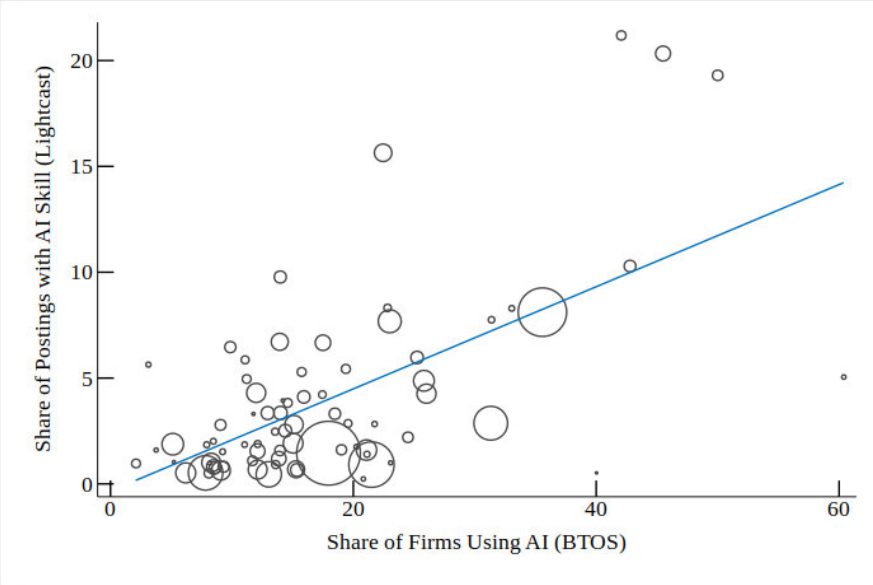
Notes: This figure displays actual job postings that require AI skills from Indeed. The Lightcast data include postings from Indeed, among many other job posting websites, and tags listings like these as including AI-related skills. **Source:** Indeed.

postings that mention an AI-related skill requirement for each occupation, industry, and month. We consider a position to be utilizing AI if Lightcast tags it with a skill that includes the terms "artificial intelligence" or "machine learning" in its description. Figure 1 provides two examples of Indeed job postings that we identify as requiring AI or ML skills in the Lightcast data—one for an agentic AI engineer at Deloitte and another for a mechanical engineer who is required to "integrate artificial intelligence tools and technologies into daily engineering workflows".

Because we infer AI adoption from job postings, measurement error is a potential concern. First, job postings reflect the skills required of new hires and may therefore miss the extent to which AI is being used by current employees. This could understate the actual extent of AI adoption.[7] Second, postings may include AI-related terms in contexts unrelated to the skill requirements of the job. For example, Indeed also contains postings with disclaimers such as "we

[7]We might also understate AI adoption if postings use more specific terminology like "LLM" or "ChatGPT" instead of the broader terms we consider. In our baseline classification that relies on mean exposure and adoption (discussed in Section 4), this issue is less of a concern because our measure mostly relies on relative rankings rather than levels. In addition, the vast majority of AI skills are the tags "artificial intelligence" and "machine learning." Other tags are much less common.

Figure 2: CORRELATION BETWEEN LIGHTCAST AND BTOS AI ADOPTION



Note: This figure plots correlations between Lightcast- and BTOS-based AI adoption measures at the 3-digit NAICS level. The Lightcast measure is the share of job postings tagged with an AI skill from 2025:M1 to 2025:M12. The BTOS measure is the share of firms reporting having used AI in any business application over the past two weeks, averaged over 2025:M11 to 2026:M6. **Source:** Authors' calculation.

may use artificial intelligence (AI) tools to support parts of the hiring process, such as reviewing applications, analyzing resumes, or assessing responses." To the extent that such postings are tagged as including an AI skill by Lightcast, it would tend to overstate AI adoption.

To check the validity of our AI exposure measure, we compare it to AI adoption rates from the Census' Business Trends Outlook Survey (BTOS). The BTOS survey asks firms whether they have used AI in any business function over the past two weeks [Bonney et al., 2026] It is therefore closer to a ground truth measure of AI use, but unlike the Lightcast data, is only reported at the 3-digit NAICS level and has no information on adoption by occupation. To compare the two datasets, we therefore aggregate our Lightcast adoption data to the industry definitions used in the BTOS. Figure 2 reveals a robust positive correlation between the two AI adoption metrics. This is reassuring, as it establishes that our Lightcast-based adoption metric captures meaningful variation in actual AI use by businesses.

3.2 Current Population Survey

We assign each individual in the CPS an AI exposure based on their reported occupation and the indices defined in Section 3.1. When CPS occupation categories include more than one detailed O*NET occupation, we use their average exposure score. We additionally harmonize CPS occupations—which were revised to reflect the 2018 Standard Occupational Classification starting in 2020:M1—to definitions that are fully comparable from 2011:M1 to 2026:M5. For the stock and flows data of labor force status, we plot 12 month moving averages to smooth through some of the month-to-month noise and to handle seasonality.

3.3 Imputation of Missing AI Exposure and Adoption

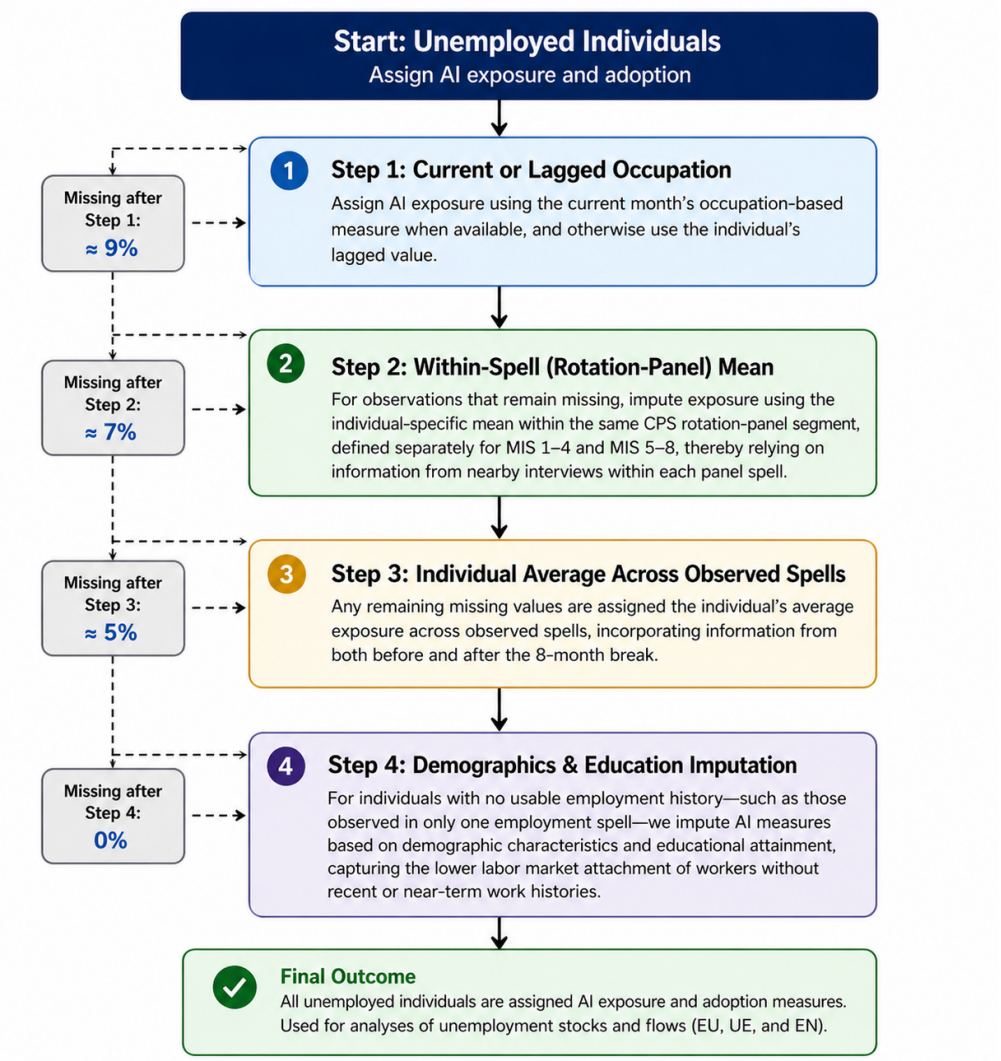
AI exposure and adoption are task-based measures and can therefore be assigned only to individuals with available occupation information. Occupation is observed for all employed workers. Unemployed individuals who were previously employed also report their past occupation, and a small fraction of nonparticipants report prior occupations as well. However, some individuals who transition from employment in the previous month to unemployment in the current month lack occupation information in the current month, likely due to measurement error. In addition, many nonparticipants who were employed in the previous month do not report their prior occupations once out of the labor force. Finally, individuals without work histories generally lack occupation information altogether.

Since our primary interest lies in the dynamics of the unemployment rate and unemployment flows, we focus on imputing AI exposure and adoption for unemployed individuals, as well as for the EU and UE flows. To maximize the use of each individual's reported AI exposure while minimizing imputation, we implement the following procedure. The same approach is applied to EN flows as well. Figure 3 summarizes the imputation procedure.

1. **Step 1**: We first assign AI exposure using the current month's occupation-based measure when available, and otherwise use the individual's lagged value. There are about 9% of unemployed individuals who do not have AI information after this step.

2. **Step 2**: For observations that remain missing, we impute exposure using the individual-

Figure 3: IMPUTATION PROCEDURE



specific mean within the same CPS rotation-panel segment, defined separately for MIS 1–4 and MIS 5–8, thereby relying on information from nearby interviews within each panel spell. After this step, about 7% of the individuals do not have AI information.

3. **Step 3**: Any remaining missing values are assigned the individual's average exposure across observed spells, incorporating information from both before and after the 8-month break. After this step, about 5% of the individuals do not have AI information.

4. **Step 4**: For individuals with no usable employment history—such as those observed in only one employment spell—we impute AI measures based on demographic characteristics and educational attainment, capturing the lower labor market attachment of workers without recent or near-term work histories. After this step, all unemployed individuals are assigned an AI exposure.

It is important to note that this procedure yields unemployment stocks and flows that are largely internally consistent. We consider two cases—with and without Step 4.[8] Omitting the final step and excluding unemployed individuals without occupation information does not materially affect the main results.[9]

3.4 Vacancies

We compute vacancies by occupation using data from JOLTS and Lightcast. We start with the level of JOLTS openings by 2-digit industry i in month t, denoted V_{it}. Because JOLTS does not contain information on occupations, we allocate industry-level openings to occupations in proportion to the share of Lightcast job postings for occupation o in industry i, denoted L_{oit}. Finally, we sum across industries to recover the count of vacancies by occupation,

[8] After step 3 of the adjustment, 99 percent of those without occupation information are new entrants to the labor force, and about 55 percent new entrants do not have occupation information and hence AI exposure or adoption scores. We impute AI exposure and adoption for for these individuals in step 4. Overall, the flows and stock data are robust to the final imputation. Individuals missing AI scores account for about 0.2 pp of the total 0.7 pp rise in the unemployment rate between 2022:M11 and 2026:M5, reflecting the increase in new entrants' unemployment rate over this period. The composition of the high-exposure or high-adoption group among new entrants does not change materially around the step 4 imputation, and the unemployment rates of new entrants with and without work experience both rose by similar amounts. After step 4 imputation, the share of the low-exposure/low-adoption group increases, further reducing the share of the high-exposure/high-adoption group. Overall, individuals with missing AI scores are unlikely to create significant bias in our analyses, regardless of whether they are included or excluded.

[9] We then extend this methodology to assign AI scores to individuals in nonparticipation (Appendix Section D).

$$V_{ot} = \sum_{i \in I} \frac{L_{oit}}{\sum_{o \in O} L_{oit}} \times V_{it}^{JOLTS}. \tag{3.1}$$

This imputation produces a series of vacancies by occupation that is fully consistent with the level of job openings from JOLTS when aggregated.

3.5 Categorization of AI Exposure and Adoption

We classify workers into distinct categories so as to identify those most heavily affected by LLMs and to capture differential effects of LLMs on labor market outcomes.

To begin with, AI exposure, defined as the share of current occupation's tasks exposed to AI, ranges between 0 and 100%. We consider four exposure groups based on the share of current occupation's tasks exposed to AI: below 25%, 25–49%, 50–74%, and 75% or higher. Note that this measure differs from employment-weighted quartiles. For instance, while the 75–100% bin contains approximately 20 occupations, these account for less than 25% of total employment. Instead of using employment-weighted quartiles, we adopt this approach because AI may have had particularly large effects on changes in within-job activities and job switches among the most exposed occupations such as software developers.

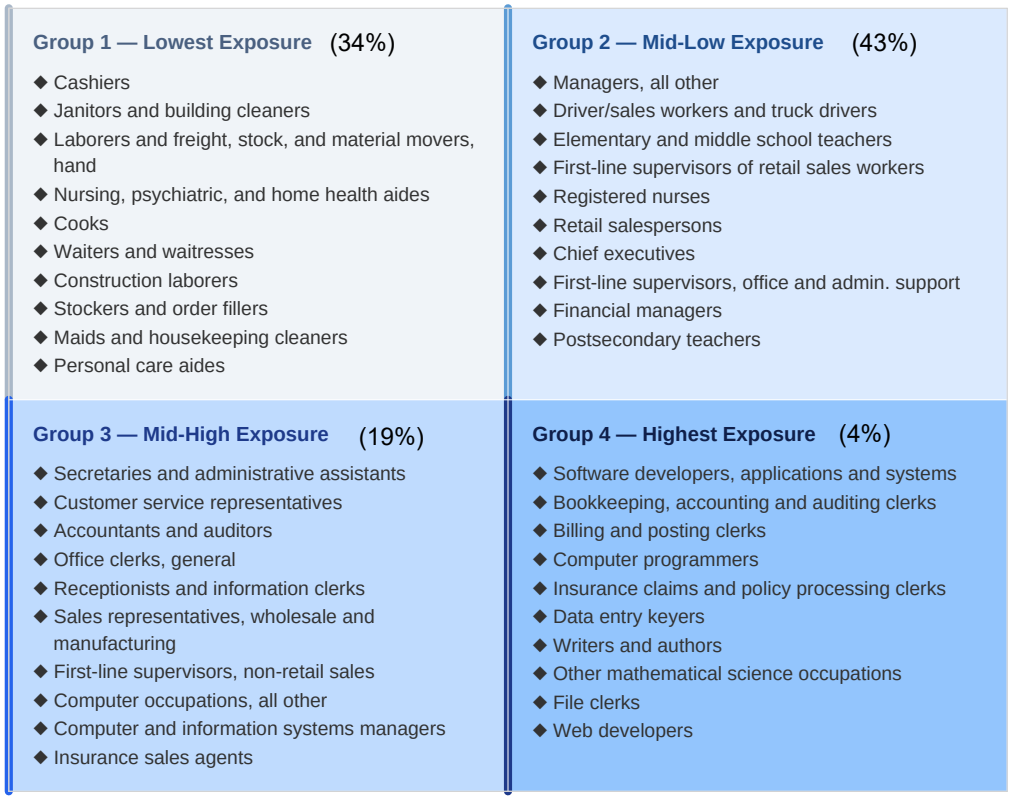
Next, we consider the joint categorization of four groups based on whether their 2025 AI exposure and adoption rates are above or below the mean.[10] The purpose of this approach is to examine the interaction effects of AI exposure and adoption on labor market outcomes.[11]

Figure 4 presents the top ten occupations by AI exposure score, ranked by employment weight. Group 1 consists of occupations requiring physical or interpersonal tasks—janitors, construction laborers, cooks, and personal care aides. Group 2 is more heterogeneous, mixing managerial and professional roles (chief executives, financial managers) with frontline service workers (truck drivers, registered nurses). Group 3 shifts toward clerical and sales work—secretaries, office

[10] For AI exposure, considering the longer sample period starting in 2012 does not change the result. For AI adoption, the median of AI adoption observed since 2012 also yields a similar result. We also consider alternative choices—specifically a higher cutoff for AI adoption—in this two-by-two classification to examine potential nonlinear effects of AI adoption on job-findings and job-losses rates. The results are qualitatively similar overall.

[11] We also consider four quartiles of AI adoption based on the share of AI-related terms in Lightcast job postings in 2025 (Appendix Section A): below the first quartile, between the first and second quartiles, between the second and third quartiles, and the top quartile. We focus specifically on 2025 because AI adoption rose sharply during 2024 and 2025. Before this period, adoption rates remained low, resulting in a highly skewed distribution for the entire sample period. (See Table A.2 for the list of representative occupations/industries.)

Figure 4: REPRESENTATIVE OCCUPATIONS BY AI EXPOSURE



Notes: This figure reports the ten most representative occupations in each of the four AI exposure groups and the exposure score is reported in Table A.1. Numbers in parentheses are the group's share in the labor force. **Source**: Authors' calculation.

clerks, receptionists, and insurance agents. Group 4 is dominated by workers whose core tasks involve processing text, code, or data: software developers, programmers, data entry keyers, and billing clerks. Notably, Group 4 includes both highly skilled tech workers and routine clerical occupations.

Figure 5 shows the largest occupation–industry cells within each exposure–adoption quadrant. The low-exposure/low-adoption (LL) group is concentrated in food service, construction, and education—waiters, cooks, carpenters, and school teachers. This LL group reflects the physically grounded or interpersonally intensive nature of the work, while the high-exposure/high-adoption (HH) group is dominated by knowledge-intensive professionals in legal services, finance, technology and higher education—software developers, lawyers, financial managers and postsecondary teachers. Since HH and LL workers constitute roughly three-quarters of the labor force, task structure and actual AI use largely move together.

Figure 5: REPRESENTATIVE OCCUPATIONS BY AI EXPOSURE/ADOPTION GROUP



Notes: This figure reports the ten most representative occupations in each of the four AI exposure/adoption groups and the exposure score is reported in Table A.3. Numbers in parentheses are the group's share in the labor force.
Source: Authors' calculation.

The off-diagonal groups are also notable. The low-exposure/high-adoption (LH) group—printing press operators, electrical power-line installers, social workers, and entertainment workers—registers low exposure yet high adoption. This category captures occupations where employers are in-

creasingly adopting AI despite their physical or in-person nature, with ambiguous implications for whether AI augments tasks or destroys jobs.

The high-exposure/low-adoption (HL) group—registered nurses, real estate agents, construction managers, and retail supervisors—comprises occupations whose tasks are conducive to AI but where employers have not adopted AI commensurately. This may reflect sector-specific barriers such as features of production. Although some occupation–industry pairs in off-diagonal cells may reflect measurement error in AI exposure and adoption, the off-diagonal groups cluster systematically by industry and capture imperfect sorting between AI exposure and adoption.

Notably, the education sector illustrates how the joint classification captures meaningful heterogeneity even within a single occupational category. Elementary and secondary school teachers fall in LL, reflecting both low adoption and relatively low exposure, while postsecondary teachers land in HH, consistent with greater institutional access to AI tools in university settings. Counselors in colleges and universities fall in the LH group—low exposure but high adoption—suggesting AI is deployed to support administrative and advising workflows without substituting for the in-person counseling task. Teacher assistants fall in HL—meaningfully exposed to AI through tasks such as grading and tallying exams, yet low adoption reflects the in-person, hands-on nature of their work. In sum, jointly classifying workers by AI exposure and adoption reveals heterogeneity that a single-dimensional measure would mask, including meaningful differences in displacement risk and task augmentation even among occupations that appear similar along either dimension alone.

3.6 Observable Attributes of Workers by AI Exposure and Adoption

We briefly discuss observable attributes of workers by AI exposure and adoption.[12] Workers with higher AI exposure and adoption are substantially more educated, older, and better paid.[13] The HH group is the most educated and least young.

Labor market status and outcomes also vary across AI groups. Overall, telework, earnings, and hours rise with AI exposure and adoption.[14] Job tenure is non-monotonic, peaking in the

[12] See Appendix Section F for more details.

[13] The education gradient is steeper along the adoption dimension.

[14] The pre-COVID telework gradient confirms that remote work was already concentrated in high-AI occupations before the pandemic, as also discussed by[Lambert and Yannick, 2026]. Despite the correlation, we find that

middle quartiles and declining at the top, pointing to greater job switching in the AI-intensive occupations. The LH group also exhibits the highest self-employment rate, perhaps reflecting independent contractors in occupations with low AI exposure where AI tools have nevertheless been widely integrated.

After the introduction of LLMs, the composition of unemployment shifted. The share of college graduates among the unemployed rose with AI exposure and adoption, as did the share of prime-age workers in higher exposure/adoption groups — suggesting that LLMs' effects are not confined to young college graduates. Indeed, the composition of permanent job losers aged 25–54 and 55–64 shifted toward the high-exposure/high-adoption group, more so among college graduates. A similar compositional shift toward the high-exposure group is observed among new entrants aged 25–54, particularly those with some college, an associate degree, or a college degree.[15]

Metro areas with high AI exposure tend to have high AI adoption, though local factors — industry specialization, workforce composition, and institutional context — can drive gaps between the two. Since HH jobs are geographically concentrated in large metro areas, reallocation pressure may vary considerably across local labor markets.[16]

4 Stocks and Flows by AI Exposure and Adoption

4.1 Stock and Flows of Unemployment and Employment

This section reports the estimated stocks and flows of workers by AI exposure and AI adoption.

Unemployment rate, job finding and job losses Table 1 reports mean unemployment rates and flow rates — employment-to-unemployment (EU) and unemployment-to-employment (UE) by AI exposure and their joint classification. First, as shown in Panel A, higher AI exposure is generally associated with lower unemployment rates, primarily reflecting lower separation rates.

teleworkability has different effects on job loss and job finding than AI exposure and adoption (Appendix Section C.1).

[15] This pattern is robust to the full imputation.

