AI and Labor-Market Reallocation*

Hie Joo Ahn† Nicholas A. Carollo‡

April 15, 2026

Abstract

This paper documents how artificial intelligence has affected stocks and flows dynamics of the U.S. labor market. Combining CPS and JOLTS data with occupation-level AI exposure measures from OpenAI and AI adoption measures derived from Lightcast job postings, we decompose 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-switch rates than other groups, while their within-job activity switching has increased noticeably relative to others. The AI factors tend to raise wage growth, and AI exposure is also associated with higher hours worked, but these effects have attenuated since 2023, suggesting that LLM-driven reallocation has operated through task reorganization within firms and a reconfiguration of labor input, alongside weaker demand for workers highly affected by 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 rate recovered from the trend components of unemployment inflows and outflows by AI exposure and adoption is estimated to have stepped up by about a quarter 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, mismatch, natural rate of unemployment, reallocation

*PRELIMINARY DRAFT. 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.

†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 remains deeply contested. 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 labor market reallocation. In this paper, we examine the effects of LLMs on aggregate labor market dynamics and functioning, focusing on their effects 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 microdata, JOLTS, and O*NET—with nontraditional data sources compiled by OpenAI, Anthropic, and Lightcast.

Our approach of measuring AI's effects has a few key aspects. First, we match individuals in the CPS microdata with the detailed occupation measures of AI exposure and AI adoption at the job. Specifically, we draw on occupation-level scores from Eloundou et al. [2023] and Handa et al. [2025], which characterize the extent to which generative AI can perform the tasks associated with each occupation. For AI adoption, we compute the share of positions that explicitly require AI-related skills from Lightcast job postings, capturing the extent to which firms have actually integrated AI into production. In other words, AI exposure captures the supply of skills that can be theoretically replaced by AI, and AI adoption measures actual demand of AI-related skills. By assigning AI exposure and adoption scores to each individual in the CPS, we can directly assess how AI has affected labor market stocks, flows, and outcomes in a comprehensive manner.

A key challenge in the matching process is missing occupation information. While all employed individuals report occupations, some unemployed individuals do not—particularly new entrants or those without recent work histories. For E-to-U (EU), U-to-E (UE), and E-to-N (EN) transitions between $t-1$ and t, we assign AI exposure and adoption using current-month (t) information when available, and otherwise use lagged ($t-1$) values. This leaves fewer than 10 percent of unemployed individuals without AI information. For these cases, we impute values using the individual-specific mean within the same CPS rotation-panel segment (MIS 1–4 and MIS 5–8), and then assign the individual's average across observed spells. After these steps, about 5 percent of unemployed individuals remain unmatched, primarily those without work histories or persistent non-employment. We further assign group averages based on demographic and educational characteristics of those who are employed for only one month during the survey. We use information from workers with only one employment spell to incorporate those with the weakest job attachment. This final step has negligible effects on the unemployment rate and on EU, UE, and EN flows.

We apply the same procedure to the remaining labor market flows and the labor force participation rate, although these measures are more affected by the final imputation step. We interpret the imputed AI exposure and adoption for individuals without occupation information as reflecting their potential exposure when employed or searching for a job.

Next, we estimate the market tightness by AI exposure and adoption to analyze the interaction between labor supply and demand as well as the labor market slack and functioning. For this, we combine individual-level stocks and flows by AI exposure and adoption with occupation-level vacancy data constructed from JOLTS and Lightcast.

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 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 jobs—the reduction in separations more than offsets this effect.

Second, since the introduction of LLMs, however, job-finding rates have declined more for workers in high–AI-exposure and high–AI-adoption jobs than for others. A similar pattern is observed in job-to-job transitions, while job-loss rates show little differential change across

groups. Consistent with this pattern, 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., 2025a, Hosseini and Lichtinger, 2025].

Third, LLMs have reshaped labor allocation within firms. The rate of within-job activity switching rises disproportionately for workers in high–AI-exposure and high–AI-adoption jobs, pointing to increased task reorganization. Consistent with this, we 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 suggest 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 job-finding rates, 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. We show that the cross-sectional dispersion in the hirability index has increased noticeably since 2023, 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.

Fifth, the overall empirical evidence points to upward pressure on the natural rate of unemployment in the post-LLM period. To quantify these effects, we use EU and UE rates for four worker groups defined by AI exposure and adoption (above/below mean exposure interacted with above/below median adoption), and estimate their trend components using a Bayesian trend–cycle decomposition. We then map these trends into group-specific natural rates using

Shimer [2012]'s steady-state model of unemployment. Finally, we aggregate across groups using labor-force shares to obtain the overall natural rate and its evolution in the post-LLM period. Reflecting the modest weakening of the job-finding rate for high-exposure workers relative to low-exposure workers, the natural rate is estimated to have risen by about a quarter percentage point between December 2022 and February 2026, albeit with considerable uncertainty surrounding this estimate.

This paper distinguishes itself from the existing literature along several dimensions. While most prior work examines AI's effects on specific occupations, firms, or demographic groups [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., 2025a, Hosseini and Lichtinger, 2025, Crane and Soto, 2026], this paper provides the first systematic analysis of how AI shapes aggregate U.S. labor market dynamics and functioning, drawing on a comprehensive examination of labor force flows and job creation by AI exposure and adoption. Rather than focusing on a single dimension of labor market outcomes—such as employment, productivity, or wages—our approach decomposes all worker transition rates by AI exposure, providing a more complete picture of how AI has reshaped the aggregate labor market.

This paper is also the first to construct individual-level hirability and separability indices to assess the effects of AI on job-finding and job-loss propensities, and their contribution to reallocation pressure. These indices extend the occupation-level adaptive capacity index of Manning and Aguirre [2026] to the individual level. Another contribution is to incorporate AI exposure and adoption directly into the estimation of the natural rate of unemployment, allowing us to assess the effects of LLMs on labor market functioning.

Section 2 describes the datasets used for the empirical analyses. Section 3 reports demographics and labor market outcomes by AI exposure and adoption, focusing on stocks and flows of workers and jobs. Section 4 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 5 discusses implications for the natural rate of unemployment and the trend labor force participation rate.

Related literature 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]. Moving beyond task exposure, [Manning and Aguirre, 2026] develop an occupation-level adaptive capacity index that measures how capable workers are of adjusting to job displacement. 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 exposure for specific occupations and demographic groups. 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]. 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., 2025a, Hosseini and Lichtinger, 2025] and reduced employment growth for coders in particular [Crane and Soto, 2026]. Consistent with our findings, these negative employment effects appear to be the result of reduced hiring rather than increased separations. Our paper complements this microeconomic evidence by providing the first systematic analysis of AI's influence on the aggregate labor market, in particular, its implications for structural unemployment.

In addition to employment effects and reallocation pressure, our analysis also speaks to the effect of AI on structural productivity. Although a number of experimental and quasi-experimental studies have shown that access to LLMs increase the productivity of individual workers across a variety of domains [Brynjolfsson et al., 2025b, Cui et al., 2024, Dell'Acqua et al., 2023, Noy and Zhang, 2023], the implications of task-level gains for aggregate productivity remain highly uncertain. Task-accounting approaches, for example, suggest a wide range of potential contributions to TFP growth ranging from essentially zero [Acemoglu, 2025] to as much as percentage point over the next decade [Filippucci et al., 2025].

Finally, our paper contributes to a broader macro-labor literature on labor market reallocation and the effects of technological change. 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 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.[1] 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. By showing that AI is the primary source of dispersion in individual-level hirability, we demonstrate that LLMs can be an important driver of reallocation pressure.

2 Data

We combine microdata from the Current Population Survey with information on occupations' AI exposure, demand for AI-related skills, and vacancy counts.

2.1 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

[1] **?** adopts the measure of Neumann and Topel [1991] for the empirical investigation of wage Phillips curve.

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 expose to occupations as described above. Beyond providing an independent measure of AI exposure, the Anthropic data also allow us to distinguish between applications that primarily automate versus augment human labor based on user interaction patterns identified by Anthropic.

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 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 postings that mention an AI-related skill requirement for each occupation 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.

2.2 Current Population Survey

We assign each individual in the CPS an AI exposure based on their reported occupation and the indices defined in Section 2.1. 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 2025:M12.