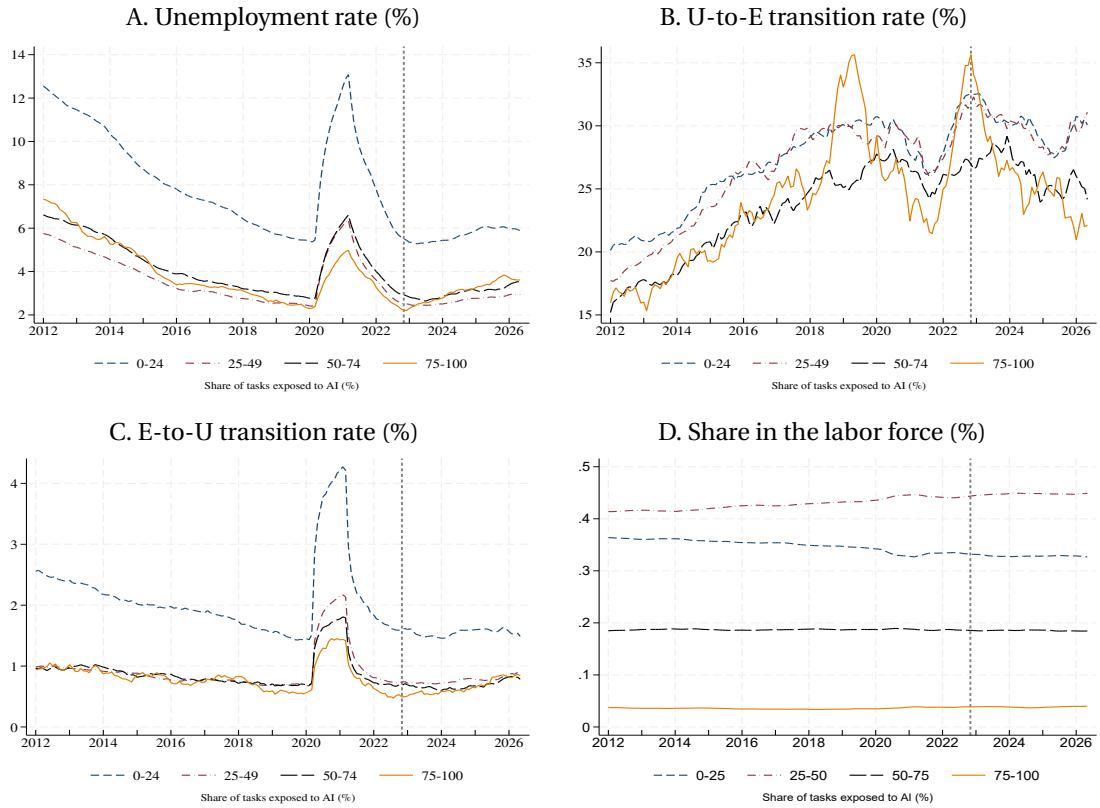
[16] Spatial mismatch and geographic mobility frictions could slow reallocation, though thick local labor markets may ease job transitions for some workers.

Table 1: AVERAGES OF UNEMPLOYMENT RATE AND SELECTED FLOWS RATES (2012:M1 – 2026:M5)

Group	EU rate	UE rate	Unemployment rate	Job Switch rate	Task Switch rate
Panel A: AI Exposure					
Exposure 1	2.0	27.7	7.6	2.7	1.1
Exposure 2	0.9	27.3	3.4	1.9	1.1
Exposure 3	0.8	24.0	3.9	2.0	1.1
Exposure 4	0.8	24.4	3.7	1.9	1.1
Panel B: Exposure × Adoption					
Low–Low	1.8	28.5	6.7	2.6	1.1
Low–High	1.5	25.3	6.2	2.4	1.1
High–Low	0.8	26.4	3.2	1.9	1.1
High–High	0.7	22.9	3.4	1.9	1.1

Note to table: The exposure groups are defined by the share of tasks in a worker's current occupation that are exposed to AI: below 25%, 25–49%, 50–74%, and 75% or higher. The exposure × adoption groups are defined jointly by AI exposure and adoption—above versus below the mean of AI exposure and AI adoption in 2025. **Source**: Authors' calculation.

Figure 6: UNEMPLOYMENT RATE, UE AND EU RATES AND LABOR FORCE SHARE (AI EXPOSURE)



Notes: This figure displays the unemployment rate (Panel A), UE rate (Panel B), EU rate (Panel C), and labor force shares (Panel D) for four AI exposure groups (12-month moving average). Blue dashed line: 0–25; red dashed-dot line: 25–50; black long-dashed line: 50–75; solid yellow line: 75–100. **Source**: Authors' calculation

Meanwhile, higher AI exposure is associated with lower job-finding probabilities, likely reflecting high-skilled workers' lengthy job-search time. Workers with low AI exposure tend to have less stable employment and higher turnover, which raises their job-finding rates accompanied by high job-loss rates, as well as the job-switch rate. Overall, a similar pattern is found in the joint classification (Panel B). The high-exposure/high-adoption group shows the lowest EU rate, UE rate and job-switch rate, while low-exposure/low-adoption group has the highest EU rate and UE rate, as well as job switch rate. Noticeably among the same exposure group, higher exposure tends to lower EU rate as well as UE rate perhaps reflecting less churns in the high-adoption jobs on average.

Figure 6 displays the time series by AI exposure.[17] Several patterns emerge. First, the unem-

[17]We focus on UE and EU flows since our primary focus is changes in the unemployment rate. Reallocation is also measured with flows per capita concept—the monthly rate of labor force transition, which includes flows associated with nonparticipation (N). Since N-to-E flows, U-to-N and N-to-U flows require substantial imputation

ployment rate of the top AI-exposure group rose noticeably (1.4 percentage point, yellow line in Panel A) after the introduction of LLMs (dotted vertical line at the end of 2022), increasing more than other groups (0.4 pp in the lower exposure groups, and 0.6 pp in 50-75% group).[18] This is mainly driven by the hiring margin: The job-finding probabilities of workers with the highest AI exposure declined persistently and more than those of other groups (yellow line in Panel B) following the introduction of LLMs. However, the separation rate of top exposure group also rose marginally faster than the other groups, but the difference is not as noticeable as in the job-finding probabilities (yellow line in the left figure of Panel C). The labor force shares of the four exposure group are 34%, 43%, 19% and 4%, from bottom to top. The labor force shares of the top two exposure groups are little changed, while that of the bottom group has declined steadily and that of the second-lowest group has risen gradually (Panel D).

In addition, it is important to note that the unemployment rates of lower-exposure groups are more cyclically sensitive than those of higher-exposure groups, particularly the top. In this sense, the faster increase in the unemployment rate of workers with high AI exposure likely reflects factors specific to them such as the introduction of LLMs. The larger rise in unemployment, steeper decline in job-finding rates, and a little faster separation rates therefore reflect a disproportionate decline and shift in labor demand for high-AI-exposure workers following the introduction of LLMs.

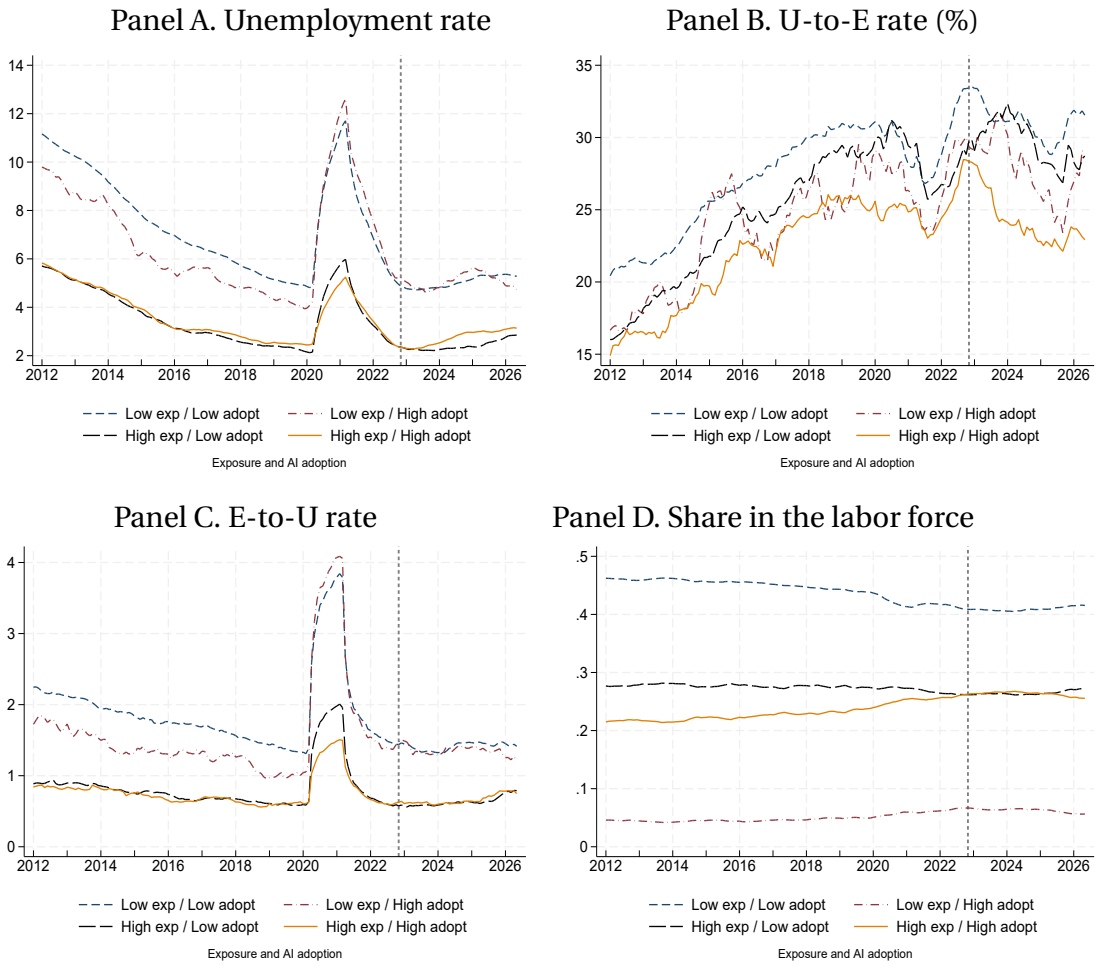
Figure 7 displays the analogous data for four groups defined jointly by AI exposure and adoption. The high-exposure/high-adoption group shows patterns similar to the top exposure group. Specifically, the unemployment rate of the HH group rose by slightly below 1 percentage point between end-2022 and May 2026, while the unemployment rate of the low-exposure/low-adoption group rose by 0.4 percentage point (Panel A).[19] The unemployment rate of high-exposure/low-adoption group rose by about half percentage point, while that of the low-exposure/high-adoption edged lower. The HH and HL groups comprise about a quarter of

for individuals without occupation information or work histories, we do not consider these flows in our baseline analysis. However, we provide these flows in Appendix Section D and the flows per capita in Figure A.7.

[18] With the fully imputed data, the unemployment rate of bottom exposure group rises more (+0.9 pp) due to the increased unemployment rate of new entrants, but still is below the rise of top exposure group. The overall patterns are robust to the full imputation (Appendix Section B).

[19] Using a higher threshold (85th percentile instead of the mean) to define the HH group—which covers roughly 20 percent of the labor force—the unemployment rate increase is somewhat larger at 1.2 percentage points (Figure A.4).

Figure 7: UNEMPLOYMENT RATE AND FLOWS BY AI EXPOSURE AND ADOPTION



Note: This figure displays the unemployment rate (Panel A), UE rate (Panel B), EU rate (Panel C), and labor force shares (Panel D) for four AI exposure/adoption groups (12-month moving average). "Exp" stands for exposure, and "adopt" stands for adoption. Blue dashed line: Low-exp/Low-adopt; red dashed-dot line: Low-exp/High-adopt; black long-dashed line: High-exp/Low-adopt; solid yellow line: High-exp/High-adopt. **Source:** Authors' calculation.

the labor force, and the LL group's share is slightly above 40% and the LH group less than 10%.

The increase in the unemployment rate for the HH group is primarily driven by a steady decline in the UE rate from the end of 2022, a pattern that is much less pronounced in the other groups (yellow line in Panel D). Meanwhile, the EU rates edged up among the high-exposure groups in 2025 and 2026, but overall do not exhibit pronounced differences across the four groups. This pattern confirms that the increase in the unemployment rate for the high–exposure/high–adoption group is driven mainly by a decline in the UE rate.[20]

Among high-exposure workers, higher adoption is associated with a more adverse hiring margin. The job-finding rate of the HH group is lower and declined more than that of the HL group, contributing to a faster increase in the unemployment rate of the HH group relative to the HL group. This suggests that AI adoption has differential effects among high-exposure workers, and that the interaction between AI exposure and adoption can be important for understanding LLMs' effects on labor market outcomes.

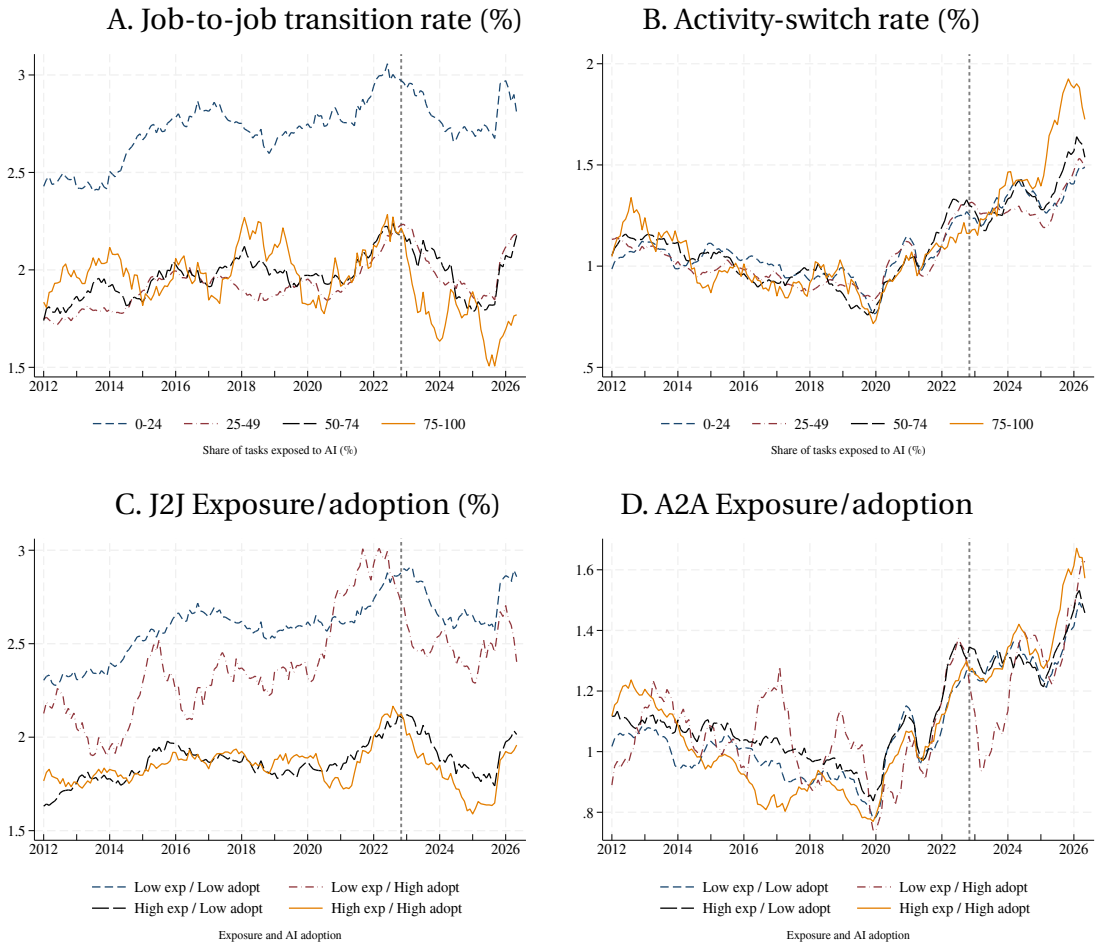
Switches of employers and job activities Figure 8 presents the job-switch rate (Panel A) and the activity-switch rate (Panel B) by AI exposure and adoption. The CPS asks respondents whether they, if employed in two consecutive months, changed employer since last month, which is used to compute the job-to-job transition rate. The CPS also asks respondents who have not changed an employer whether they experienced changes in their work activities, which is used to compute the activity switch rate.

Consistent with reduced hiring among workers with high AI exposure and adoption, job-switch rates—another indicator of labor demand—declined more for high-exposure/adoption workers than for other groups. Job-switch rates fell across all groups considered, but the decline is more pronounced among the top exposure group and the HH group. Although overall job-switch rates began to recover in the second half of 2025, they continued to decline for the top-exposure group.

Meanwhile, the activity-switch rate paints a markedly different picture of LLMs' reallocation. Across the four exposure groups, activity-switch rates are similar through 2024. However, starting around 2025, the highest-exposure group exhibits a noticeably faster increase in activity-switch

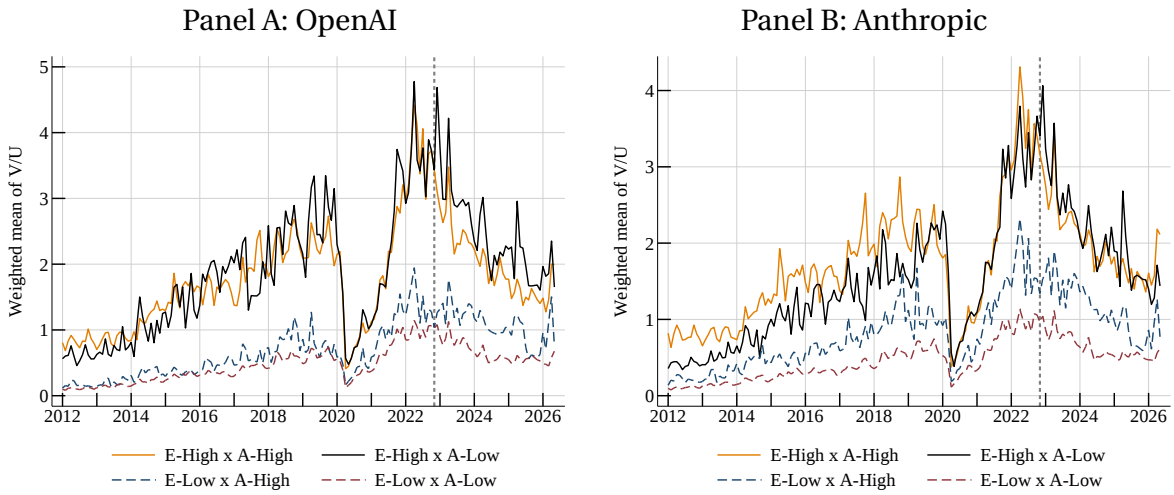
[20]The overall conclusion is similar once we break the data by the quartile of AI adoption. See Appendix Section B for more details.

Figure 8: SWITCHES OF JOBS AND ACTIVITIES BY AI EXPOSURE



Notes: This figure displays job-to-job transition rates by AI exposure (Panel A), activity-switch rates by AI exposure (Panel B), job-to-job transition rates by AI exposure and adoption (Panel C), and activity-switch rates by AI exposure and adoption (Panel D) (12-month moving average). For Panels A and B: blue dashed line: 0–25; red dashed-dot line: 25–50; black long-dashed line: 50–75; yellow solid line: 75–100. For Panels C and D: blue dashed line: Low-Exp/Low-Adopt; red dashed-dot line: Low-Exp/High-Adopt; black long-dashed line: High-Exp/Low-Adopt; yellow solid line: High-Exp/High-Adopt. **Source**: Authors' calculation

Figure 9: MARKET TIGHTNESS BY AI-EXPOSURE AND ADOPTION



Notes: Panels A and B display market tightness by AI exposure and adoption, with the AI exposure measure constructed using the OpenAI and Anthropic data, respectively. Blue dashed line: Low-Exp/Low-Adopt; red dashed-dot line: Low-Exp/High-Adopt; black long-dashed line: High-Exp/Low-Adopt; yellow solid line: High-Exp/High-Adopt. **Source:** Authors' calculation.

rates than the other groups (Panel C). It is less pronounced when using the coarse joint categorization, but the HH group's activity switch rate picked up more than others (Panel D). This observation suggests that the task reorganization is concentrated in the top AI exposure and HH groups. In sum, LLM-driven reallocation has increasingly taken place within establishments and within current jobs.

Overall, LLMs appear to have reduced reallocation via job-to-job transitions while raising internal task reorganization. This reflects weaker demand for HH workers in a low-hiring environment and a buildup of reallocation pressure absorbed through task reorganization within jobs.

4.2 Market Tightness by AI Exposure and Adoption

We compute market tightness for four groups of workers defined by above- and below-mean AI exposure and adoption (LL, LH, HL, and HH). We consider the joint categorization incorporating both the supply and demand sides of AI intensity, as market tightness captures both dimensions of the labor market. Figure 9, Panel A reports estimates based on OpenAI's AI exposure measure and Panel B uses the Anthropic measure.

In general, workers with high AI exposure exhibit higher market tightness, consistent with their lower unemployment rates.[21] However, a notable divergence emerges over time. Market tightness for high-exposure workers—particularly those with both high exposure and high adoption—has remained below its pre-COVID level since 2023, while market tightness for low-exposure workers has largely returned to around its pre-COVID level. This suggests that the evolution of tightness across groups cannot be fully explained by aggregate labor-market conditions.

Instead, market tightness among high-exposure workers points to factors associated with AI and LLMs. While these workers continue to experience relatively favorable labor market conditions in cross section, the persistent shortfall relative to pre-COVID levels indicates weaker vacancy creation in high-exposure occupations. Potential explanations for this phenomenon include firms adjusting labor demand through task reorganization and partial substitution of labor with LLM technologies. At the same time, it could also reflect an interaction between macroeconomic uncertainty and uncertainty about the progress of AI models and their ability to substitute for labor, even among firms that have yet to widely integrate these technologies.

The evidence from market tightness since the end of 2022 is broadly consistent with patterns in unemployment flows, job-switching, and activity-switching rates. This points to uneven labor market effects of LLMs and a potential buildup of reallocation pressure, whereby workers are unable to move to sectors with better reemployment prospects due to mismatch or other frictions. That said, the recent recovery in job-finding and job-switching rates for high-exposure high-adoption groups shown in Figure 8, along with an increase in market tightness among these workers could signal some improvement in labor demand in the first half of 2026. This, in turn, could reflect reduced uncertainty about both the macroeconomic environment and the capabilities and limitations of LLMs. To the extent that the low-hiring, low-firing environment may itself be an endogenous response to uncertainty, policies that reduce uncertainty may help alleviate reallocation pressure and facilitate more efficient labor market adjustment.

[21] This pattern is driven by lower separation rates of high-exposure workers, notwithstanding their lower job-finding rates due to lower matching efficiency.

4.3 Wage Growth and Hours Worked

We discuss effects of AI exposure and adoption on wage growth and hours worked, as these indicators also measure labor demand.[22]

A. Wage growths This section analyzes the effects of AI factors on wage growth in the post-LLM era. We measure individual-level wage growth as 12-month changes in log hourly earnings, following the approach of the Atlanta Fed Wage Tracker.[23]

Consider a regression model for the wage growth:

$$\begin{aligned}\pi_{it} = {} & \beta_1 AIX_{it} + \beta_2 \left(AIX_{it} \times \text{Post-LLM}_t\right) \\ & + \beta_3 AIA_{it} + \beta_4 \left(AIA_{it} \times \text{Post-LLM}_t\right) \\ & + \beta_5 \theta_{it} + \beta_6 age_{it} + \beta_7 age_{it}^2 \\ & + \gamma' X_{it} + \lambda_t + \delta_i + \varepsilon_{it},\end{aligned} \tag{4.1}$$

where π_{it} is the 12-month nominal hourly wage growth of individual i at time t, θ_{it} denotes market tightness of the individual's industry at t, age_{it} and age_{it}^2 are the individual's age and its square, and X_{it} includes gender and education controls.[24]

Columns (1) and (2) of Table 2 reports the regression results. Both AI exposure and adoption raise wage growth, but their interactions with the post-LLM period are negative and statistically significant. Workers in more AI-exposed occupations and higher-adoption industries experience faster wage growth, perhaps reflecting productivity gains, though this effect attenuates after the introduction of LLMs. While AI initially complements high-exposure workers and boosts their wages, the diffusion of LLM technologies may compress wage growth by facilitating task substitution or reducing demand for these workers. The inclusion of time fixed effects and industry-level market tightness suggests that these negative effects are unlikely to reflect changes macroeconomic conditions or sector-specific developments other than AI. Nonetheless, the net effects of AI exposure and adoption on wage growth remain positive.

[22] Effects of AI intensity on the level of wages are reported in Section C.2.

[23] See `https://www.atlantafed.org/research-and-data/data/wage-growth-tracker` for more details.

[24] For wage growth, we observe only one change when an individual reports hourly wages in MIS 4 and 8. Earnings weights (`EARNWT`) are used for all wage regressions.

Table 2: DETERMINANTS OF WAGE GROWTH

	(1)	(2)	(3)	(4)
	Time FE	Time+Ind FE	JtJ Baseline	JtJ Post-LLM
AI Exposure (std.)	0.794***	0.785***	0.560***	0.650***
	(0.0972)	(0.105)	(0.0980)	(0.108)
AI Adoption (std.)	2.269***	1.822***	0.419	1.409***
	(0.593)	(0.564)	(0.266)	(0.537)
AI Exposure × Post-LLM	-0.399**	-0.412**		-0.497**
	(0.199)	(0.199)		(0.205)
AI Adoption × Post-LLM	-1.850***	-1.569**		-1.229**
	(0.675)	(0.646)		(0.620)
AI Exposure × Job-to-job			2.291***	1.242
			(0.702)	(0.837)
AI Adoption × Job-to-job			1.616	13.93*
			(2.972)	(7.921)
AI Exposure × Job-to-job × Post-LLM				3.275*
				(1.753)
AI Adoption × Job-to-job × Post-LLM				-16.74**
				(8.475)
Market Tightness (θ_t)	0.0125**	0.0101*	0.0130**	0.0125**
	(0.00586)	(0.00589)	(0.00618)	(0.00617)
Age	-0.265***	-0.407***	-0.392***	-0.391***
	(0.0315)	(0.0331)	(0.0345)	(0.0345)
Age^2	0.00181***	0.00321***	0.00307***	0.00306***
	(0.000357)	(0.000372)	(0.000386)	(0.000386)
Female	-1.765***	-1.318***	-1.192***	-1.180***
	(0.172)	(0.196)	(0.202)	(0.202)
High School	1.567***	1.292***	1.112***	1.104***
	(0.271)	(0.275)	(0.287)	(0.287)
Some College / Associate	1.333***	0.961***	0.872***	0.857***
	(0.277)	(0.284)	(0.296)	(0.296)
College Graduate	2.394***	1.875***	1.694***	1.674***
	(0.330)	(0.341)	(0.352)	(0.352)
Observations	301209	301209	272286	272286
R-squared	0.00504	0.00754	0.00738	0.00750

Notes: Standard errors in parentheses. JtJ stands for job-to-job transitions.

Time fixed effects included in all columns.

Industry fixed effects included in columns (2)–(4).

Standard errors clustered at individual level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

We further analyze the differential effects of AI exposure and adoption on job switchers and incumbents using Columns (3) and (4) of Table 2, without and with post-LLM interactions respectively. Workers with higher AI exposure experience stronger wage gains when switching jobs. While LLMs compressed wage premia for both AI exposure and adoption among incumbents, job switchers with higher AI exposure continued to experience stronger wage growth. In contrast, higher AI adoption dampens wage growth even among job switchers in the post-LLM era. These opposing effects suggest growing wage disparities across the gradation of AI exposure and adoption, aligning with Freund and Mann [2026], who find widening wage gaps between job switchers and incumbents among AI-exposed workers, reflecting reallocation toward more productive tasks at new employers.

Hours worked Next, we examine AI's effects on hours worked. We use the same regression specification as equation (4.1), replacing wage growth with hours worked as the dependent variable. Table 3 reports the results.[25]

Overall, AI exposure is positively and statistically significantly associated with hours worked (line 1), indicating that workers in more AI-exposed occupations tend to supply more labor and work longer hours. In contrast, AI adoption has a negative baseline effect on hours worked (line 3), consistent with automation reducing labor input.

Importantly, these effects again attenuate during the LLM period (2023 onward), as reflected in the statistically significant offsetting interaction terms (lines 2 and 4). In particular, the positive interaction between AI adoption and the LLM period (line 4) may signal a shift from labor-substituting to more labor-complementary effects along the adoption margin, while the labor-supply effect associated with AI exposure diminished (line 2). Overall, these empirical results suggest reduced demand for high-exposure workers, while LLM technologies appear to have altered how AI exposure and adoption translate into labor input decisions.

In sum, AI's upside effects on nominal wages and hours worked attenuated from the introduction of LLMs, all suggesting that LLM-driven reallocation has occurred through reconfiguration of labor input alongside a decline in demand for workers with high AI exposure and adoption jobs.

[25] Final weights (`WTFINL`) are used for all hours regressions. We control for month-in-sample to address potential rotation group bias; the results are robust to this correction.

Table 3: DETERMINANTS OF HOURS WORKED

	(1)	(2)
	Time FE	Time + Industry FE
AI Exposure (std.)	0.749***	0.655***
	(0.00997)	(0.0107)
AI Adoption (std.)	-0.292***	-0.307***
	(0.0186)	(0.0187)
AI Exposure × Post-LLM	-0.140***	-0.140***
	(0.0211)	(0.0206)
AI Adoption × Post-LLM	0.164***	0.177***
	(0.0247)	(0.0245)
Market Tightness (θ_t)	0.00648***	0.00441***
	(0.000371)	(0.000364)
Age	0.977***	0.880***
	(0.00454)	(0.00448)
Age^2	-0.0107***	-0.00966***
	(0.0000523)	(0.0000515)
Female	-4.458***	-3.630***
	(0.0185)	(0.0204)
High School	2.627***	2.411***
	(0.0383)	(0.0376)
Some College / Associate	2.302***	2.194***
	(0.0398)	(0.0394)
College Graduate	3.445***	3.575***
	(0.0397)	(0.0403)
Observations	7872212	7872212
R-squared	0.0837	0.107

Standard errors in parentheses

Time fixed effects included in all columns.

Industry fixed effects included in column (2).

Standard errors clustered at individual level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5 Hirability and Separability: AI-driven Reallocation Pressure

This section analyzes the effects of AI exposure and adoption on job finding and job loss. While the preceding analysis documents differences in labor market outcomes across AI exposure and adoption groups, these patterns may merely reflect factors correlated with AI intensity, such as education. To address this issue, we estimate the effects of AI exposure and adoption on individuals' propensity of job finding and job loss, controlling for observable characteristics. Based on these estimates, we construct an individual-level *hirability* index that captures an unemployed individual's latent propensity to find a job. Analogously, we construct an individual-level *separability* index that captures an employed individual's latent propensity to lose a job. Second, using these indices we assess the extent to which AI exposure and adoption account for the dispersion in job-finding and job-loss prospects across individuals.

5.1 Model

Regression models Consider the baseline pooled regression:

$$\begin{aligned} UE_{it} &= \alpha + \beta_1 AIX_{it} + \beta_2 AIA_{it} + \beta_3 AIA_{it}^2 + \beta_4 \theta_{i,t-1} \\ &+ \beta_5 age_{it} + \beta_6 age_{it}^2 + \beta_7 sex_i + \beta_8 edu_i + \beta_9 mis_{it} + \varepsilon_{it}, \end{aligned} \tag{5.1}$$

where UE_{it} is an indicator for transitions from unemployment to employment between $t-1$ and t, equal to 1 if the individual finds a job and 0 if the individual remains unemployed or exits the labor force (including cases with missing status in the following month); mis_{it} denotes the month-in-sample indicator; AIX_{it} measures AI exposure (standardized), while AIA_{it} and AIA_{it}^2 denote the AI adoption rate (standardized) and its squared term, respectively; age_{it} and age_{it}^2 capture age and its quadratic term; sex_{it} indicates gender; and edu_{it} represents education categories (less than high school, high school graduate, some college or associate degree, and college graduate); ε_{it} is the error term. We consider mis_{it} to control potential rotation group bias in the reported employment and unemployment status (Ahn and Hamilton, 2022). We include a squared term for AI adoption to account for the nonlinear relationship between the UE rate and

AI adoption (Figure 6).[26] Importantly, we include $\theta_{i,t-1}$, the lagged market tightness (vacancies-to-unemployment ratio) of the individual's industry, to capture industry-level business-cycle conditions and changes in supply-demand balance, thereby distinguishing the effects of LLMs from those of aggregate and industry-level economic fluctuations.[27] We restrict the sample to the period from 2023 onward to capture the effects of LLMs on job finding and job loss, as well as the consequences of increased AI adoption during this period.[28]

After we estimate the model, we recover the fitted value $\widehat{UE}_{it}$ based on the estimated coefficients for all i and t. We call $\widehat{UE}_{it}$ *"the hirability index."* To assess the contribution of AI exposure and AI adoption, we compare the hirability index with "AI-purged" hirability index $\widehat{UE}_{it}^{wo}$, which we define as follows:

$$\widehat{UE}_{it}^{wo} = \widehat{UE}_{it} - \widehat{\beta_1}\, AIX_{it} - \widehat{\beta_2}\, AIA_{it} - \widehat{\beta_3}\, AIA_{it}^2. \tag{5.2}$$

We also construct analogous measures for EU transitions, defined as an indicator that takes the value one if a worker becomes unemployed, and zero if the worker remains employed, exits the labor force, or becomes missing from the survey. Using the same pooled regression in equation (5.2), we construct the *separability* index, denoted $\widehat{EU}_{it}$. We then apply equation (5.2) to obtain the "AI-purged" separability index, denoted $\widehat{EU}_{it}^{wo}$.

Cross-sectional dispersion over time We then measure the dispersion of the hirability index across individuals over time to capture the reallocation pressure that individuals face. Previous studies (e.g., Lilien [1982] and Jackman and Roper [1987]) adopted dispersion measures to assess the effects of unresolved reallocation pressure on the unemployment rate. Higher dispersion implies greater disparities in reemployment prospects across sectors, reflecting more difficulties in moving to sectors with stronger demand and higher market tightness.

[26] We do not consider the squared term of AI exposure due to its high correlation with the linear term, which would introduce multicollinearity.

[27] Including time fixed effects do not change the result.

[28] We use the individual final weight (WTFINL) for the estimation. We also consider including time fixed effects in the regression, but the result remains robust showing a little wider dispersion in the UE rates across AI exposure and adoption.

The weighted cross-sectional standard deviation is

$$sd_t^w(\text{UE}) = \sqrt{\frac{\sum_i w_{it}\left(\widehat{UE}_{it} - \mu_t^w(UE)\right)^2}{\sum_i w_{it}}}, \tag{5.3}$$

where w_{it} is the individual weight at t and hence the cross-sectional mean of the individual-level hirability index as follows:

$$\mu_t^w(UE) = \frac{\sum_i w_{it}\widehat{UE}_{it}}{\sum_i w_{it}}. \tag{5.4}$$

Analogously, the weighted cross-sectional standard deviation of "AI-purged" hirability index is written as follows:

$$sd_t^{wo}(UE) = \sqrt{\frac{\sum_i w_{it}\left(\widehat{UE}_{it}^{wo} - \mu_t^{wo}(UE)\right)^2}{\sum_i w_{it}}}, \tag{5.5}$$

where the cross-sectional mean of "AI-purged" hirability index is defined as:

$$\mu_t^{wo}(UE) = \frac{\sum_i w_{it}\widehat{UE}_{it}^{wo}}{\sum_i w_{it}}. \tag{5.6}$$

5.2 Estimated Coefficients

Table 4 reports the coefficient estimates of equation 5.2 for job-finding (UE) and job-separation (EU) probabilities. A few observations stand out.

To begin with, AI exposure lower job-finding rate as well as job-loss rates with statistical significance (Line 1, Columns 1 and 3). Workers with high AI exposure are less likely to transition between employment and unemployment, implying that more task exposure to AI can reduce labor market churns on the margin of UE and EU transition rates.

Meanwhile, AI adoption lowers job-finding rates but raises job-loss rates, suggesting reduced job creation and greater job destruction (Line 2 of Columns 1 and 3). However, the squared term shows the opposite signs: unlike the mean effects, at higher levels of adoption, separation rates decline and job-finding rates rise, potentially reflecting positive net effects of LLMs on job creation.[29]

[29]Lagged market tightness is negatively associated with the separation rate but statistically insignificant for job finding. Demographic controls display expected patterns.

Table 4: DETERMINANTS OF HIRABILITY AND SEPARABILITY

	Hirability (UE)		Separability (EU)	
	(1)	(2)	(3)	(4)
	Continuous	2×2 Group	Continuous	2×2 Group
AI Exposure (std.)	-0.0108***		-0.00245***	
	(0.00317)		(0.000133)	
AI Adoption (std.)	-0.0212***		0.000716***	
	(0.00450)		(0.000152)	
AI Adoption (std.)^2	0.00109***		-0.0000325***	
	(0.000366)		(0.00000977)	
Group 2 (LH)		0.0761**		0.000569
		(0.0362)		(0.00146)
Group 3 (HL)		-0.0233***		-0.00444***
		(0.00681)		(0.000244)
Group 4 (HH)		-0.0765***		-0.00416***
		(0.00899)		(0.000288)
Market Tightness (θ_{t-1})	-0.0000510	-0.0000325	-0.0000145***	-0.0000134***
	(0.000118)	(0.000116)	(0.00000230)	(0.00000230)
Age	0.000281	0.000377	-0.000778***	-0.000757***
	(0.000893)	(0.000894)	(0.0000476)	(0.0000477)
Age^2	-0.0000190*	-0.0000195*	0.00000680***	0.00000660***
	(0.0000101)	(0.0000101)	(0.000000496)	(0.000000496)
Female	0.00327	0.00250	-0.000210	-0.000374*
	(0.00535)	(0.00532)	(0.000219)	(0.000219)
High School	-0.0187**	-0.0194**	-0.00419***	-0.00447***
	(0.00839)	(0.00837)	(0.000624)	(0.000623)
Some College / Associate	0.00816	0.00705	-0.00673***	-0.00710***
	(0.00916)	(0.00910)	(0.000614)	(0.000613)
College Graduate	-0.000531	0.000664	-0.00869***	-0.00897***
	(0.00963)	(0.00959)	(0.000599)	(0.000597)
Observations	43778	43778	1232459	1232459
R-squared	0.00550	0.00622	0.00309	0.00306

Standard errors in parentheses

Constant and MIS fixed effects not reported.

Standard errors clustered at individual level.

Sample: 2023–2026.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Next, we examine the effects of the interaction between exposure and adoption on job finding. We find that the negative effects on job finding are largest among the HH group, followed by the HL group (Column 2 in Table 4). However, a similar pattern is not observed in the separation margin. This suggests that the interaction between AI exposure and adoption has a statistically significant negative effect on the hiring margin.[30]

In sum, both AI exposure and adoption reduce job-finding probabilities, with mixed effects on the separation margin. In particular, AI adoption lowers job-finding rates and increases job-loss rates, implying a net job-destructive effect despite its nonlinear effects on both margins. These results suggest that the diffusion of LLMs is associated with the destruction of jobs with tasks highly exposed to AI and with greater AI adoption, with the adjustment operating primarily through the hiring margin, though the separation margin also plays some role.

5.3 Estimated Indices and Their Dispersion

Figure 10 reports the hirability index. Panel A plots the mean hirability index by AI exposure. The index is lower for individuals with higher AI exposure. Moreover, the top exposure group exhibits a pronounced decline after the introduction of LLMs (vertical dashed line), with a particularly sharp drop in 2024, while the lower exposure groups remain relatively stable.

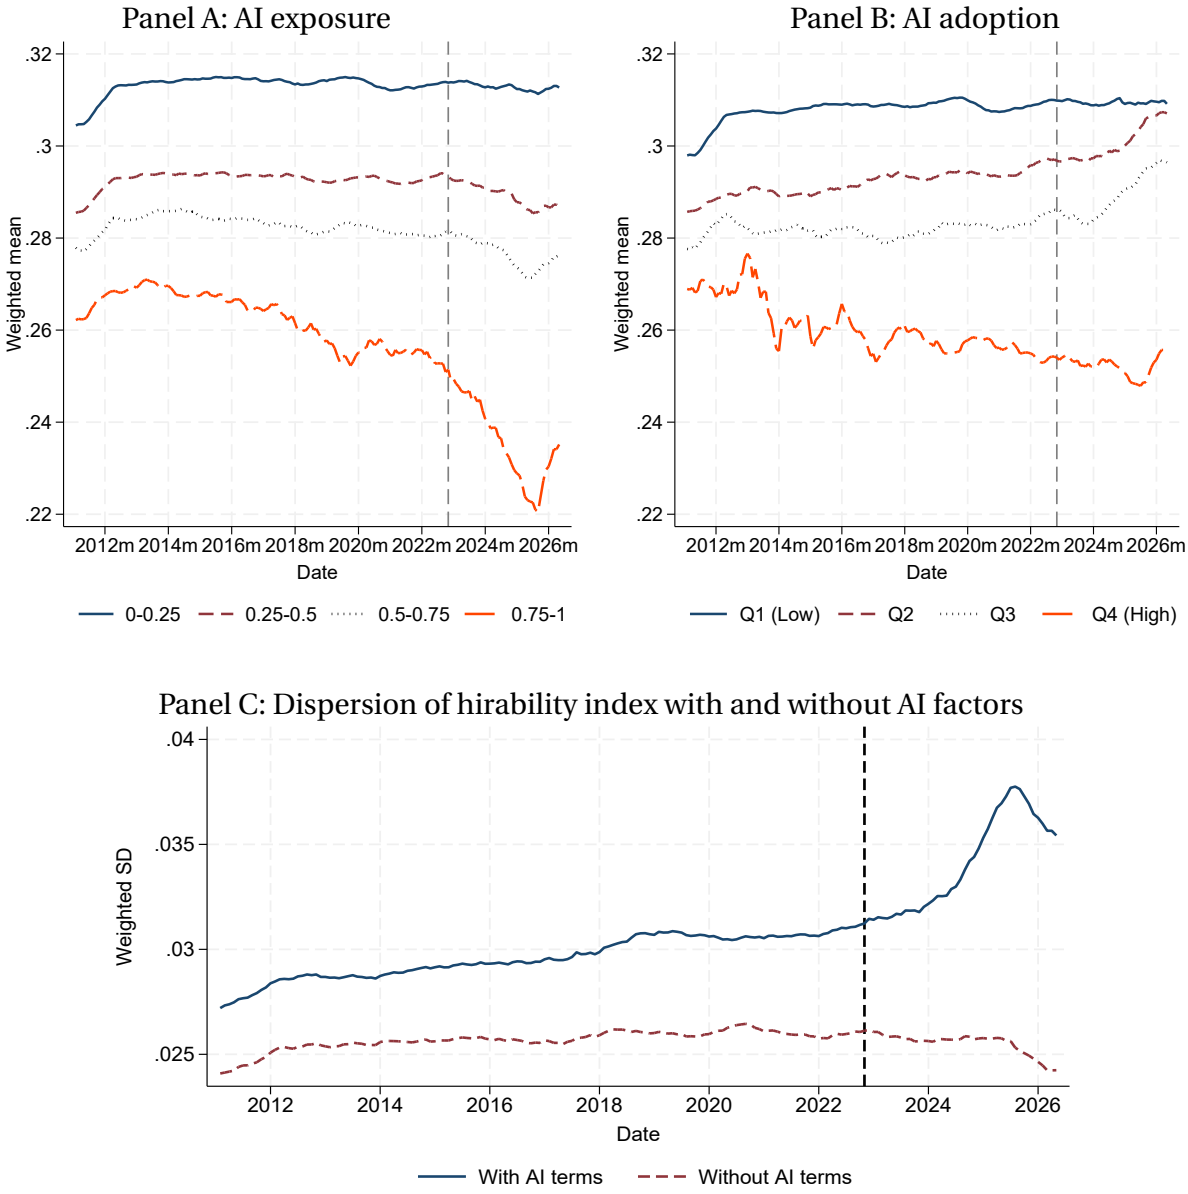
Panel B displays the mean hirability index by AI adoption quartile (2025 based).[31] The index is also lower for individuals in higher AI adoption quartiles. Notably, while the top adoption quartile exhibits an overall decline in hirability including the post-LLM period, the middle two quartiles have risen steadily since 2023. This divergence becomes more pronounced in 2024 and 2025. Meanwhile, the hirability of the top exposure and HH groups recovered somewhat from 2025.

Panel C displays the standard deviation of the hirability index ($sd_t^w(UE)$, solid blue line) and its AI-purged counterpart ($sd_t^{wo}(UE)$, red dashed line). Prior to 2023, $sd_t^w(UE)$ remains largely flat, then rises steadily and accelerates from 2024 onward, indicating growing reallocation

[30] A similar pattern is observed when we interact the four groups of AI exposure with the quartiles of AI adoption.

[31] As briefly mentioned in Section 3, we consider four quartiles of AI adoption based on the share of AI-related terms in Lightcast job postings in 2025: below the first quartile, between the first and second quartiles, between the second and third quartiles, and the top quartile. We focus specifically on 2025 because AI adoption rose sharply during 2024 and 2025.