Individuals without prior work experience—such as young labor force entrants—are missing occupational information and therefore cannot be assigned an occupation-based AI exposure. Yet it is precisely these workers who may be most affected by AI [Brynjolfsson et al., 2025a]. We therefore we impute the AI-exposure score for individuals who do not report an occupation using the average exposure of other individuals with the same gender, age, and education.[2] This allows us to define exposure for the full population, a necessary step when studying macroeconomic aggregates. AI exposoure ranges from 0 (no exposure) to 1 (full exposure), similar to a probability. By multiplying the score (and one minus the score) by the individual sampling weight, we construct stocks and flows for fully exposed (high-exposure) and non-exposed (low-exposure) workers. This approach preserves accounting identities for stocks and flows within exposure groups and in the aggregate.

2.3 Vacancies

We compute vacancies by occupation using data from the Job Openings and Labor Turnover Survey (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 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,

$$V_{ot} = \sum_{i \in I} \frac{L_{oit}}{\sum_{o \in O} L_{oit}} \times V_{it}^{JOLTS}. \tag{2.1}$$

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

2.4 Imputations of missing AI exposure and adoption

AI exposure and adoption are task-based measures and therefore can be assigned only to individuals with available occupation information. Occupation data are observed for all employed workers, but only partially for others. In particular, unemployed individuals who were pre-

[2] We further consider cohort, but this extension does not change the result materially.

viously employed report their past occupation, and a small fraction of individuals out of the labor force also report prior occupation. In addition, some workers who were employed in the previous month but become unemployed in the current month lack occupation information, likely reflecting measurement error.

We adopt two imputation procedures.

For the unemployment rate and the EU, UE, and EN flows, we apply the following imputation procedure to maximize the use of each individual's reported AI exposure with minimum imputation.

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 9.5% of unemployed individuals who do not have AI information.

2. Step 2: For observations that remain missing, we impute exposure using the individual-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, all unemployed individuals are assigned AI exposure measures. After this step, about 5% of the individuals do not have AI information.

It is important to note that this procedure produces flows and stock of unemployment that are largely internally consistent. Therefore, for the analyses on unemployment, we use this imputation procedure without the small fraction of unemployed individuals without AI information. Without step 3, the results do not change significantly.[3]

The remaining unemployed individuals without AI information are the ones who do not find a job or have work histories in the adjacent months, and hence are likely to be those who continue to be unemployed or transition from and to nonparticipation. We consider the following approach for these individuals:

[3]This observation is consistent with the view that workers tend to return back to the previous industries or occupation after displacement (Hobijn et al.)

1. 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.

We then extend this methodology to assign AI measures to observationally similar individuals in nonparticipation.

2.5 Categorization of AI exposure and adoption

We classify workers into 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. For example, the 75-100% bin contains around 20 occupations, but these make up 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, telemarketers, and various clerical workers.

In addition, we define four adoption groups based on the share of AI-related terms in Lightcast job postings: below the first quartile, between the first and second quartiles, between the second and third quartiles, and the top quartile. These categorizations reveal the effects of AI more clearly than traditional approaches.

Last, we plot 12 month moving averages to smooth through some of the month-to-month noise and to handle seasonality.

3 Stocks and flows by AI exposure and adoption

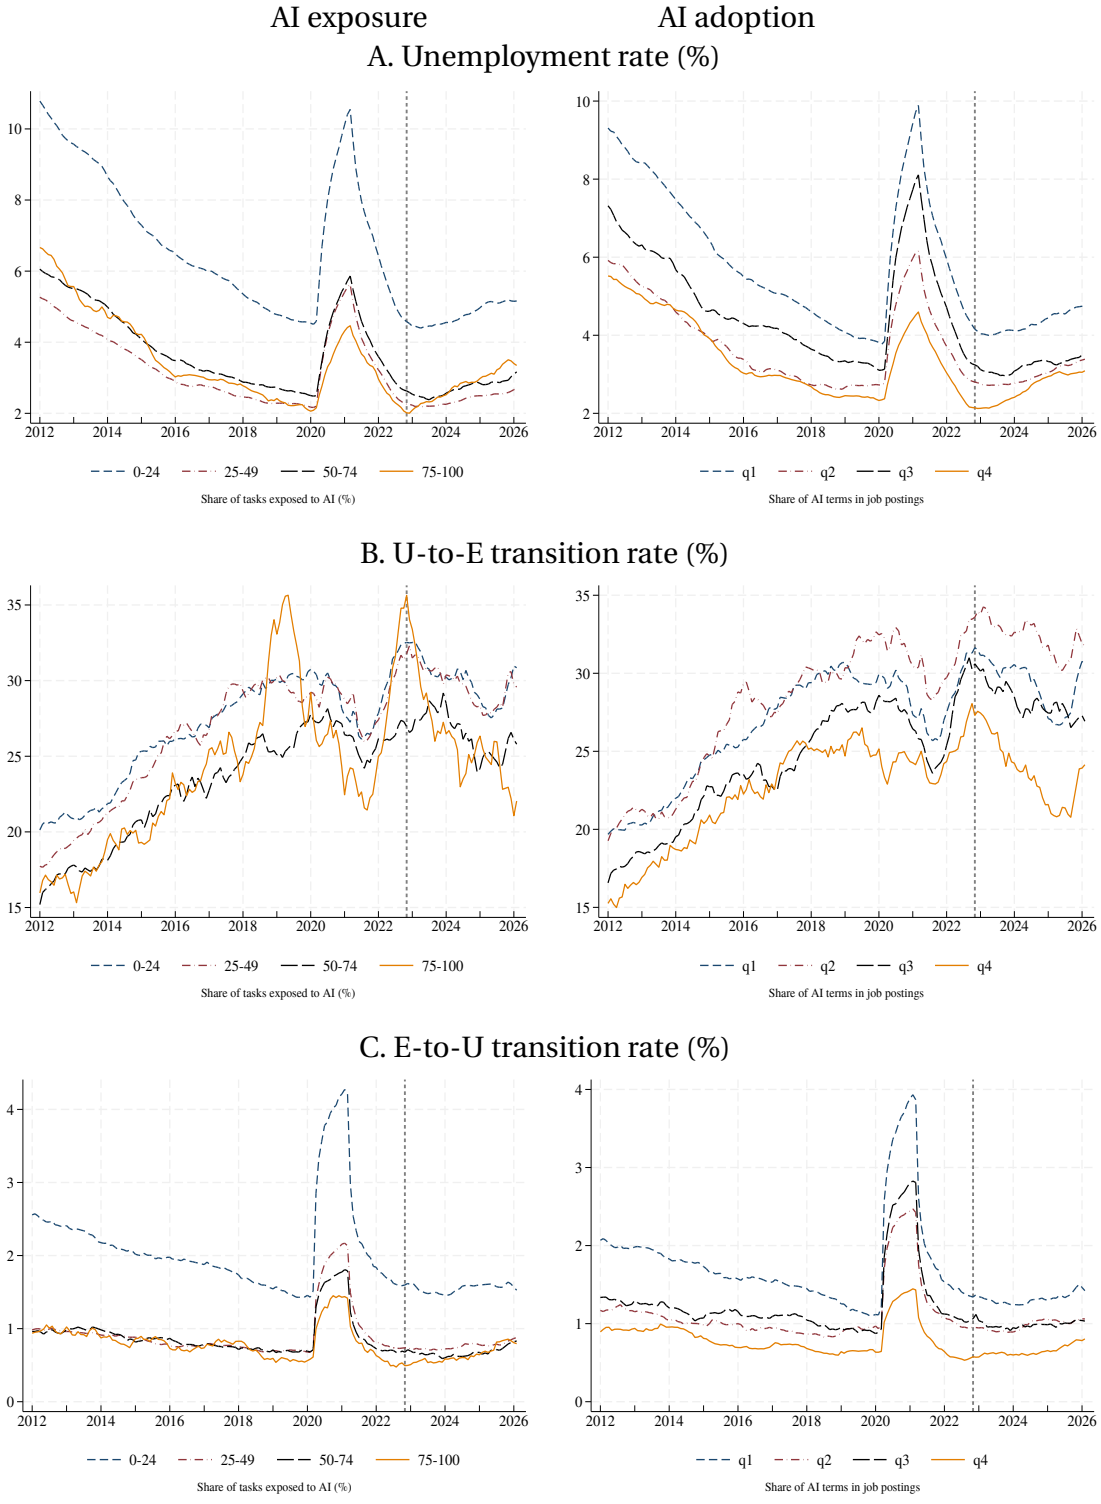
3.1 Stock and flows of unemployment and employment

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

Unemployment rate, job findings and job losses Panel A of Figure 1 displays the unemployment rate for four groups defined by AI exposure (left) and AI adoption (right).[4]

[4] Other statistics, including the LFPR, EPOP, and flows involving nonparticipation, are reported in the appendix.