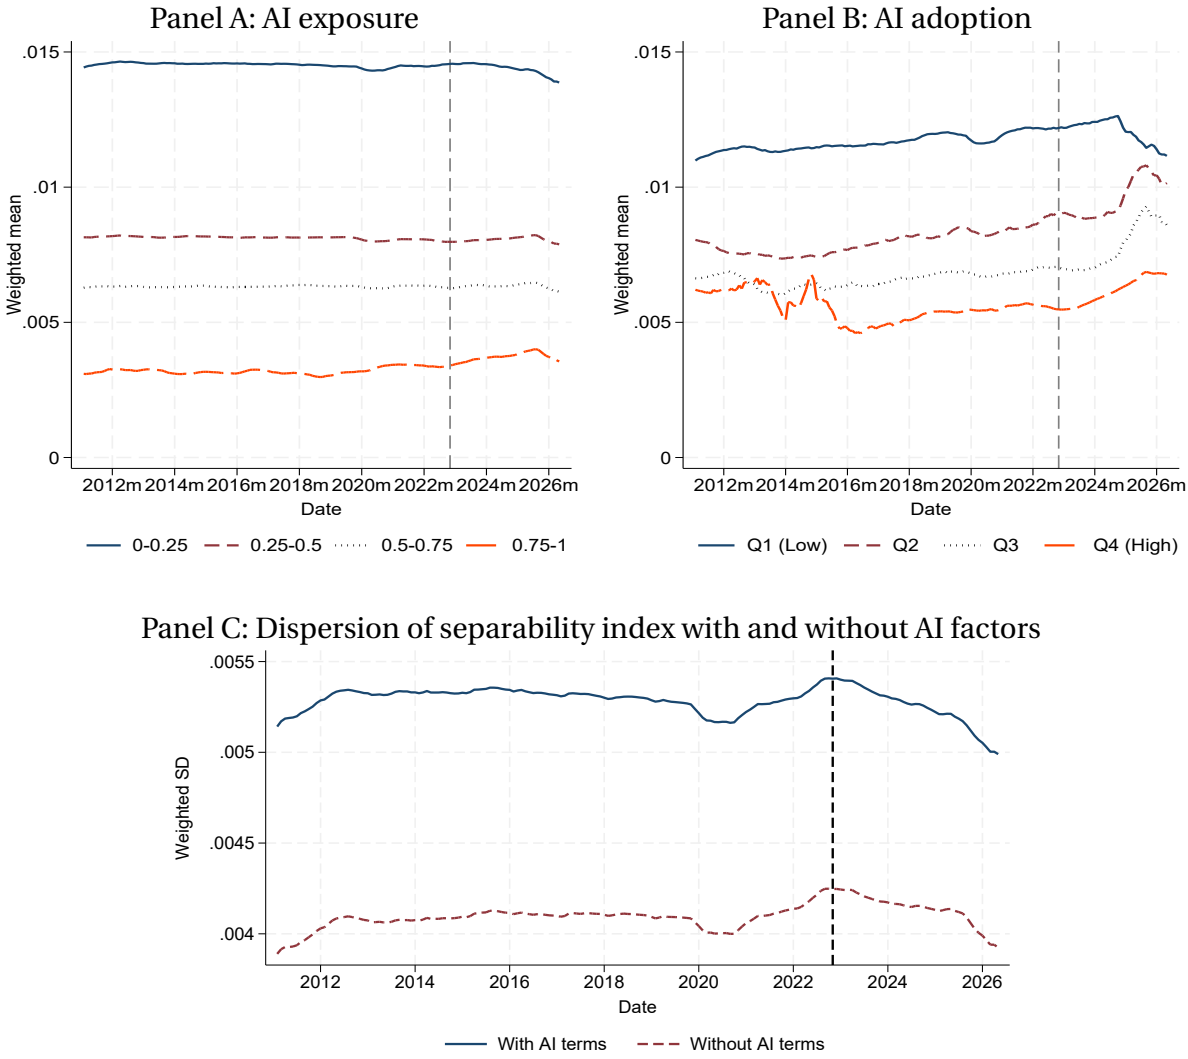
Figure 10: HIRABILITY INDEX BY AI EXPOSURE AND ADOPTION



Note: Panel A displays the mean hirability index by AI exposure score. Panel B displays the mean hirability index by the quartile of AI adoption (thresholds measured in 2025). Panel C plots the dispersion of hirability index with AI exposure and adoption (solid blue line) and that without AI exposure and adoption (dashed red line). Centered 13-month moving averages are plotted to smooth noises in the estimates. **Source:** Authors' calculation.

pressure following the introduction of large language models. By contrast, the dispersion of the AI-purged hirability index ($sd_t^{wo}(UE)$) remains essentially unchanged throughout. This pattern suggests that AI is the main driver of the widening dispersion of hirability and the associated increase in reallocation pressure.

Figure 11: SEPARABILITY INDEX BY THE QUARTILE OF AI EXPOSURE AND AI ADOPTION



Note: Panel A displays the mean separability index by AI exposure score. Panel B displays the mean separability index by the quartile of AI adoption (thresholds measured in 2025). Panel C plots the dispersion of hirability index with AI exposure and adoption (solid blue line) and that without AI exposure and adoption (dashed red line). Centered 13-month moving averages are plotted to smooth noises in the estimates. **Source:** Authors' calculation.

Figure 11 plots the separability index ($sd_t^w(EU)$) by quartiles of AI exposure and AI adoption. In contrast to the hirability index, the separability index by AI exposure shows no notable change following the introduction of LLMs (vertical dashed line).

Figure 11 plots the dispersion of the separation index. Relative to the dispersion of the hirability index, the dispersion of the separation index shows no noticeable change after 2024:Q4 and remains broadly similar to that of the separation index constructed without AI-related factors.

These results suggest that AI-driven reallocation pressures operate primarily through the hiring margin rather than through separations. In addition, the recovery in hirability of the top exposure/adoption group has lowered AI-driven dispersion, suggesting that improved demand for these workers can alleviate reallocation pressure and reduce structural unemployment.

6 Implications for the Natural Rate of Unemployment

The overall empirical evidence points to LLM-driven buildup of reallocation pressure and upward pressure on the structural unemployment rate—and, by extension, the natural rate of unemployment—in the post-LLM period.[32] To assess the effects of LLMs on the natural rate of unemployment, we employ two approaches: one based on a theoretical model of reallocation, and the other based on a statistical model combined with two-state model of the unemployment rate (Shimer [2012]).

6.1 Implications for Structural Unemployment

Jackman and Roper [1987] characterize structural unemployment as arising from mismatch between job seekers and vacancies. Building on this, Şahin et al. [2014] demonstrates that increased mismatch is measured with the dispersion of market tightness across sectors, reflecting heightened frictions in reallocating workers to sectors with better reemployment prospects. We extend Jackman and Roper's model to link structural unemployment to the dispersion of individual job-finding prospects—the hirability— and to identify the extent to which AI factors contributed to changes in structural unemployment.

According to Jackman and Roper [1987], structural unemployment (SU) is defined as

$$SU \equiv \frac{1}{2}\sum_i \left| U_i - \frac{U}{V} V_i \right|, \tag{6.1}$$

where U_i and V_i denote unemployment and vacancies in sector i, and $U = \sum_i U_i$, $V = \sum_i V_i$. The

[32] Note that the structural unemployment rate is a component of the natural rate of unemployment, which includes both frictional and structural unemployment (Daly et al., 2012).

time subscript t is suppressed for notational simplicity. Define shares

$$u_i = \frac{U_i}{U}, \qquad v_i = \frac{V_i}{V}.$$

We then have

$$SU = \frac{U}{2} \sum_i |u_i - v_i|, \tag{6.2}$$

which suggests that mismatch between the distribution of job seekers and job vacancies across sectors determines structural unemployment.

We rewrite equation (6.2) as a function of the dispersion in job-finding probabilities. First, define sectoral and aggregate tightness as follows:

$$\theta_i = \frac{V_i}{U_i}, \qquad \theta = \frac{V}{U}.$$

With $V_i = \theta_i U_i$, equation (6.2) is rewritten as:

$$SU = \frac{U}{2\theta} \sum_i u_i |\theta_i - \theta|. \tag{6.3}$$

Assuming a Cobb–Douglas constant-returns matching function

$$M = AU^{1-\alpha} V^{\alpha}, \qquad 0 < \alpha < 1,$$

we obtain

$$f_i = A\theta_i^{\alpha}, \qquad f = A\theta^{\alpha}.$$

Taking a first-order approximation around θ, we obtain

$$\theta_i - \theta \approx \frac{\theta}{\alpha f}(f_i - f). \tag{6.4}$$

Substituting equation (6.4) into equation (6.3), we have

$$SU \approx \frac{U}{2\alpha f} \sum_i u_i |f_i - f|. \tag{6.5}$$

Now, define the unemployment-weighted variance

$$Var_i(f_i) = \sum_i u_i (f_i - f)^2.$$

Using the approximation between mean absolute deviation and standard deviation,

$$\sum_i u_i |f_i - f| \approx \sqrt{\frac{2}{\pi}} \sqrt{Var_i(f_i)},$$

we arrive at

$$SU \approx \frac{U}{\alpha f \sqrt{2\pi}} \sqrt{Var_u(f_i)}. \tag{6.6}$$

Equation (6.6) provides a micro-founded link between the dispersion of individual job-finding prospects and structural unemployment. Note that our goal is to assess the extent to which LLMs have changed the structural unemployment rate. To do so, we estimate the difference between $\sqrt{\text{Var}_i(f_i)}$ and $\sqrt{\text{Var}_i(f_i^{wo})}$, where the latter is the standard deviation with AI factors purged. The standard deviations of individual-level hirability with and without AI factors — displayed in Figure 10 — proxy for $\sqrt{\widehat{\text{Var}_i(f_i)}}$ and $\sqrt{\widehat{\text{Var}_i(f_i^{wo})}}$, respectively.

According to Figure 10, the AI factors account for +0.005 in the dispersion of hirability, with an additional widening of +0.008 after 2022:M11 reflecting LLMs' effects, for a total AI effect of +0.013. Using these estimates, we compute the implied change in the structural unemployment rate.

Table 5 reports the implied change in the structural unemployment rate for three values of the matching elasticity α and three levels of dispersion. For the calibration, we set the unemployment rate to 4% and the aggregate job-finding probability to 0.3. We consider three cases of matching elasticity to vacancies $\alpha = 0.3$, 0.5 and 0.7.

Before the advent of LLMs, AI-driven dispersion ($\Delta\sigma = 0.005$) raised the structural unemployment rate by 0.04–0.09 percentage points depending on α. The LLM contribution alone ($\Delta\sigma = 0.008$) adds a further 0.06–0.14 percentage points, and the total AI effect ($\Delta\sigma = 0.013$) implies a boost of 0.10–0.23 percentage points. Our preferred baseline is $\alpha = 0.3$ as shown in recent empirical studies (e.g., Figura and Waller, 2024).[33] We therefore evaluate that LLMs likely

[33]The estimated structural unemployment rates are naturally larger under lower matching elasticity, as weaker substitutability between worker types amplifies the unemployment consequences of mismatch.

Table 5: IMPLIED CHANGE IN STRUCTURAL UNEMPLOYMENT RATE (PERCENTAGE POINT)

	Matching Elasticity α		
	$\alpha = 0.3$	$\alpha = 0.5$	$\alpha = 0.7$
Before LLM ($\Delta\sigma = 0.005$)	0.09	0.05	0.04
LLM contribution ($\Delta\sigma = 0.008$)	0.14	0.09	0.06
Total AI effect ($\Delta\sigma = 0.013$)	0.23	0.14	0.10

Notes: Entries are percentage point changes in the structural unemployment rate, computed as $\Delta SU = u \cdot \Delta\sigma / (\alpha f \sqrt{2\pi})$ with $u = 4\%$, $f = 0.3$, and $\Delta\sigma$ denoting the change in the dispersion of job-finding probabilities. **Source**: Authors' calculation.

boosted structural unemployment rate by 0.14 percentage point.

6.2 Statistical Model

We statistically estimate the natural rate of unemployment for four groups defined jointly by AI exposure and adoption. To do so, we recover the trend components of each group's unemployment inflow and outflow rates using a Bayesian trend-cycle decomposition, then feed these trends into Shimer's two-state model of unemployment rate to recover each group's natural rate. We then aggregate these estimates using labor-force shares to obtain the aggregate natural rate.[34] This framework allows us to account for both the trend and cyclical components of worker flows in the pre-LLMs period and to control for cyclical fluctuations during the LLM era — both of which are critical when assessing the effects of LLMs on the natural rate.

For each group $g \in \{LL, LH, HL, HH\}$ and each labor market flow $j \in \{\text{EU}, \text{UE}\}$, the observed monthly flow rate $y^j_{g,t}$ is decomposed into a stochastic trend $\tau^j_{g,t}$, a stochastic cycle $c^j_{g,t}$, and a measurement error:

$$y^j_{g,t} = \tau^j_{g,t} + c^j_{g,t} + \varepsilon^j_{g,t}, \qquad \varepsilon^j_{g,t} \sim \mathcal{N}\left(0, \sigma^{2,j}_{\varepsilon,g}\right). \tag{6.7}$$

The sample period is 2013:M1-2026:M5. Observations during the COVID-19 swing (2020:M3–2021:M12) are treated as missing and excluded from likelihood evaluation.[35]

[34] This approach has been adopted by Tasci [2012], Crump et al. [2019], Ahn [2023].

[35] This treatment prevents the large, transitory COVID shock from persistently distorting inference on the trend and cyclical components, following Ng [2021]. We consider narrower windows for the COVID-19 swing, but the main conclusion does not change meaningfully.

6.2.1 Trend: Smooth Trend

We model the trend components of flow rates as integrated random walks, a specification that produces smooth trends and flexibly captures changing momentum.

EU rate:

$$\tau^{\text{EU}}_{g,t} = \tau^{\text{EU}}_{g,t-1} + \beta^{\text{EU}}_{g,t-1}, \tag{6.8}$$

$$\beta^{\text{EU}}_{g,t} = \beta^{\text{EU}}_{g,t-1} + \zeta^{\text{EU}}_{g,t}, \qquad \zeta^{\text{EU}}_{g,t} \sim \mathcal{N}\left(0, \sigma^{2,\text{EU}}_{\zeta,g}\right). \tag{6.9}$$

UE rate:

$$\tau^{\text{UE}}_{g,t} = \tau^{\text{UE}}_{g,t-1} + \beta^{\text{UE}}_{g,t-1}, \tag{6.10}$$

$$\beta^{\text{UE}}_{g,t} = \beta^{\text{UE}}_{g,t-1} + \zeta^{\text{UE}}_{g,t}, \qquad \zeta^{\text{UE}}_{g,t} \sim \mathcal{N}\left(0, \sigma^{2,\text{UE}}_{\zeta,g}\right). \tag{6.11}$$

6.2.2 Cycle

For the cyclical component, we use Harvey [1985]'s trigonometric stochastic cycle.[36]

Defining the auxiliary variable $c^{j*}_{g,t}$, the pair $\left(c^{j}_{g,t}, c^{j*}_{g,t}\right)$ evolves as

$$\begin{pmatrix} c^{j}_{g,t} \\ c^{j*}_{g,t} \end{pmatrix} = \rho \begin{pmatrix} \cos\lambda & \sin\lambda \\ -\sin\lambda & \cos\lambda \end{pmatrix} \begin{pmatrix} c^{j}_{g,t-1} \\ c^{j*}_{g,t-1} \end{pmatrix} + \begin{pmatrix} \kappa^{j}_{g,t} \\ \kappa^{j*}_{g,t} \end{pmatrix}, \tag{6.12}$$

where $\lambda = 2\pi/p$ is the fundamental frequency and $\kappa^{j}_{g,t}, \kappa^{j*}_{g,t} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^{2,j}_{\kappa,g})$. We set $p = 36$ months and $\rho = 0.90$ is the damping factor. For the ease of computation and transparency given a relatively short sample period, we calibrate p and ρ. The calibrated value implies that a length of cyclical component is 3 years during the sample period, conditional on the trend.

[36] We could alternatively model the cycles using stationary AR(2) processes, but doing so would require informative priors to separate the cyclical component from the trend given the relatively short sample period. We view the trigonometric approach as more restrictive, but also more transparent regarding the prior assumptions imposed on the dynamics of the cyclical component.

6.2.3 State-Space Representation

Collecting the latent states as $\alpha^j_{g,t} = \left(\tau^j_{g,t},\, \beta^j_{g,t},\, c^j_{g,t},\, c^{j*}_{g,t}\right)'$, the system takes the linear Gaussian state-space form

$$y^j_{g,t} = Z\,\alpha^j_{g,t} + \varepsilon^j_{g,t}, \tag{6.13}$$

$$\alpha^j_{g,t} = T\,\alpha^j_{g,t-1} + u^j_{g,t}, \tag{6.14}$$

with system matrices

$$Z = \begin{pmatrix} 1 & 0 & 1 & 0 \end{pmatrix}, \qquad T = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \rho\cos\lambda & \rho\sin\lambda \\ 0 & 0 & -\rho\sin\lambda & \rho\cos\lambda \end{pmatrix},$$

and state-disturbance covariance

$$\mathrm{Var}(u^j_{g,t}) = Q^j_g = \mathrm{diag}\left(0,\, \sigma^{2,j}_{\zeta,g},\, \sigma^{2,j}_{\kappa,g},\, \sigma^{2,j}_{\kappa,g}\right),$$

where the zero in the first position reflects the smooth-trend restriction ($\sigma^{2,j}_{\eta,g} = 0$) for both flows. The Harvey cycle block is initialized at its stationary covariance, $\sigma^{2,j}_{\kappa,g}/(1-\rho^2)\cdot I_2$; the trend and slope states are initialized diffusely. The model is estimated with a Bayesian method. Section E details the estimation including priors and posterior simulations.

6.3 Natural Rate of Unemployment

6.3.1 Group-Level Estimate

We feed the trend UE and EU rates for group g into Shimer's two-state model of unemployment rate to recover the natural rate of unemployment for that group as follows:

$$u^*_{g,t} = \frac{100\cdot\tau^{\mathrm{EU}}_{g,t}}{\tau^{\mathrm{EU}}_{g,t} + \tau^{\mathrm{UE}}_{g,t}}, \tag{6.15}$$

where $\tau^{\mathrm{EU}}_{g,t}$ and $\tau^{\mathrm{UE}}_{g,t}$ are the trend components of the EU and UE flow rates for group g. Equation (6.15) is evaluated draw-by-draw so that posterior uncertainty in both trend components propagates into uncertainty about $u^*_{g,t}$.[37]

6.3.2 Aggregate Estimate

Groups are aggregated using time-varying labor-force shares $\omega_{g,t}$ (normalized to sum to one each period):

$$u^*_{\mathrm{agg},t} = \sum_{g=1}^{4} \omega_{g,t}\, u^*_{g,t}. \tag{6.16}$$

Aggregation is performed draw-by-draw, so the posterior distribution of $u^*_{\mathrm{agg},t}$ reflects uncertainty in every group's trend and time variation in labor-force composition.

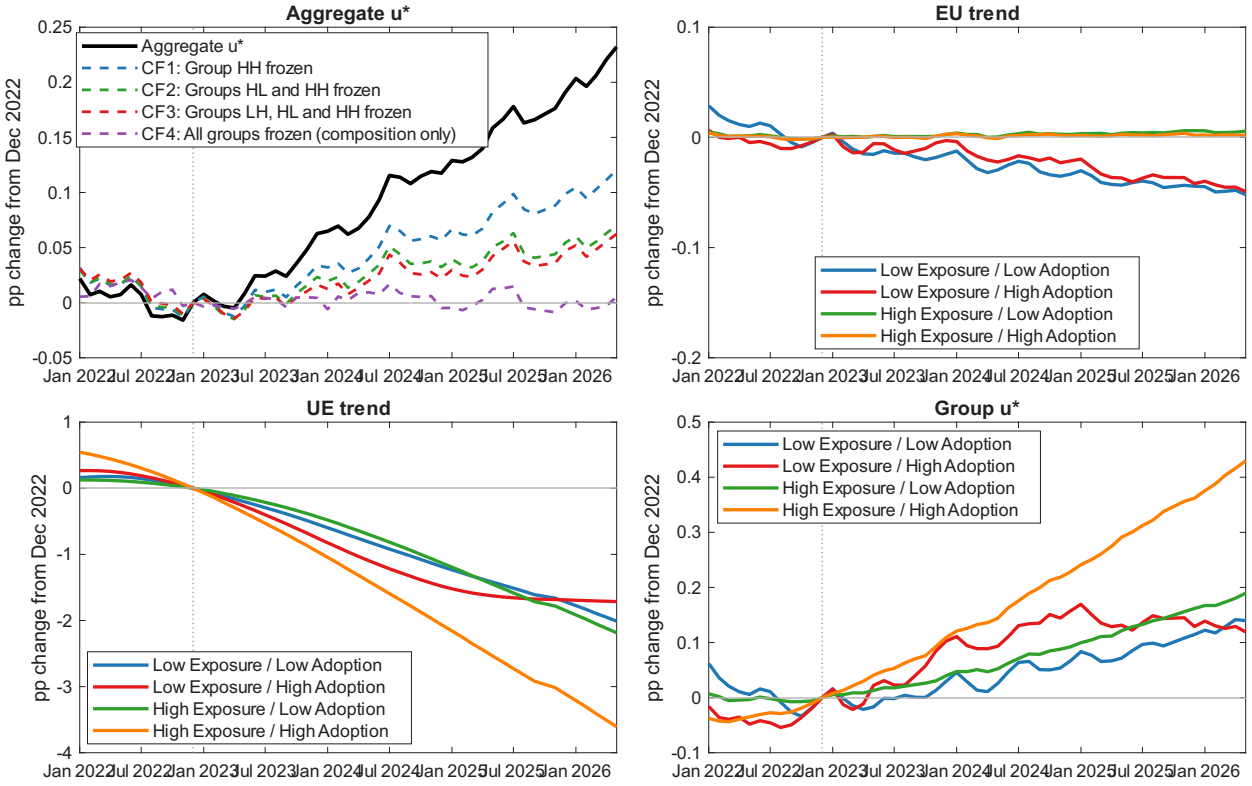
6.3.3 Counterfactuals

To assess the contribution of AI exposure/adoption to the rise in the natural rate, we construct a series of counterfactual aggregate natural rates. Each counterfactual freezes the natural rate of one or more groups at the end of 2022, while allowing the remaining groups to evolve freely. The gap between the actual aggregate natural rate and a given counterfactual thus measures the contribution of the frozen group(s) to the overall change.

Specifically, the first counterfactual freezes only the high-exposure/high-adoption group (Group 4, HH), isolating the role of workers most intensively exposed to and engaged with AI. The second counterfactual additionally freezes the high-exposure/low-adoption group (Group 3, HL), capturing the combined effect of all high-exposure workers. The third counterfactual further freezes the low-exposure/high-adoption group (Group 2, LH), so that only the low-exposure/low-adoption group (Group 1, LL) is permitted to evolve. A fourth counterfactual freezes the natural

[37] The steady-state model should approximate the unemployment rate for each group. However, measurement error in flows tends to understate the unemployment rate, with the degree of understatement varying across groups. Excluding the COVID period (2020:M3–2021:M12), the gap between the model-implied and observed unemployment rate is stable over our sample (2022:M1–2026:M5). We therefore adjust each group's natural rate by a constant correction factor. This factor does not affect labor-force shares across groups and hence does not alter the contribution of each group's AI exposure and adoption to changes in the natural rate. An alternative approach would be to adopt "raking" adjusting the flows to match each group's stock unemployment rate. In the interest of transparency, we opt for the simpler constant adjustment rather than modifying further the underlying flow and stock data.

Figure 12: TREND UNEMPLOYMENT FLOWS AND U-STAR SINCE LLMS



Notes: The upper-left panel displays the aggregate natural rate and a counterfactual series; the upper-right and lower-left panels show the EU and UE trend rates for the four groups, respectively; and the lower-right panel displays the estimated natural rate of each group. **Source**: Authors' calculation.

rates of all four groups simultaneously at their December 2022 levels, so that the only source of variation in the aggregate natural rate is the time-varying labor-force shares of the groups. This counterfactual therefore isolates the contribution of compositional shifts in the labor force to overall changes in the natural rate.

6.4 Estimation results

Figure 12 zooms in on the estimates for the period following the introduction of LLMs.[38] For ease of comparison, all values are normalized to zero in December 2022.

The upper-left panel displays the aggregate natural rate (solid line) alongside the four counterfactuals (dashed line). The purple dashed line denotes the counterfactual natural rate when all groups are held at their December 2022 levels, reflecting compositional shifts alone. The red

[38] See Section E.3 for the background estimation results.

dashed line additionally allows the LL group's natural rate to vary. The green dashed line further allows the LH group to vary. The blue dashed line additionally frees the HL group. Finally, the gap between the solid black line (actual aggregate natural rate) and the blue dashed line captures the contribution of the HH group.

The HH group alone accounts for a little more than 0.1 percentage point (12 bp) of the rise in the natural rate, while the HL group contributes approximately 5 basis points. Taken together, high-exposure workers raise the natural rate by about 0.15 percentage point (17 bp). This upward pressure is primarily driven by a faster decline in the trend job-finding rate among the HH group (the orange line, bottom-left panel) relative to the much slower downtrend in the other groups. Meanwhile, the LL group contributes roughly 5 basis points, likely reflecting modest upward pressure from factors such as increased immigration, as lower-skilled immigrants tend to face less efficient job searches during the labor market assimilation process. The LH group's contribution is negligible largely due to their small share in the labor force.

In sum, this result also reinforces the importance of hiring margin in the upside pressure on the natural rate driven by LLMs. The contribution of higher AI-exposure group to the natural rate aligns with the increase in structural unemployment implied by the calibration based on Jackman and Roper [1987]. Nonetheless, the estimates are subject to considerable uncertainty as shown in the sizable posterior intervals in each group's estimate (Figure A.9). We therefore interpret our results as indicating potential upside risks to the natural rate of unemployment associated with LLMs.

7 Conclusion

In this paper, we provide a comprehensive assessment of how the rapid diffusion of AI—particularly large language models—has affected the aggregate labor market. By integrating traditional labor market data with measures of AI exposure and adoption and tracing their effects through labor market stocks and flows, we uncover evidence of rising reallocation pressure and structural unemployment. Since the introduction of LLMs, workers with high AI exposure and adoption have experienced a disproportionate decline in job-finding rates and job-to-job mobility, alongside weaker market tightness relative to the pre-COVID level, indicating that adjustment has

occurred primarily along the hiring margin. The rise in within-job activity switching among highly exposed workers suggests that firms are reorganizing tasks rather than replacing workers, consistent with a reconfiguration of production processes and reduced demand for workers in high-AI-intensity jobs.

These patterns point to a buildup of reallocation pressure in the post-LLM period. The sharp rise in the dispersion of job-finding propensities—largely driven by AI-related factors—suggests increased disparities in reemployment prospects, which our flow-based framework translates into modest upward pressure on the natural rate of unemployment since late 2022. Recent signs of recovery in job-finding and job-switching rates for HH workers, along with rising market tightness among them, could signal some improvement in labor demand, possibly reflecting reduced uncertainty about both the macroeconomic environment and the capabilities of LLMs. To the extent that the low-hiring, low-firing environment is itself an endogenous response to uncertainty, policies that reduce uncertainty may help alleviate reallocation pressure and facilitate more efficient labor market adjustment.

References

Jaison R. Abel, Richard Deitz, and Natalia Emanuel. Survey evidence on firms' use of artificial intelligence, 2025. Federal Reserve Bank of New York, unpublished survey evidence.

Katharine G. Abraham and Lawrence F. Katz. Cyclical unemployment: Sectoral shifts or aggregate disturbances? *Journal of Political Economy*, 94(3):507–522, 1986. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/1833046.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019. doi: 10.1257/jep.33.2.3. URL https://doi.org/10.1257/jep.33.2.3.

Daron Acemoglu, David Autor, Jonathon Hazell, and Pascual Restrepo. Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):293–340, 2022.

Hie Joo Ahn. Duration structure of unemployment hazards and the trend unemployment rate. *Journal of Economic Dynamics and Control*, 151:104670, 2023. ISSN 0165-1889. doi: 10.1016/j.jedc.2023.104670.

Hie Joo Ahn and James D. Hamilton. Measuring labor-force participation and the incidence and duration of unemployment. *Review of Economic Dynamics*, 44:1–32, 2022. ISSN 1094-2025. doi: https://doi.org/10.1016/j.red.2021.04.005. URL https://www.sciencedirect.com/science/article/pii/S109420252100034X.

Lukas Althoff and Hugo Reichardt. Task-specific technical change and comparative advantage. Working Paper, June 2026. URL https://hugoreichardt.com/pdf/tstc_compadvantage.pdf.

David Autor and Neil Thompson. Expertise. *Journal of the European Economic Association*, 23(4):1203–1271, 2025. doi: 10.1093/jeea/jvaf023. URL https://doi.org/10.1093/jeea/jvaf023.

Salome Baslandze, Erik Brynjolfsson, Daniel Rock, and Chad Syverson. Artificial intelligence, productivity, and the workforce: Evidence from corporate executives. Technical Report 34984, National Bureau of Economic Research, 2026. URL https://www.nber.org/papers/w34984.

Kathryn Bonney, Cory L Breaux, Emin Dinlersoz, Lucia S. Foster, John C Haltiwanger, and Aditya A Pande. The microstructure of ai diffusion: Evidence from firms, business functions, and worker tasks. NBER Working Paper No. 35141, 2026.

Erik Brynjolfsson, Tom Mitchell, and Daniel Rock. What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108:43–47, 2018.

Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen. Canaries in the coal mine?: Six facts about the recent employment effects of artificial intelligence. Working Paper, 2025.

C. K. Carter and R. Kohn. On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553, 1994. doi: 10.1093/biomet/81.3.541.

Leland D. Crane and Paul E. Soto. Ai and coder employment: Compiling the evidence. Finance and Economics Discussion Series (FEDS): https://doi.org/10.17016/FEDS.2026.018, 2026.

Richard K. Crump, Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin. A Unified Approach to Measuring u. NBER Working Papers 25930, National Bureau of Economic Research, Inc, June 2019. URL https://ideas.repec.org/p/nbr/nberwo/25930.html.

Mary C. Daly, Bart Hobijn, AysegÃ¼l Sahin, and Robert G. Valletta. A search and matching approach to labor markets: Did the natural rate of unemployment rise? *Journal of Economic Perspectives*, 26(3):3–26, Summer 2012. doi: 10.1257/jep.26.3.3. URL https://ideas.repec.org/a/aea/jecper/v26y2012i3p3-26.html.

J. Durbin and S. J. Koopman. A simple and efficient simulation smoother for state space time series analysis. *Biometrika*, 89(3):603–615, 2002. doi: 10.1093/biomet/89.3.603.

Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. Gpts and gpts: An early look at the labor market impact potential of large language models. arXiv:2303.10130, 2023.

Edward W. Felten, Manav Raj, and Robert Seamans. Occupational heterogeneity in exposure to generative ai. Working Paper, 2023.

Andrew Figura and Chris Waller. What does the beveridge curve tell us about the likelihood of soft landings? *Journal of Economic Dynamics and Control*, 169:104957, 2024. ISSN 0165-1889. doi: https://doi.org/10.1016/j.jedc.2024.104957. URL https://www.sciencedirect.com/science/article/pii/S0165188924001490. Heterogeneity and Macroeconomics of Labor Markets.

Lukas B. Freund and Lukas F. Mann. Job transformation, specialization, and the labor market effects of AI. Institute Working Paper 114, Federal Reserve Bank of Minneapolis and CESifo, 2026. URL https://doi.org/10.21034/iwp.114.

Menaka Hampole, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller. Artificial intelligence and the labor market. *NBER Working Paper No. 33509*, 2025.

Kunal Handa, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, Kevin K. Troy, Dario Amodei, Jared Kaplan, Jack Clark, and Deep Ganguli. Which economic tasks are performed with ai? evidence from millions of claude conversations, 2025. URL https://arxiv.org/abs/2503.04761.

A. C. Harvey. Trends and cycles in macroeconomic time series. *Journal of Business & Economic Statistics*, 3(3):216–227, 1985. doi: 10.1080/07350015.1985.10509453.

Seyed M. Hosseini and Guy Lichtinger. Generative ai as seniority-biased technological change: Evidence from u.s. résumé and job posting data. Working Paper, 2025.

Seyed Mahdi Hosseini Maasoum and Guy Lichtinger. Generative ai and occupational entry barriers: The labor-supply channel of technological change. Technical report, SSRN, January 2026. Available at SSRN: https://ssrn.com/abstract=6059674.

Richard Jackman and Stephen Roper. Structural unemployment. *Oxford Bulletin of Economics and Statistics*, 49(1):9–36, 1987.

Peter J. Lambert and Schindler Yannick. The broken ladder: Ai, remote work, and early-career hiring. http://dx.doi.org/10.2139/ssrn.6787638, 2026.

David M. Lilien. Sectoral shifts and cyclical unemployment. *Journal of Political Economy*, 90(4): 777–793, 1982. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/1831352.

Sam J. Manning and TomÃ¡s Aguirre. How adaptable are american workers to ai-induced job displacement? NBER Working Papers 34705, National Bureau of Economic Research, Inc, Jan 2026. URL https://ideas.repec.org/p/nbr/nberwo/34705.html.

George R. Neumann and Robert H. Topel. Employment Risk, Diversification, and Unemployment. *The Quarterly Journal of Economics*, 106(4):1341–1365, 1991. URL https://ideas.repec.org/a/oup/qjecon/v106y1991i4p1341-1365..html.

Serena Ng. Modeling macroeconomic variations after covid-19. Working Paper 29060, National Bureau of Economic Research, 2021. URL https://nber.org.

Ellen R. Rissman. Wage growth and sectoral shifts : Phillips curve redux. *Journal of Monetary Economics*, 31(3):395–416, June 1993. doi: None. URL https://ideas.repec.org/a/eee/moneco/v31y1993i3p395-416.html.

Robert Shimer. Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 15(2):127–148, 2012. ISSN 1094-2025. doi: https://doi.org/10.1016/j.red.2012.02.001. URL https://www.sciencedirect.com/science/article/pii/S1094202512000063.

Murat Tasci. The ins and outs of unemployment in the long run: unemployment flows and the natural rate. Working Papers 1224, Federal Reserve Bank of Cleveland, 2012. URL https://ideas.repec.org/p/fip/fedcwp/1224.html.

The Budget Lab at Yale. Artificial intelligence and labor market outcomes. https://budgetlab.yale.edu, 2025a. Report finding no relationship between AI exposure and recent employment or unemployment changes.

The Budget Lab at Yale. Evaluating the impact of ai on the labor market: Occupational composition and employment trends. https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market, 2025b. Finds that the occupational mix of the U.S. labor market is changing somewhat faster than historically.

Kiran Tomlinson, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. Working with ai: Measuring the occupational implications of generative ai. Working Paper, 2025.

Lee C. Tucker. You're (not) hired: Artificial intelligence and early career hiring in the quarterly workforce indicators. Census Bureau Working Paper CES-26-27, 2026.

Michael Webb. The impact of artificial intelligence on the labor market. Working Paper, 2020.

Michelle Yin, Hoa Vu, and Claudia Persico. How (un)stable are llm occupational exposure scores? evidence from multi-model replication. NBER Working Paper No. 35110, 2026.

Ayşegül Şahin, Joseph Song, Giorgio Topa, and Giovanni L. Violante. Mismatch unemployment. *American Economic Review*, 104(11):3529–3564, November 2014. doi: 10.1257/aer.104.11.3529.

A Measurement of AI exposure and adoption

This section discusses the most representative occupations by AI exposure (Table A.1), by AI adoption (Table A.2), and joint categorization of AI exposure and adoption (Table A.3).

The occupation–industry pairs by AI adoption show a clear gradient from manual and service work at the bottom to knowledge-intensive professions at the top. Quartile 1, the least exposed, is concentrated in food service, construction, and retail—waiters, cooks, cashiers, and construction laborers working predominantly in eating and drinking establishments and the construction sector. Quartile 2 shifts toward the public sector and healthcare: elementary and secondary school teachers, registered nurses, and teacher assistants account for most of the largest cells, with these workers employed heavily in schools and hospitals. Quartile 3 occupations span a more heterogeneous set of industries—real estate brokerage, accounting, agriculture, and some assembly work. Quartile 4 is dominated by high-skill professional and knowledge workers: software developers, lawyers, financial managers, management analysts, and postsecondary teachers, predominantly in computer services, legal services, finance, and universities.

B Additional Data by AI exposure and adoption

Figures A.1 and Figures A.2 display the fully-imputed data by AI exposure and the joint classification, respectively.

Figure A.3 displays the unemployment rate, EU/UE flow rates and labor force share by the quartile of AI adoption. Overall, the patterns are similar to those by AI exposure. One difference is that the separation rate of the second adoption quartile rose faster than other groups, pushing up their unemployment rate alongside the rise in the top quartile. This pickup is driven primarily by workers with AI exposure scores above 50, and especially those above 75.[39] However, this group

[39]Among workers in the second adoption quartile, separation rates of those with AI exposure below 50 were unchanged between 2022 and 2025, while those with exposure scores of 50–75 more than doubled (from 0.4% to 1.1%) and those above 75 tripled (from 0.4% to 1.3%). These groups, constituting about 10% of the second quartile group and 2% of the labor force, include computer programmers, database administrators, data entry keyers, and

Table A.1: TOP 20 OCCUPATIONS BY AI GROUP (EXPOSURE)

Group	Occupation	AI Exposure
Group 1		
1	Cashiers	0.236
1	Janitors and building cleaners	0.000
1	Laborers and freight, stock, and material movers, hand	0.042
1	Nursing, psychiatric, and home health aides	0.076
1	Cooks	0.091
1	Waiters and waitresses	0.184
1	Construction laborers	0.050
1	Stockers and order fillers	0.182
1	Maids and housekeeping cleaners	0.000
1	Personal care aides	0.222
Group 2		
2	Managers, all other	0.482
2	Driver/sales workers and truck drivers	0.326
2	Elementary and middle school teachers	0.327
2	First-line supervisors of retail sales workers	0.486
2	Registered nurses	0.380
2	Retail salespersons	0.342
2	Chief executives	0.472
2	First-line supervisors of office and administrative support workers	0.490
2	Financial managers	0.485
2	Postsecondary teachers	0.486
Group 3		
3	Secretaries and administrative assistants	0.664
3	Customer service representatives	0.568
3	Accountants and auditors	0.560
3	Office clerks, general	0.576
3	Receptionists and information clerks	0.533
3	Sales representatives, wholesale and manufacturing	0.539
3	First-line supervisors of non-retail sales workers	0.537
3	Computer occupations, all other	0.736
3	Computer and information systems managers	0.530
3	Insurance sales agents	0.576
Group 4		
4	Software developers, applications and systems software	0.873
4	Bookkeeping, accounting and auditing clerks	0.802
4	Billing and posting clerks	0.774
4	Computer programmers	0.950
4	Insurance claims and policy processing clerks	0.833
4	Data entry keyers	0.893
4	Writers and authors	0.877
4	Other mathematical science occupations	0.783
4	File Clerks	0.794
4	Web developers	0.811

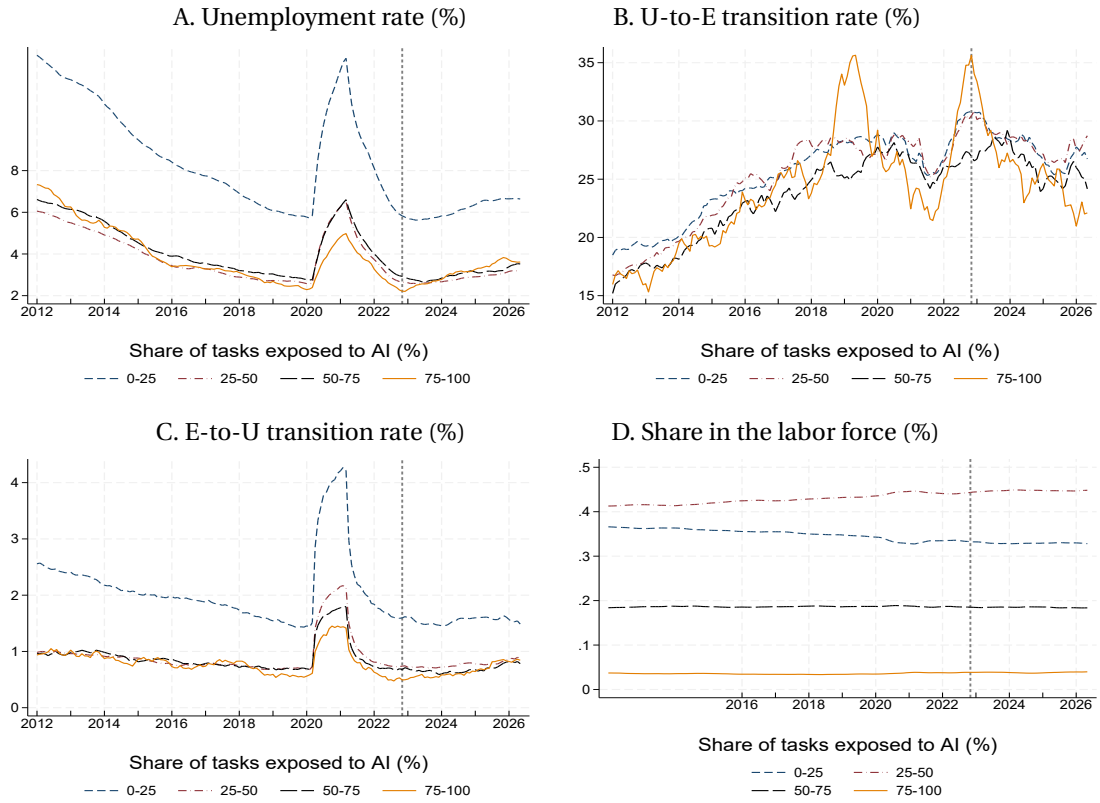
Table A.2: TOP 10 OCCUPATION–INDUSTRY CELLS BY AI ADOPTION

Group	Occupation	Industry	AI Adoption
Group 1			
1	Waiters and waitresses	eating and drinking places	0.002
1	Cooks	eating and drinking places	0.002
1	Driver/sales workers and truck drivers	trucking service	0.005
1	Grounds maintenance workers	landscape and horticultural services	0.003
1	Food service managers	eating and drinking places	0.002
1	Fast food and counter workers	eating and drinking places	0.002
1	Janitors and building cleaners	services to dwellings and other buildings	0.003
1	Couriers and messengers	trucking service	0.005
1	Childcare workers	child day care services	0.003
1	Cashiers	eating and drinking places	0.002
Group 2			
2	Elementary and middle school teachers	elementary and secondary schools	0.008
2	Construction laborers	all construction	0.007
2	Registered nurses	hospitals	0.007
2	Teacher assistants	elementary and secondary schools	0.008
2	Construction managers	all construction	0.007
2	Carpenters	all construction	0.007
2	Secondary school teachers	elementary and secondary schools	0.008
2	Managers, all other	all construction	0.007
2	Electricians	all construction	0.007
2	First-line supervisors of construction trades and extraction workers	all construction	0.007
Group 3			
3	Real estate brokers and sales agents	real estate, including real estate-insurance offices	0.010
3	Property, real estate, and community association managers	real estate, including real estate-insurance offices	0.010
3	Nursing, psychiatric, and home health aides	health services, n.e.c	0.009
3	Accountants and auditors	accounting, auditing, and bookkeeping services	0.014
3	Other agricultural workers	agricultural production, crops	0.009
3	Other teachers and instructors	educational services, n.e.c	0.013
3	Registered nurses	health services, n.e.c	0.009
3	Farmers, ranchers, and other agricultural managers	agricultural production, crops	0.009
3	Registered nurses	hospitals	0.008
3	Personal care aides	health services, n.e.c	0.009
Group 4			
4	Software developers, applications and systems software	computer and data processing services	0.062
4	Postsecondary teachers	colleges and universities	0.027
4	Management analysts	management and public relations services	0.035
4	Insurance sales agents	insurance	0.021
4	Lawyers, and judges, magistrates and other judicial workers	legal services	0.017
4	Managers, all other	computer and data processing services	0.062
4	Financial managers	banking	0.030
4	Computer occupations, all other	computer and data processing services	0.062
4	Managers, all other	management and public relations services	0.035
4	Civil engineers	engineering, architectural, and surveying services	0.025

Table A.3: TOP 10 OCCUPATION–INDUSTRY CELLS BY AI EXPOSURE/ADOPTION GROUP

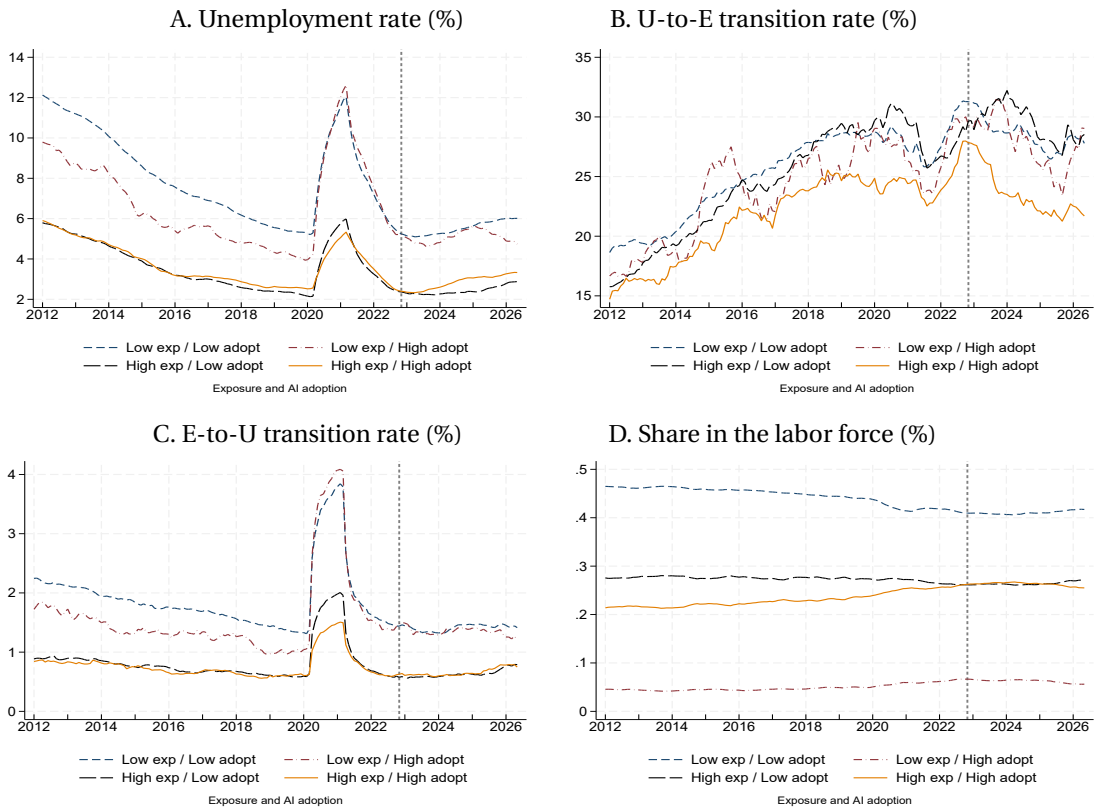
Group	Occupation	Industry	AI Exposure
Low-exposure/Low-adoption			
1	Elementary and middle school teachers	elementary and secondary schools	0.327
1	Construction laborers	all construction	0.050
1	Waiters and waitresses	eating and drinking places	0.184
1	Driver/sales workers and truck drivers	trucking service	0.326
1	Cooks	eating and drinking places	0.091
1	Grounds maintenance workers	landscape and horticultural services	0.047
1	Carpenters	all construction	0.144
1	Nursing, psychiatric, and home health aides	health services, n.e.c	0.076
1	Fast food and counter workers	eating and drinking places	0.097
1	Couriers and messengers	trucking service	0.310
Low-exposure/High-adoption			
2	Other assemblers and fabricators	motor vehicles and motor vehicle equipment	0.115
2	Personal care aides	administration of human resources programs	0.222
2	Counselors	colleges and universities	0.242
2	Recreation and fitness workers	miscellaneous entertainment and recreation services	0.126
2	Electrical power-line installers and repairers	electric light and power	0.000
2	Social workers	administration of human resources programs	0.286
2	Printing press operators	printing, publishing, and allied industries, except newspapers	0.287
2	Athletes, coaches, umpires, and related workers	miscellaneous entertainment and recreation services	0.204
2	Janitors and building cleaners	colleges and universities	0.000
2	Other entertainment attendants and related workers	miscellaneous entertainment and recreation services	0.180
High-exposure/Low-adoption			
3	Registered nurses	hospitals	0.380
3	Construction managers	all construction	0.444
3	Food service managers	eating and drinking places	0.417
3	Real estate brokers and sales agents	real estate, including real estate-insurance offices	0.469
3	Teacher assistants	elementary and secondary schools	0.388
3	Managers, all other	all construction	0.482
3	First-line supervisors of construction trades and extraction workers	all construction	0.417
3	Property, real estate, and community association managers	real estate, including real estate-insurance offices	0.443
3	Registered nurses	health services, n.e.c	0.380
3	First-line supervisors of retail sales workers	grocery stores	0.486
High-exposure/High-adoption			
4	Software developers, applications and systems software	computer and data processing services	0.873
4	Postsecondary teachers	colleges and universities	0.486
4	Lawyers, and judges, magistrates and other judicial workers	legal services	0.460
4	Management analysts	management and public relations services	0.405
4	Insurance sales agents	insurance	0.576
4	Managers, all other	computer and data processing services	0.482
4	Accountants and auditors	accounting, auditing, and bookkeeping services	0.560
4	Financial managers	banking	0.485
4	Paralegals and legal assistants	legal services	0.525
4	Computer occupations, all other	computer and data processing services	0.736

Figure A.1: UNEMPLOYMENT RATE, UE AND EU RATES AND LABOR FORCE SHARES BY AI EXPOSURE (FULL IMPUTATION)



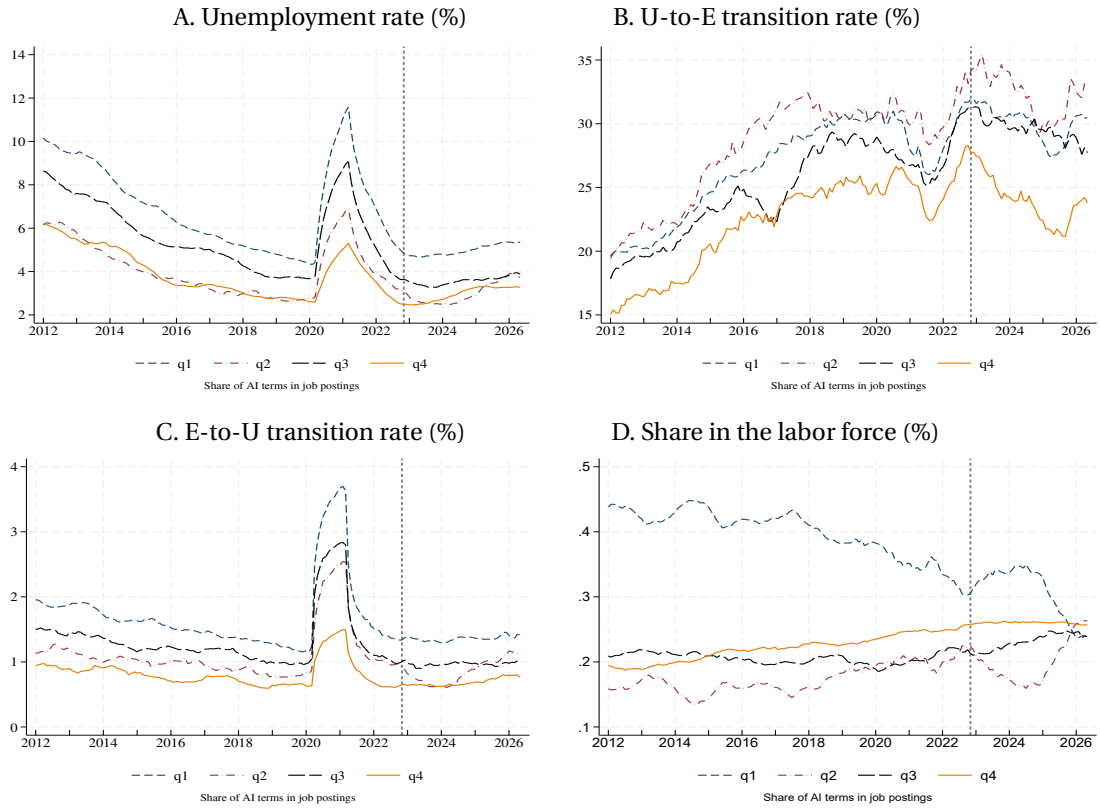
Note: This figure displays the unemployment rate (Panel A), UE rate (Panel B), EU rate (Panel C), and labor force shares (Panel D) for four AI exposure groups with full imputation (12-month moving average). Blue dashed line: 0-0.25; red dashed-dot line: 0.25-0.5; black long-dashed line: 0.5-0.75; solid yellow line: 0.75-1. **Source:** Authors' calculation.

Figure A.2: UNEMPLOYMENT RATE, UE AND EU RATES AND LABOR FORCE SHARES BY AI EXPOSURE/ADOPTION (FULL IMPUTATION)



Note: This figure displays the unemployment rate (Panel A), UE rate (Panel B), EU rate (Panel C), and labor force shares (Panel D) for four AI exposure/adoption groups (12-month moving average). "Exp" stands for exposure, and "adopt" stands for adoption. Blue dashed line: Low-exp/Low-adopt; red dashed-dot line: Low-exp/High-adopt; black long-dashed line: High-exp/Low-adopt; solid yellow line: High-exp/High-adopt. **Source:** Authors' calculation.

Figure A.3: UNEMPLOYMENT RATE, UE AND EU RATES AND LABOR FORCE SHARES BY AI ADOPTION



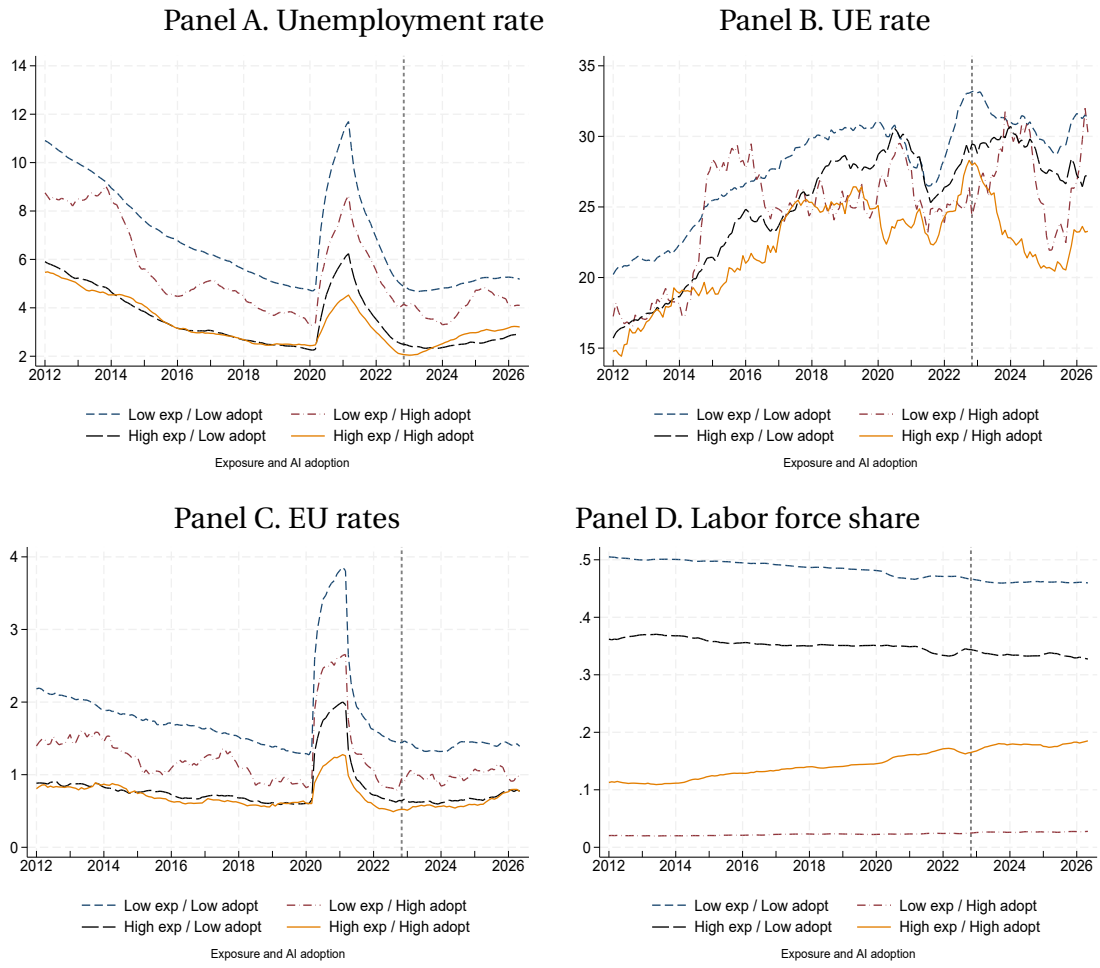
Note: This figure displays the unemployment rate (Panel A), UE rate (Panel B), EU rate (Panel C), and labor force shares (Panel D) for four AI adoption quartile groups (12-month moving average). Blue dashed line: first quartile; red dashed-dot line: second quartile; black long-dashed line: third quartile; solid yellow line: fourth quartile.
Source: Authors' calculation.

constitutes a small fraction in the labor force (2%) and the second quartile's job-finding rate remains the highest across groups and improved substantially in 2024–2025, partially offsetting the effect of the higher separation rate on the unemployment rate.

Figure A.4 displays the unemployment rate and EU/UE flow rates of the joint exposure/adoption categories when the adoption threshold is set at the 85th percentile.

software developers in construction, department stores, and health services, among others.

Figure A.4: UNEMPLOYMENT RATE AND FLOWS BY AI EXPOSURE (MEAN) AND ADOPTION (85%)



Notes: E-High (E-Low) denotes high (low) AI exposure, and A-High (A-Low) denotes high (low) AI adoption. For the AI exposure, high and low means above and below the mean AI exposure. For AI adoption, high and low means above and below the 85 percentile AI adoption rate in 2025. Panel A shows the unemployment rate; Panel B shows UE rates; Panel C shows EU rates; and Panel D shows labor force shares. **Source:** Authors' calculation.

C Additional Regression Results

C.1 Teleworkability and AI Intensity

We examine whether teleworkability has effects on labor market outcomes that are distinct from those of AI exposure and adoption. Consider the following model:

$$y_{it} = \alpha + \beta_1 \text{Exposure}_i + \beta_2 \text{Adoption}_i + \beta_3 \text{Telework}_{it} + \gamma X_{it} + \lambda_t + \delta_j + \varepsilon_{it}, \tag{C.1}$$

where y_{it} is one of four outcomes: the unemployment-to-employment (UE) and employment-to-unemployment (EU) transition indicators, usual hours worked, and the log hourly wage. Controls X_{it} include local labor market tightness, age, age squared, sex, and education; λ_t and δ_j are month and industry fixed effects. Standard errors are clustered at the individual level. The sample covers November 2022 through May 2026 and excludes agricultural occupations.

We include telework alongside the AI measures because AI-exposed occupations are disproportionately teleworkable, so raw associations between AI and outcomes could partly reflect remote-work amenities rather than AI itself. The results in Table A.4 show that the two channels are distinct and often pull in opposite directions.

For job loss, all three are stabilizing: telework reduces the monthly EU probability by 0.9 percentage points, and higher AI exposure and adoption are each associated with lower separation rates. For re-employment, however, the pattern diverges. Telework strongly accelerates job-finding, whereas both AI exposure and AI adoption are associated with *slower* re-employment among the unemployed—suggesting that occupational AI exposure signals churn that complicates matching, not merely the kind of remote-friendly work that facilitates search.

For hours, exposure and adoption pull in opposite directions: exposure is associated with 2.4 more hours per week, consistent with task complementarity raising labor demand at the intensive margin, while adoption is associated with 3.0 fewer hours, consistent with automation substituting for routine task time. The telework hours premium is small (0.5 hours), indicating

Table A.4: AI EXPOSURE, ADOPTION, AND LABOR MARKET OUTCOMES

	(1) UE transition	(2) EU transition	(3) Hours worked	(4) Log wage
AI exposure	-0.147***	-0.00527***	2.368***	0.115***
	(0.0233)	(0.000538)	(0.0985)	(0.00518)
AI adoption	-0.351***	0.00924***	-3.004***	0.780***
	(0.0912)	(0.00234)	(0.443)	(0.0551)
Telework for pay	0.731***	-0.00897***	0.542***	0.0833***
	(0.00661)	(0.000171)	(0.0438)	(0.00379)
Market tightness			0.00384***	0.000692***
			(0.000474)	(0.0000431)
Market tightness (lagged)	0.000110	-0.00000346**		
	(0.000124)	(0.00000156)		
Age	-0.00712***	-0.000596***	0.911***	0.0265***
	(0.00147)	(0.0000418)	(0.00847)	(0.000374)
Age squared	0.0000667***	0.00000522***	-0.00997***	-0.000261***
	(0.0000166)	(0.000000437)	(0.0000967)	(0.00000429)
Female	0.0302***	0.000682***	-3.197***	-0.130***
	(0.00804)	(0.000186)	(0.0390)	(0.00223)
High-school graduation	-0.0162	-0.00383***	3.065***	0.107***
	(0.0141)	(0.000571)	(0.0801)	(0.00316)
Some college/associate degree	-0.00766	-0.00490***	2.686***	0.146***
	(0.0149)	(0.000569)	(0.0838)	(0.00338)
College graduation	-0.0737***	-0.00523***	3.467***	0.338***
	(0.0154)	(0.000565)	(0.0849)	(0.00395)
Observations	20706	1105216	1494911	188096
R-squared	0.140	0.00638	0.104	0.321

Standard errors in parentheses

Standard errors clustered at the individual level (cpsidp).

Month and industry fixed effects included in all columns.

Sample: 2022:M11–2026:M5.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

the hours effects are driven by AI rather than remote-work selection.

For wages, all three carry positive and significant premia, but their magnitudes differ sharply. The adoption premium exceeds both the exposure premium and the telework premium. The persistence of large AI premia net of telework argues against the interpretation that AI-correlated earnings simply reflect the amenity value of remote work; the returns to actively using AI tools substantially exceed those to merely working in an AI-exposed occupation or working remotely.

C.2 Mincer regression for wage levels

We employ the Mincer regression to analyze effects of AI exposure and adoption and their effects in the post LLM era.

The Mincer wage regression is written as follows:

$$\ln w_{it} = \alpha + \beta_1 \,\text{Educ}_i + \beta_2 \,\text{Exper}_{it} + \beta_3 \,\text{Exper}_{it}^2 + \beta_4 \,\text{AIX}_{it} + \beta_5 \,\text{AIA}_{it} \tag{C.2}$$
$$+ \beta_6 \,(\text{AIX}_{it} \times \text{Post-LLM}_t) + \beta_7 \,(\text{AIA}_{it} \times \text{Post-LLM}_t) + \gamma X_{it} + \lambda_t + \delta_{j(i)} + \varepsilon_{it},$$

where $\ln w_{it}$ is the log hourly wage of individual i in month t; Educ_i is a vector of education dummies (high school, some college, and college graduation, relative to less than high school); $\text{Exper}_{it} = \text{age}_{it} - \text{YrsSchool}_i - 6$ is potential labor market experience; Post-LLM_t is an indicator equal to one from November 2022 onward; X_{it} includes gender; and λ_t and $\delta_{j(i)}$ are time and industry fixed effects. The AI exposure (AIX_{it}) and AI adoption (AIA_{it}) are standardized based on their respective moments estimated from 2023 onward, as the adoption measure is skewed and shows more variations from 2023. The sample period is 2012 onward, as the AI adoption is available from 2012.[40]

The baseline (Column 1, Table A.5) recovers standard Mincer patterns. Column (2) adds AI exposure and adoption, both of which carry large and significant wage premia conditional on education, experience, sex, and time and industry fixed effects. Column (3) reports the

[40] Earnings weights (`EARNWT`) are used for all wage regressions. Note that AIX_{it} varies over time as individual i changes occupation.

Table A.5: MINCER WAGE REGRESSION

	(1)	(2)	(3)
	Baseline	AI Factors	Post-LLM Interactions
AI Exposure (std.)		0.0405***	0.0427***
		(0.000568)	(0.000650)
AI Adoption (std.)		0.0356***	0.0679***
		(0.00197)	(0.00500)
AI Exposure (std.) × Post-LLM			-0.0146***
			(0.00121)
AI Adoption (std.) × Post-LLM			-0.0428***
			(0.00534)
Experience	0.0207***	0.0205***	0.0205***
	(0.000118)	(0.000118)	(0.000118)
Experience^2	-0.000320***	-0.000317***	-0.000318***
	(0.00000240)	(0.00000240)	(0.00000240)
Female	-0.120***	-0.135***	-0.135***
	(0.00116)	(0.00120)	(0.00120)
High School	0.138***	0.127***	0.127***
	(0.00159)	(0.00160)	(0.00159)
Some College / Associate	0.202***	0.182***	0.181***
	(0.00168)	(0.00170)	(0.00170)
College Graduate	0.442***	0.410***	0.409***
	(0.00203)	(0.00208)	(0.00208)
Observations	998812	987897	987897
R-squared	0.395	0.400	0.400

Standard errors in parentheses

Time and industry fixed effects included in all columns.

Standard errors clustered at individual level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

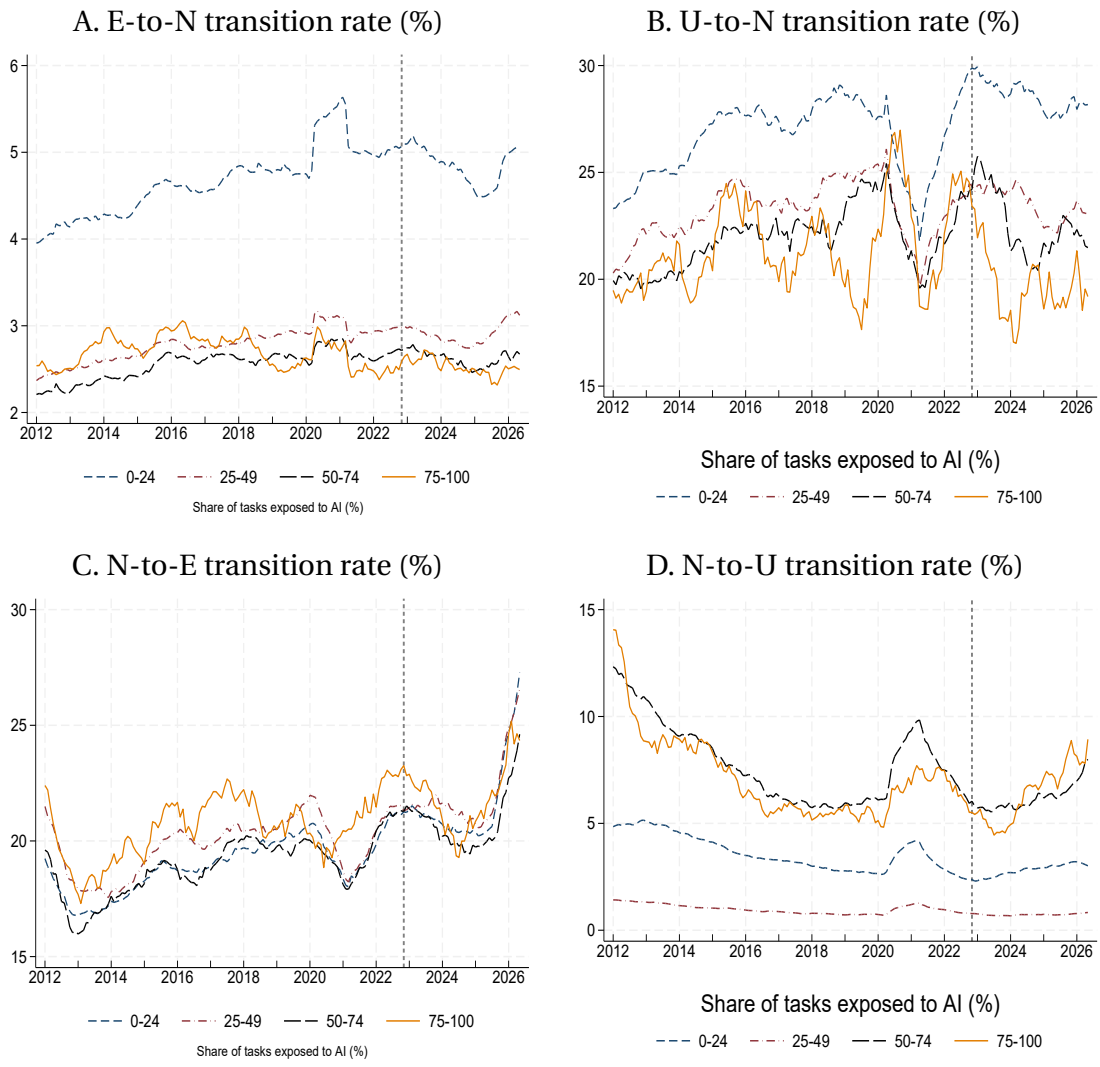
full results: The post-LLM interaction terms are negative and significant, suggesting negative effects of LLMs on the level of wages, though the net effects of each AI factor remain still positive (Column 3).[41]

[41] In principle, hourly wages are observed twice for each individual. We include individual fixed effects—dropping the time-invariant education and sex controls—and find that the baseline results hold: workers in more AI-intensive occupations earn higher wages, but this premium compressed in the post-LLM era.

D Stocks and Flows in the Participation Margin

We further examine the labor force participation rate and the employment-to-population ratio, along with flows between employment and nonparticipation and between unemployment and nonparticipation. These stocks and flows by AI exposure are affected by the imputation for missing occupation (the last step—the imputation based on demographic and education attributes—in particular).

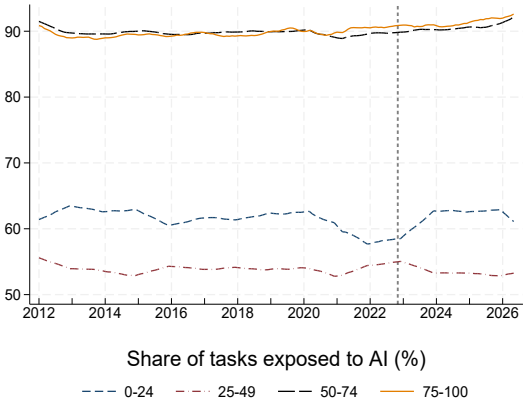
Figure A.5: EN, UN, NE AND NU FLOWS RATES BY AI EXPOSURE



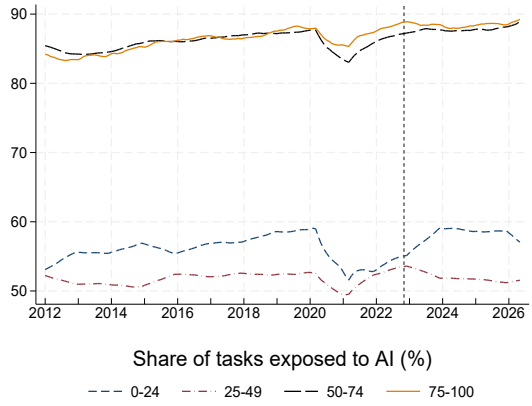
Notes: This figure displays the EN rate (Panel A), UN rate (Panel B), NE rate (Panel C), and NU rate (Panel D) for four AI exposure groups (12-month moving average). Blue dashed line: 0–25; red dashed-dot line: 25–50; black long-dashed line: 50–75; solid yellow line: 75–100. **Source**: Authors' calculation

Figure A.6: LFPR AND EPOP RATIO

A. Labor force participation rate (%)



B. Employment-to-population ratio (%)



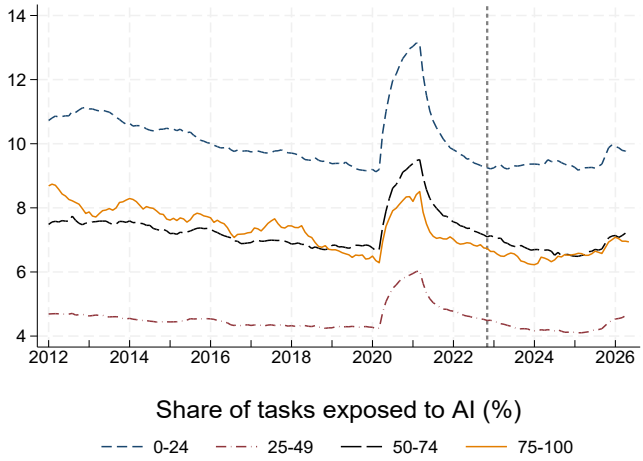
Note to figure: Panels A and B show LFPR and the EPOP ratio by AI exposure. Blue dashed line: 0–25; red dashed-dot line: 25–50; black long-dashed line: 50–75; solid yellow line: 75–100. **Source**: Authors' calculation

Figure A.5 displays the EN and UN rates by AI exposure (left panels) and AI adoption (right panels), capturing outflows from the labor force. Note that the EN rates are less affected by the imputation than the UN rates. Overall, both EN and UN rates are lower for workers with higher AI exposure and higher AI adoption. Higher AI exposure and adoption is associated with a lower likelihood of exiting the labor force overall.

Figure A.5 displays the NE and NU rates by AI exposure and adoption. Workers with higher AI exposure and adoption tend to exhibit higher NE and NU rates, reflecting stronger labor force attachment. Notably, these rates declined over time among the top-quartile adoption group. The relationship between AI factors and flows into the labor force is also nonlinear: both NE and NU rates are higher in the lowest-intensity group than in the middle groups, likely reflecting frequent transitions between participation and nonparticipation among the least-exposed workers.

Figure A.6 displays the LFPR and EPOP of workers by AI exposure and adoption. Reflecting the nonlinear relationship between NE/NU rates and AI factors, both LFPR and EPOP initially decline with increases in AI exposure or adoption and then rise, resulting in the highest levels for workers in the top quartiles of AI exposure and adoption. Following the COVID pandemic, LFPR and EPOP stepped down among workers in the top AI-adoption quartile, consistent with the reduced NU and NE rates for this group.

Figure A.7: WORKER REALLOCATION RATE: FLOWS PER CAPITA BY AI EXPOSURE (12-MONTH MOVING AVERAGE)



Note to figure: This figure displays the monthly worker reallocation rate by AI exposure. The unit for Y-axis is percent. **Source**: Authors' calculation

AI's effects on the worker reallocation So far, we have examined six different flow rates. To provide a comprehensive measure of worker reallocation and its magnitude, we compute the fraction of all six labor force status transitions between $t-1$ and t out of the civilian noninstitutional population at $t-1$, which we refer to as the monthly worker reallocation rate or the *monthly flows per capita*.

Figure A.7 displays the worker reallocation rate by AI exposure. The relationship between reallocation and AI exposure is nonlinear. Workers with the lowest AI exposure exhibit greater reallocation, reflecting churn associated with job instability. Meanwhile, the top-exposure group also reallocates more than the middle groups (those with 25–49% AI exposure), suggesting stronger job and labor force attachment — these workers are more likely to find jobs or return from nonparticipation.

Notably, the monthly flows per capita picked up among high-exposure workers — particularly the top exposure group — reflecting increased reallocation pressure (Figure A.7).

E Bayesian Estimation of U-star

E.1 Priors

Since $\sigma^{2,j}_{\eta,g} = 0$ for both flows, only the slope, cycle, and observation-noise variances are estimated. Each flow j is estimated independently, so all variance parameters carry a flow index j; the same prior form is used for both flows. To account for differences in scale across groups, the prior on $\sigma^{2,j}_{\zeta,g}$ is scaled by the ratio of each group's empirical EU rate's variance (computed outside the COVID-19 window) to that of group 1:

$$r_g = \frac{\widehat{\mathrm{Var}}(y^{\mathrm{EU}}_g)}{\widehat{\mathrm{Var}}(y^{\mathrm{EU}}_1)}.$$

The priors are:

$$\sigma^{2,j}_{\zeta,g} \sim \mathrm{IG}\big(a_\zeta,\ b_\zeta \cdot r_g\big) \qquad \text{(slope innovation)}, \tag{E.1}$$

$$\sigma^{2,j}_{\kappa,g} \sim \mathrm{IG}(a_\kappa,\ b_\kappa) \qquad \text{(cycle innovation)}, \tag{E.2}$$

$$\sigma^{2,j}_{\varepsilon,g} \sim \mathrm{IG}(a_\varepsilon,\ b_\varepsilon) \qquad \text{(observation noise)}, \tag{E.3}$$

where $\mathrm{IG}(a, b)$ denotes the inverse-gamma with shape a and scale b (prior mean $= b/(a-1)$). The same hyperparameter values are used for both flows; see Table A.6.

Table A.6: PRIOR HYPERPARAMETERS (IDENTICAL FOR $j \in \{\mathrm{EU}, \mathrm{UE}\}$)

Parameter	Description	Shape (a)	Scale (b)	Prior mean
$\sigma^{2,j}_{\zeta,g}$	Slope innovation	30	$0.05 \times r_g$	$0.0017\, r_g$
$\sigma^{2,j}_{\kappa,g}$	Cycle innovation	2	0.10	0.10
$\sigma^{2,j}_{\varepsilon,g}$	Observation noise	10	0.10	0.011

E.2 Posterior Simulation

The posterior is approximated by a Markov chain Monte Carlo (MCMC) algorithm that alternates between two blocks.

State draws. Conditional on the current variance parameters, the full trajectory $\{\alpha^j_{g,t}\}_{t=1}^T$ is drawn jointly using the forward-filter backward-sampler (FFBS) of Carter and Kohn [1994], a special case of the simulation smoother of Durbin and Koopman [2002]. The Kalman filter is run forward to obtain the sequence of filtered means and covariances $\{a_{t|t}, P_{t|t}\}$; states are then drawn recursively from right to left by conditioning each α_t on α_{t+1} and Y_t. The two trend states $(\tau^j_{g,t}, \beta^j_{g,t})$ are initialized diffusely; the AR(2) cycle state block $(c^j_{g,t}, c^j_{g,t-1})$ is initialized at its stationary covariance, obtained by solving the discrete Lyapunov equation. The AR(2) coefficients ϕ_1 and ϕ_2 are treated as fixed (calibrated) parameters.

Variance draws. Given the state draws, the disturbances $\zeta^j_{g,t}$, $\kappa^j_{g,t}$, and $\varepsilon^j_{g,t}$ are observed, and each variance parameter is drawn from its conjugate inverse-gamma full conditional:

$$\sigma^{2,j}_{\zeta,g} \mid \cdot \sim \mathrm{IG}\left(a_\zeta + \tfrac{T-1}{2},\; b_\zeta r_g + \tfrac{1}{2}\sum_{t=2}^{T}(\zeta^j_{g,t})^2\right),$$
$$\sigma^{2,j}_{\kappa,g} \mid \cdot \sim \mathrm{IG}\left(a_\kappa + \tfrac{2(T-1)}{2},\; b_\kappa + \tfrac{1}{2}\sum_{t=2}^{T}\left[(\kappa^j_{g,t})^2 + (\kappa^{j*}_{g,t})^2\right]\right),$$
$$\sigma^{2,j}_{\varepsilon,g} \mid \cdot \sim \mathrm{IG}\left(a_\varepsilon + \tfrac{T_{\mathrm{obs}}}{2},\; b_\varepsilon + \tfrac{1}{2}\sum_{t\in\mathcal{T}}(\varepsilon^j_{g,t})^2\right),$$

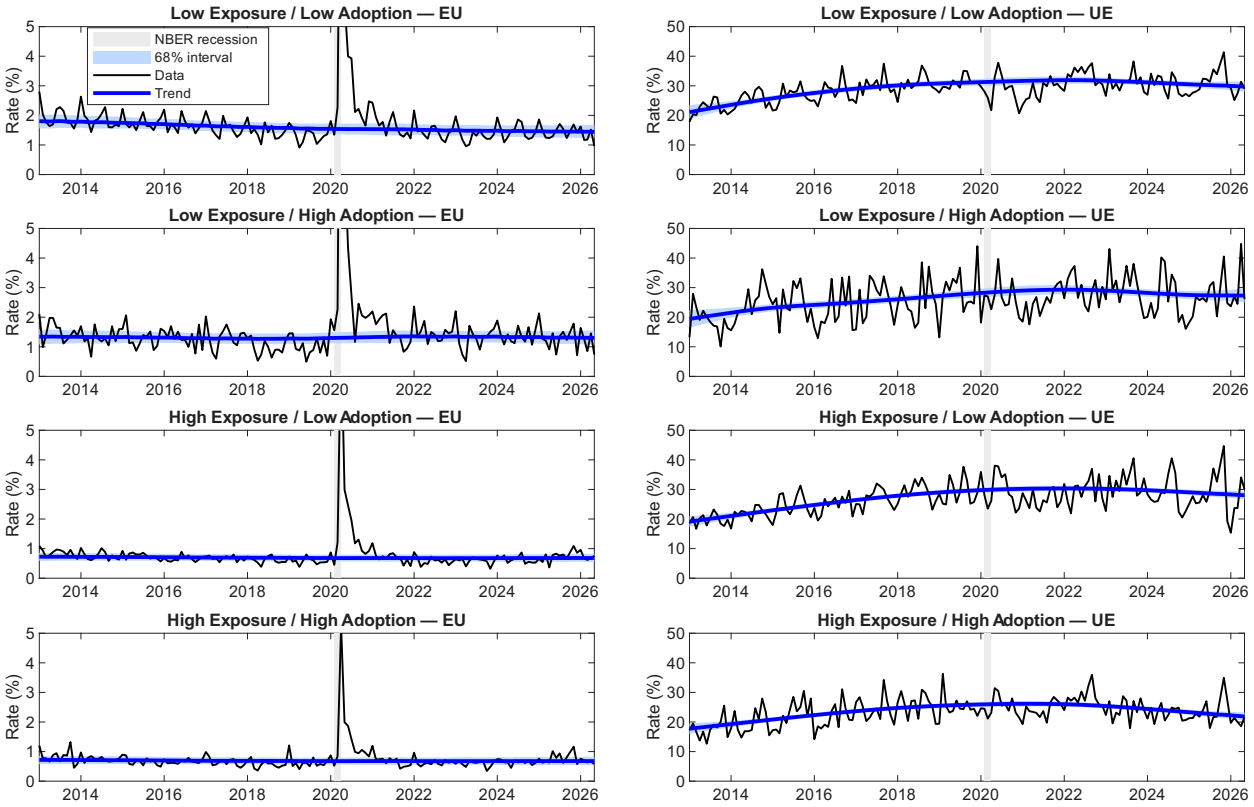
where $\mathcal{T}$ denotes the set of non-missing observation dates.

The sampler runs for 6,000 iterations, discarding the first 1,500 as burn-in and retaining every 5th draw, yielding $N = 900$ posterior draws. Posterior means and 68% credible intervals are reported from these draws.

E.3 Trend Estimates

Figures A.8 and A.9 display the estimated EU and UE rates and the natural rate estimates of the four AI exposure/adoption groups, respectively.

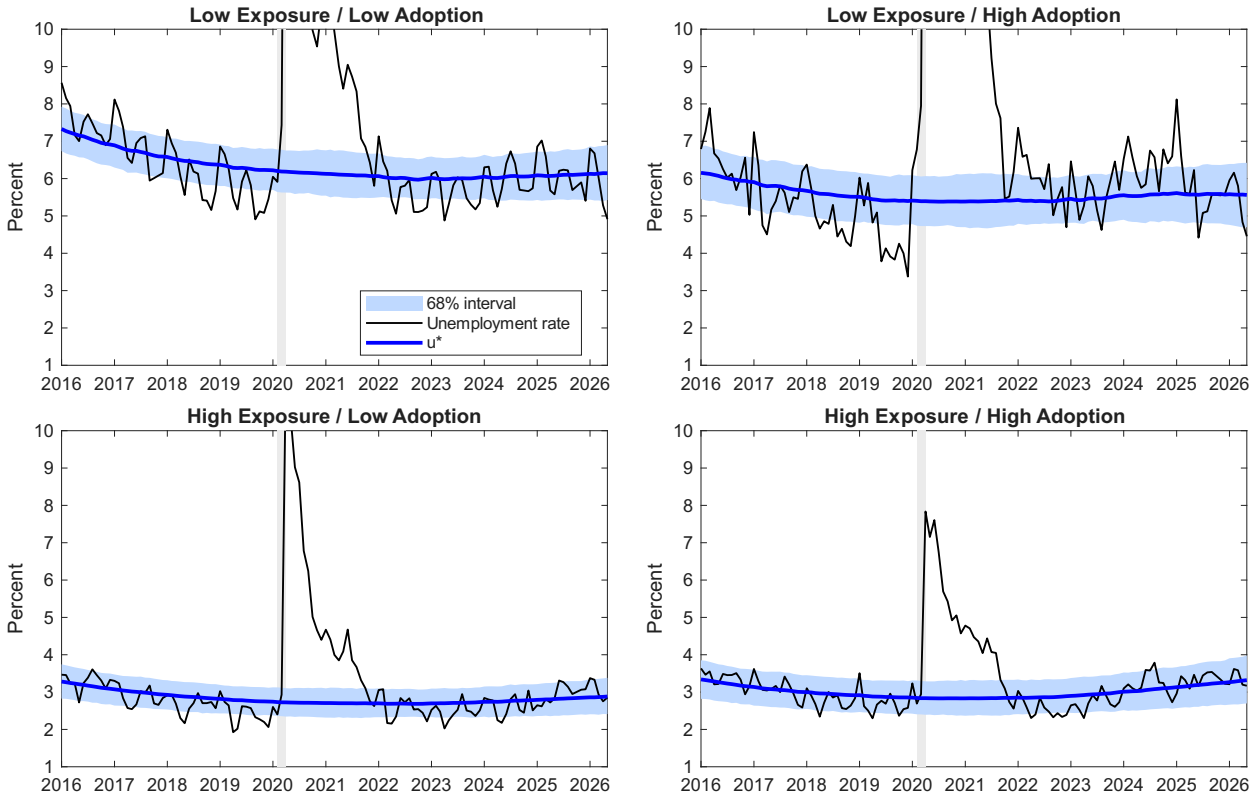
Figure A.8: EU AND UE TRENDS BY AI EXPOSURE AND ADOPTION



Note to figure: This figure displays unemployment inflow (EU) and outflow (UE) rates by AI exposure and adoption. The black lines represent the observed data, the blue lines denote the posterior median estimates, and the shaded blue areas indicate the 68 percent posterior intervals.

Note to figure: Authors' calculation.

Figure A.9: U-STAR ESTIMATE BY AI EXPOSURE AND ADOPTION



Note to figure: This figure displays the unemployment rate and the estimated natural rate for the four groups defined by AI exposure and adoption. The black lines represent the observed data, the blue lines denote the posterior median estimates, and the shaded areas indicate the 68 percent posterior intervals.

Note to figure: Authors; calculation.

F Attributes by AI Exposure and Adoption

F.1 Demographics

Table A.7 documents the demographic compositions of each group by AI exposure and adoption for 2012:M1-2026:M5.

Those with higher AI exposure are markedly more educated. Among the employed, the college graduation share rises from 13.1% in Q1 to 55.4% in Q4; among the unemployed, it rises from 8.2% to 44.5%. The share with some college or an associate degree, by contrast, is roughly flat across quartiles, suggesting that the education gradient in AI exposure is driven primarily by the college-educated rather than those with intermediate credentials.

High AI exposure is also associated with a slightly higher female share and a substantially lower share of young workers — the share aged 24 or under falls from 21.0% to 7.6% among the employed and from 32.3% to 15.0% among the unemployed as exposure increases. Prime-age workers (25–54) are correspondingly more prevalent in high-exposure jobs.

The AI adoption gradient is steeper on education but shows a gender pattern somewhat different from AI adoption. The college share among employed workers rises from 24.0% in Q1 to 74.0% in Q4, and from 13.0% to 67.8% among the unemployed — a nearly threefold increase across the adoption distribution. Unlike exposure, higher adoption is associated with a lower female share, falling from 47.5% to 36.6% among the employed and from 45.3% to 40.3% among the unemployed. Young workers are again underrepresented at higher adoption levels, while prime-age and older workers are more prevalent.

Panel C confirms that these patterns compound. The HH group is the most educated (college share of 75.2% employed, 70.0% unemployed), the least likely to be young (5.5% and 10.7%), and the low female share (36.5% and 40.2%). The LL group sits at the opposite end: least educated, most young, and with a higher female share. The LH and HL groups occupy intermediate positions, though the HL group — high exposure but low AI adoption — retains a notably high

Table A.7: DEMOGRAPHIC COMPOSITION BY AI EXPOSURE, ADOPTION, AND GROUP

Panel A: By AI exposure quartile

		Overall	Q1 (Low)	Q2	Q3	Q4 (High)
Share female	Employed	0.470	0.405	0.482	0.544	0.542
	Unemployed	0.462	0.405	0.487	0.582	0.591
Share college grad	Employed	0.384	0.131	0.511	0.498	0.554
	Unemployed	0.216	0.082	0.375	0.378	0.445
Some college/assoc.	Employed	0.274	0.288	0.258	0.288	0.269
	Unemployed	0.268	0.248	0.291	0.333	0.323
Share age ≤24	Employed	0.125	0.210	0.076	0.099	0.076
	Unemployed	0.283	0.323	0.164	0.186	0.150
Share prime-age (25–54)	Employed	0.644	0.595	0.668	0.664	0.698
	Unemployed	0.550	0.546	0.603	0.598	0.620
Share age 55+	Employed	0.231	0.195	0.256	0.237	0.226
	Unemployed	0.167	0.130	0.233	0.215	0.230

Panel B: By AI adoption quartile

		Overall	Q1 (Low)	Q2	Q3	Q4 (High)
Share female	Employed	0.470	0.475	0.508	0.451	0.366
	Unemployed	0.462	0.453	0.504	0.463	0.403
Share college grad	Employed	0.384	0.240	0.439	0.579	0.740
	Unemployed	0.216	0.130	0.319	0.463	0.678
Some college/assoc.	Employed	0.274	0.302	0.280	0.233	0.167
	Unemployed	0.268	0.278	0.296	0.266	0.197
Share age ≤24	Employed	0.125	0.170	0.090	0.067	0.057
	Unemployed	0.283	0.293	0.182	0.150	0.114
Share prime-age (25–54)	Employed	0.644	0.614	0.660	0.677	0.722
	Unemployed	0.550	0.555	0.604	0.610	0.629
Share age 55+	Employed	0.231	0.216	0.250	0.256	0.221
	Unemployed	0.167	0.151	0.214	0.240	0.257

Panel C: By AI exposure × adoption group

		Overall	Low–Low	Low–High	High–Low	High–High
Share female	Employed	0.470	0.424	0.362	0.537	0.365
	Unemployed	0.462	0.415	0.391	0.568	0.402
Share college grad	Employed	0.384	0.214	0.369	0.503	0.752
	Unemployed	0.216	0.118	0.294	0.392	0.700
Some college/assoc.	Employed	0.274	0.284	0.300	0.282	0.162
	Unemployed	0.268	0.259	0.278	0.326	0.192
Share age ≤24	Employed	0.125	0.182	0.098	0.075	0.055
	Unemployed	0.283	0.303	0.233	0.152	0.107
Share prime-age (25–54)	Employed	0.644	0.610	0.668	0.667	0.725
	Unemployed	0.550	0.551	0.569	0.613	0.632
Share age 55+	Employed	0.231	0.207	0.234	0.258	0.221
	Unemployed	0.167	0.145	0.198	0.235	0.261

Notes: Sample period 2012:M1–2026:M5.

Shares weighted by `wtfinl`. Panel C: Low–Low = low exposure & low adoption, etc.

female share (53.7% employed). The higher female share in the HL group relative to HH group suggests that high-exposure occupations with low AI adoption tend to be female-dominated, likely reflecting in-person service roles where practical barriers to AI adoption remain high.

F.2 Compositional Shift of Unemployment

Table A.8 examines how the demographic composition of unemployed workers shifted between 2012–2019 and 2023–2026, and how these changes vary across AI exposure, adoption, and group.

The most pronounced pattern across the AI exposure distribution is the sharp rise in college graduation rates, which increases monotonically with exposure: the college share among unemployed workers rose by 2.8 pp in Q1, 7.9 pp in Q2, 8.6 pp in Q3, and 23.0 pp in Q4. Correspondingly, the share with some college or an associate degree fell more steeply in higher-exposure quartiles (-11.8 pp in Q4 versus -2.6 pp in Q1), suggesting that educational upgrading of the unemployed pool is concentrated in highly AI-exposed occupations. The female share moved in the opposite direction: it was essentially unchanged in low-exposure quartiles but fell sharply in Q4 (-13.4 pp), indicating that women became a smaller fraction of the unemployed in the most AI-exposed jobs. Prime-age workers became somewhat more prevalent in high-exposure unemployment ($+5.0$ pp in Q4), while the age composition changed little in lower-exposure quartiles.

The AI adoption distribution tells a different story. The college share increased only modestly in Q1 ($+0.2$ pp) and Q4 ($+2.1$ pp) but declined substantially in Q2 (-6.6 pp) and Q3 (-11.1 pp), pointing to a non-monotonic compositional shift. The female share moved in the opposite direction to the exposure pattern: it rose in the upper adoption quartiles ($+3.9$ pp in Q3, $+4.8$ pp in Q4), suggesting that women became more represented among the unemployed in high-adoption occupations.

Panel C highlights that the Low–High group — low AI exposure but high adoption — experienced dramatic compositional shift. The college share rose quite a bit and so did the female share, while the share aged 24 or under fell by 10.5 pp and prime-age workers rose by 6.9 pp. By contrast, the high exposure and high adoption showed comparatively modest changes.

Table A.8: DEMOGRAPHIC COMPOSITION OF UNEMPLOYED WORKERS: 2012–2019 VS. 2023–2026

	Overall			Difference (2023–26 minus 2012–19) by quartile			
	2012–19	2023–26	Diff	Q1 (Low)	Q2	Q3	Q4 (High)
Panel A: By AI exposure quartile							
Share female	0.459	0.455	-0.004	0.007	-0.009	-0.035	-0.134
Share college grad	0.190	0.261	0.071	0.028	0.079	0.086	0.230
Share some college/assoc.	0.275	0.235	-0.041	-0.026	-0.058	-0.066	-0.118
Share age ≤24	0.292	0.295	0.002	0.013	0.013	0.001	-0.020
Share prime-age (25–54)	0.548	0.541	-0.006	-0.026	0.001	0.010	0.050
Share age 55+	0.160	0.164	0.004	0.013	-0.014	-0.011	-0.031
Panel B: By AI adoption quartile							
Share female	0.459	0.455	-0.004	-0.017	-0.000	0.039	0.048
Share college grad	0.190	0.261	0.071	0.002	-0.066	-0.111	0.021
Share some college/assoc.	0.275	0.235	-0.041	-0.035	-0.045	-0.009	-0.030
Share age ≤24	0.292	0.295	0.002	0.025	0.100	0.063	-0.020
Share prime-age (25–54)	0.548	0.541	-0.006	-0.023	-0.054	-0.018	0.030
Share age 55+	0.160	0.164	0.004	-0.003	-0.046	-0.045	-0.010
Panel C: By AI exposure × adoption group							
	2012–19	2023–26	Diff	Low–Low	Low–High	High–Low	High–High
Share female	0.459	0.455	-0.004	0.002	0.095	0.010	0.047
Share college grad	0.190	0.261	0.071	0.032	0.239	0.029	0.011
Share some college/assoc.	0.275	0.235	-0.041	-0.025	-0.090	-0.043	-0.028
Share age ≤24	0.292	0.295	0.002	0.013	-0.105	0.027	-0.018
Share prime-age (25–54)	0.548	0.541	-0.006	-0.022	0.069	-0.002	0.026
Share age 55+	0.160	0.164	0.004	0.009	0.035	-0.026	-0.008

Notes: Unemployed workers only. Shares weighted by `wtfinl`.

Diff = 2023–2026 minus 2012–2019. Panel C: Low–Low = low exposure & low adoption, etc.

Higher AI exposure and adoption tend to raise the share of college graduates in unemployment: even among low-exposure workers, higher adoption raises the college graduate share. It is also notable that the share of prime-age workers in unemployment rises with higher exposure and adoption while the share of young workers aged 24 or under decreases, suggesting that labor market outcomes of prime-age workers may also have been disproportionately affected by AI factors.

F.3 Reason for Unemployment

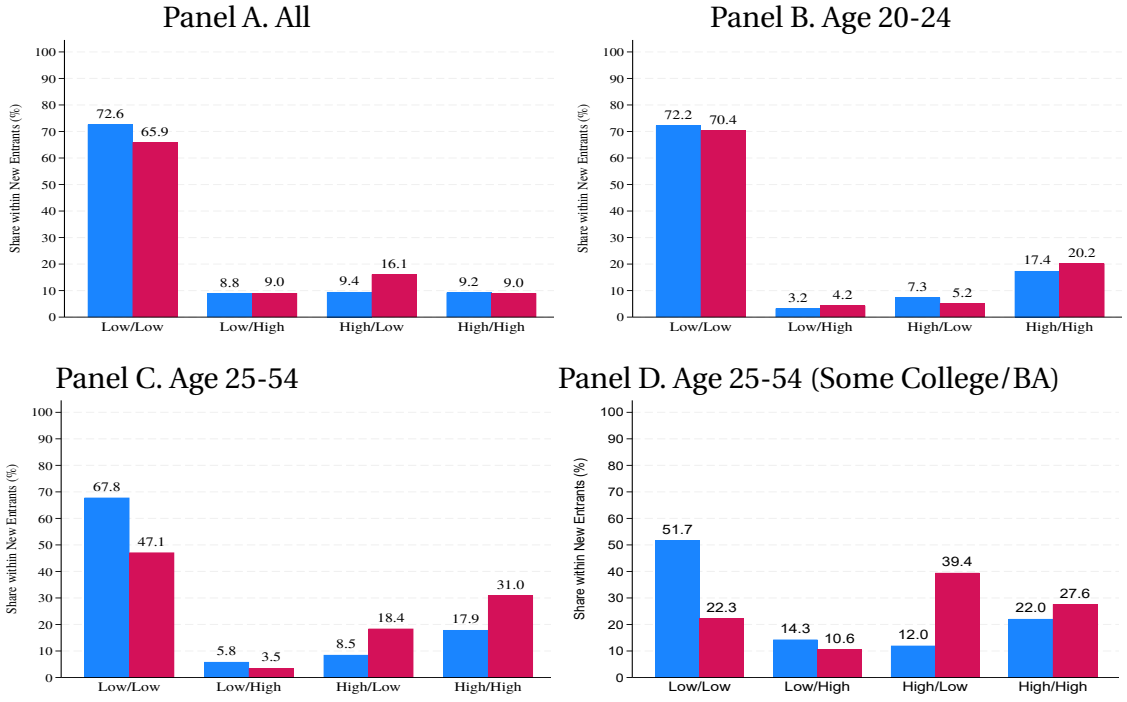
We examine whether AI exposure and adoption disproportionately affected workers searching for jobs for different reasons, focusing on compositional changes among new entrants — specifically college-educated workers aged 20–24 and prime-age workers — and among permanent job losers across age and education groups. We compare differences between 2022 and 2025 onward, as the share of new entrants can be affected by immigration and other post-pandemic economic developments.

Figure A.10 displays the average composition of new entrants to the labor force by AI exposure and adoption in 2022 and 2025–2026. Overall, no noticeable compositional changes are observed between the two periods (Panel A), and similarly so among new entrants aged 20–24. Among prime-age new entrants, however, the composition shifted toward high-exposure workers, particularly those with some college/associate degrees or college degrees. This suggests that the diffusion of LLMs likely raised unemployment among prime-age workers as well, although anecdotal and empirical evidence points more strongly to effects on early-career workers [Brynjolfsson et al., 2025, Hosseini and Lichtinger, 2025, Tucker, 2026].[42]

Figure A.11 displays the average composition of job seekers who experienced permanent separation by AI exposure and adoption in 2022 and 2025–2026. Overall, the share of high-exposure/high-adoption workers increased by 7.6 percentage points (Panel A), largely reflecting

[42]The qualitative conclusion remains largely robust, when imputing the AI exposure and adoption for new entrants who do not have any occupation information.

Figure A.10: COMPOSITION OF NEW ENTRANTS TO THE LABOR FORCE (BLUE: 2022; RED: 2025)

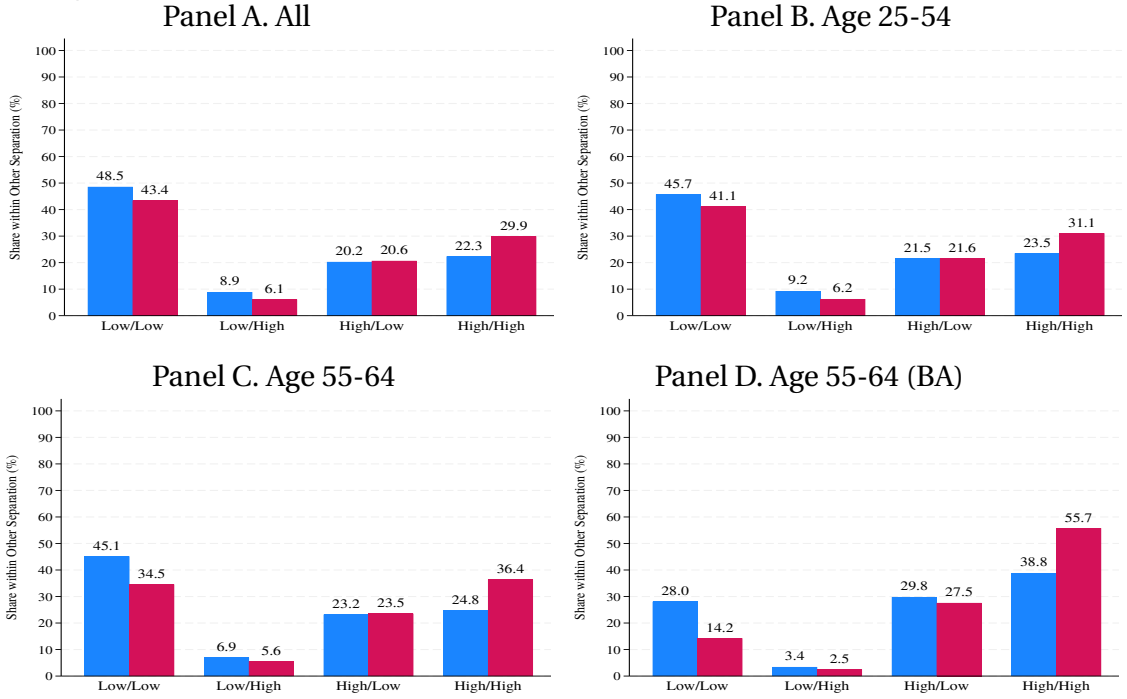


Notes: Bars display the share of each AI exposure/adoption group within new entrants (Panel A), new entrants aged 20–24 (Panel B), new entrants aged 25–54 (Panel C), and new entrants aged 25–54 with education beyond high school (Panel D). Blue bars denote 2022 shares and red bars denote 2025–2026 shares. **Source**: Authors' calculation.

a compositional shift among prime-age workers (Panel B). Notably, this increase was more pronounced among workers aged 55–64 with a college degree, suggesting that LLMs disproportionately raised job losses among those approaching retirement — again pointing to broad effects across age groups.

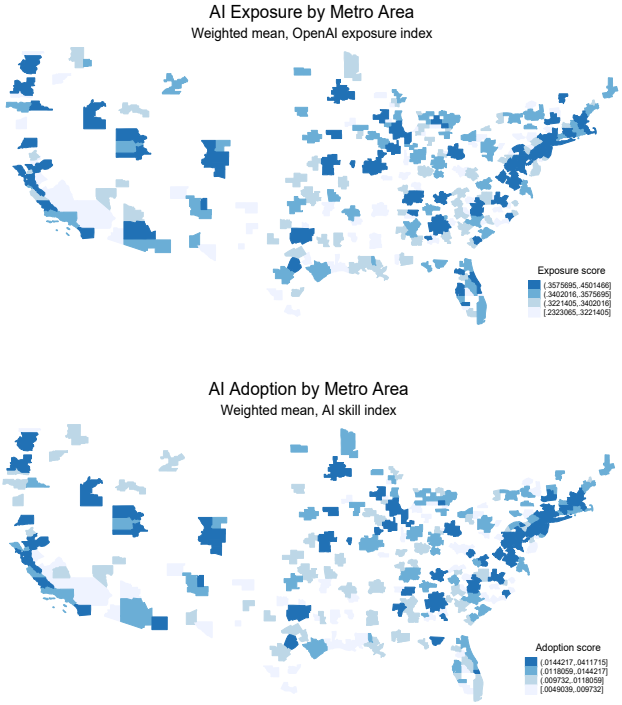
Beyond new entrants and permanent separations, LLMs do not appear to have created noticeable compositional shifts toward high-exposure/high-adoption groups. The evidence suggests that LLMs' effects extend beyond young recent college graduates to prime-age workers more broadly. Moreover, the increased prevalence of prime-age college graduates among permanent layoffs implies elevated mismatch risk and potentially long-term unemployment for these job losers, posing upside pressure on structural unemployment.

Figure A.11: COMPOSITION OF OTHER SEPARATIONS (BLUE: 2022; RED: 2025̃)



Notes: Bars display the share of each AI exposure/adoption group within new entrants (Panel A), new entrants aged 20–24 (Panel B), new entrants aged 25–54 (Panel C), and new entrants aged 25–54 with education beyond high school (Panel D). Blue bars denote 2022 shares and red bars denote 2025–2026 shares. **Source**: Authors' calculation.

Figure A.12: AI EXPOSURE AND ADOPTION BY METRO AREA



Notes: This figure displays AI exposure and adoption intensities across U.S. metropolitan areas. Darker shades indicate higher quartiles of AI intensity. The raw data are reported in Table A.9.: Authors' calculations.

Table A.9: TOP 50 METROPOLITAN AREAS BY AI EXPOSURE, 2023:M1–2026:M5

Rank	Metro Area	AI Exposure	AI Adoption
1	Columbia, MO	0.450	0.036
2	San Jose-Sunnyvale-Santa Clara, CA	0.429	0.032
3	Boulder, CO	0.427	0.027
4	Punta Gorda, FL	0.410	0.014
5	Washington-Arlington-Alexandria, DC-VA-MD-WV	0.408	0.023
6	Santa Fe, NM	0.406	0.022
7	Seattle-Tacoma-Bellevue, WA	0.402	0.023
8	Fort Collins, CO	0.399	0.021
9	San Francisco-Oakland-Hayward, CA	0.398	0.024
10	Madison, WI	0.397	0.019
11	Trenton, NJ	0.396	0.025
12	Pueblo, CO	0.394	0.041
13	California-Lexington Park, MD	0.393	0.024
14	San Luis Obispo-Paso Robles-Arroyo Grande, CA	0.389	0.018
15	Raleigh, NC	0.387	0.021
16	Bremerton-Silverdale, WA	0.385	0.030
17	Austin-Round Rock, TX	0.384	0.020
18	Boston-Cambridge-Newton, MA-NH	0.384	0.021
19	York-Hanover, PA	0.381	0.015
20	Huntsville, AL	0.381	0.018
21	Baltimore-Columbia-Towson, MD	0.381	0.017
22	Santa Rosa, CA	0.381	0.013
23	Nashville-Davidson–Murfreesboro–Franklin, TN	0.380	0.014
24	Wilmington, NC	0.378	0.029
25	Richmond, VA	0.377	0.015
26	Columbus, OH	0.376	0.016
27	Ogden-Clearfield, UT	0.375	0.015
28	Provo-Orem, UT	0.375	0.017
29	Manchester-Nashua, NH	0.374	0.016
30	San Diego-Carlsbad, CA	0.374	0.018
31	Atlanta-Sandy Springs-Roswell, GA	0.373	0.016
32	Albany-Schenectady-Troy, NY	0.373	0.015
33	Dallas-Fort Worth-Arlington, TX	0.372	0.016
34	Hartford-West Hartford-East Hartford, CT	0.370	0.015
35	Gainesville, FL	0.370	0.012
36	Denver-Aurora-Lakewood, CO	0.370	0.016
37	Knoxville, TN	0.370	0.015
38	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	0.370	0.017
39	Hilton Head Island-Bluffton-Beaufort, SC	0.368	0.012
40	Sacramento–Roseville–Arden-Arcade, CA	0.368	0.015
41	Tampa-St. Petersburg-Clearwater, FL	0.367	0.013
42	Colorado Springs, CO	0.367	0.017
43	Durham-Chapel Hill, NC	0.367	0.016
44	Kansas City, MO-KS	0.366	0.015
45	Phoenix-Mesa-Scottsdale, AZ	0.366	0.014
46	Springfield, IL	0.365	0.014
47	Charlotte-Concord-Gastonia, NC-SC	0.365	0.015
48	Minneapolis-St. Paul-Bloomington, MN-WI	0.364	0.015
49	Auburn-Opelika, AL	0.364	0.014
50	Portland-Vancouver-Hillsboro, OR-WA	0.364	0.016

Notes: Means weighted by `wtfinl`. Ranked by AI exposure.

F.4 Geography

Next, we discuss implications of spatial mismatch and sorting of AI exposure/adoption. Figure A.12 displays AI exposure and adoption by state, summarized by Table A.9. The metro-level table reveals substantial variation within states and highlights the role of specific local economies. Columbia, MO ranks first (0.450), likely reflecting the presence of the University of Missouri and a high share of professional and research workers. San Jose (0.429) and Boulder (0.427) — major technology hubs — follow closely, as do the Washington DC metro area (0.408) and Seattle (0.402).[43]

Metro areas with high AI exposure tend to have high AI adoption, though local factors—industry specialization, workforce composition, and institutional context—can drive gaps between the two at finer geographic levels. Since HH workers experience lower job-finding rates, geographic concentration of high-AI-intensity jobs may signal difficulties finding work locally, and the spatial dispersion of exposure and adoption points to potentially significant reallocation costs. Where HH jobs are geographically concentrated—for instance, in large metro areas—certain local labor markets may face greater reallocation pressure than others. Spatial mismatch, combined with frictions to geographic mobility, could slow reallocation if long-distance moves are required, though thick local labor markets may ease job transitions for some workers.

F.5 Labor Market Status and Outcomes

Table A.10 shows that labor market outcomes vary substantially across AI exposure, adoption, and exposure-adoption group, with adoption generally exhibiting stronger and more monotonic gradients than exposure. We focus on the post LLM period (2023:M1-2026:M5).

First, the share of telework increases with AI exposure and adoption. The share working from home for pay last week increases from 3.9% in Q1 to 56.8% in Q4, and the pre-COVID telework

[43] Several defense and government-oriented metros also rank highly, including Bremerton-Silverdale, WA, California-Lexington Park, MD, and Huntsville, AL. University towns such as Madison, WI, Fort Collins, CO, and Raleigh, NC also feature prominently, pointing to the role of research and education in driving AI exposure.

Table A.10: LABOR MARKET OUTCOMES BY AI EXPOSURE, ADOPTION, AND GROUP

Panel A: By AI exposure score

	Overall	Q1 (Low)	Q2	Q3	Q4 (High)
Hours worked (main job)	37.539	35.399	38.828	38.081	37.656
Hourly wage	23.621	20.800	26.986	24.345	26.486
Weekly earnings	1413.766	882.712	1660.132	1622.347	1988.709
Job tenure (years)	7.370	6.029	8.171	7.653	7.033
Share self-employed	0.102	0.089	0.132	0.061	0.067
Share government	0.134	0.087	0.173	0.131	0.079
Share telework for pay	0.209	0.039	0.227	0.383	0.568
Share telework before COVID†	0.097	0.024	0.113	0.157	0.232

Panel B: By AI adoption quartile

	Overall	Q1 (Low)	Q2	Q3	Q4 (High)
Hours worked (main job)	37.539	35.915	37.258	38.852	39.626
Hourly wage	23.621	21.103	24.202	25.844	31.575
Weekly earnings	1413.766	975.088	1229.596	1613.487	2261.457
Job tenure (years)	7.370	6.170	8.107	8.790	7.668
Share self-employed	0.102	0.101	0.089	0.116	0.106
Share government	0.134	0.090	0.172	0.166	0.136
Share telework for pay	0.209	0.065	0.146	0.286	0.481
Share telework before COVID†	0.097	0.035	0.095	0.151	0.237

Panel C: By AI exposure × adoption group

	Overall	Low–Low	Low–High	High–Low	High–High
Hours worked (main job)	37.539	36.043	37.440	38.443	39.738
Hourly wage	23.621	21.279	28.842	26.066	31.944
Weekly earnings	1413.766	972.743	1472.888	1550.849	2302.147
Job tenure (years)	7.370	6.344	8.956	8.340	7.651
Share self-employed	0.102	0.084	0.158	0.124	0.104
Share government	0.134	0.141	0.155	0.126	0.135
Share telework for pay	0.209	0.049	0.118	0.276	0.498
Share telework before COVID†	0.097	0.030	0.113	0.133	0.237

Notes: Employed workers (empstat = 1), 2023:M1–2026:M5.

Weekly earnings weighted by `earnwt`; job tenure by `jtsuppwt`; all others by `wtfinl`.

†Telework before COVID restricted to November 2023 only.

share follows the same gradient (2.1% to 22.7%), confirming that remote work was already disproportionately concentrated in high-exposure occupations before the AI era. Telework rises steeply with adoption as well, from 6.5% in Q1 to 48.1% in Q4. Even before the COVID pandemic, workers with higher AI exposure and adoption tended to telework.

Second, weekly earnings also rise markedly with exposure, from $883 in Q1 to $1,989 in Q4, though the hourly wage gradient is less monotonic, peaking in Q2 before rising again in Q4. Hourly wages rise monotonically from $21.1 in Q1 to $31.6 in Q4, and weekly earnings nearly double from $975 to $2,261. Hours worked also increase steadily with adoption (35.9 to 39.6 hours), indicating that high-adoption workers work both longer and at higher pay.

Panel C reveals an important asymmetry between the two dimensions. Workers in the Low–High group — low exposure but high adoption — earn higher hourly wages ($28.8) and have longer job tenure (9.0 years) than those in the High–Low group ($26.1 and 8.3 years), despite having lower AI exposure. This suggests that AI adoption contributes more to compensation and job stability than exposure alone. Unsurprisingly, the High–High group dominates on nearly every dimension: the highest hours (39.7), the highest hourly wages ($31.9), the highest weekly earnings ($2,302), and the highest telework share (49.8%), underscoring that jobs combining both high exposure and high adoption offer the most favorable labor market outcomes.

Job tenure is shortest among workers with the lowest AI exposure, lowest adoption, and in the Low–Low group, but does not follow a consistent gradient across either dimension. It peaks at Q2 of the exposure distribution (8.2 years) and Q3 of the adoption distribution (8.8 years), with the highest exposure and adoption quartiles showing lower tenure than the middle quartiles. The same pattern holds in Panel C: the High–High group has shorter tenure (7.7 years) than the Low–High group (9.0 years). This non-monotonicity suggests greater job turnover at the top of the AI intensity distribution, possibly reflecting increased churns in the most AI-intensive occupations.

Self-employment is lower in the highest exposure quartiles (6–7%) than in Q2 (13.2%), but highest in the Low–High group (15.8%), consistent with freelancers and independent contractors

adopting AI tools intensively despite working in lower-exposure occupations. Government employment peaks in the middle of both the exposure and adoption distributions and is lowest among the most AI-exposed workers (7.9% in Q4), while remaining relatively flat across the four exposure–adoption groups in Panel C.