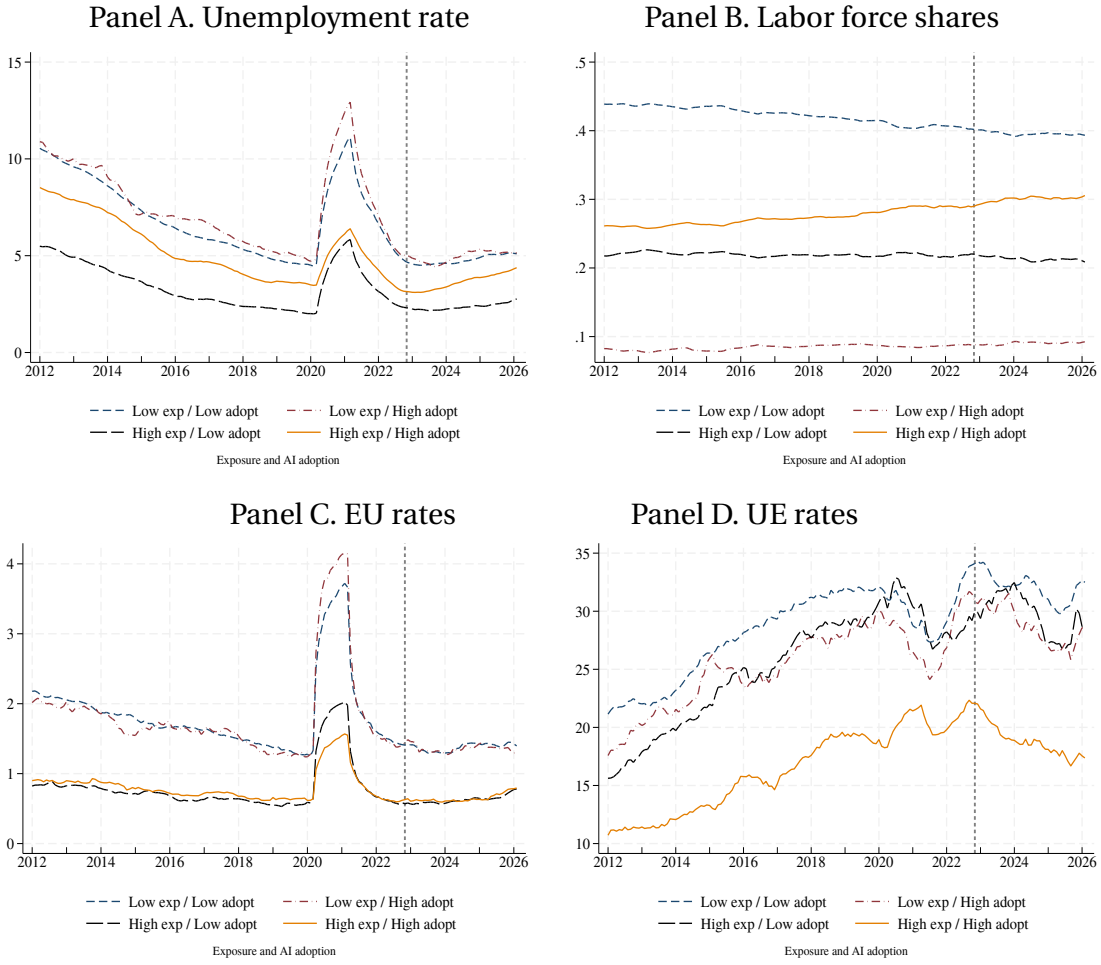
Figure 1: UNEMPLOYMENT RATE, AND UE AND EU RATES



Note to figure: Panels A, B, and C display the unemployment rate, and UE and EU transition probabilities, by AI exposure (left) and AI adoption (right).

Source: Authors' calculation

Figure 2: UNEMPLOYMENT RATE AND FLOWS BY AI EXPOSURE AND ADOPTION



Note: E-High (E-Low) denotes high (low) AI exposure, and A-High (A-Low) denotes high (low) AI adoption.
Source: Authors' calculation.

Several patterns emerge. First, both AI exposure and adoption are generally associated with lower unemployment rates, primarily through lower separation rates. Meanwhile, higher AI exposure and adoption are associated with lower job-finding probabilities. Workers with low AI exposure and adoption tend to have less stable employment and higher turnover, which raises their job-finding rates accompanied by high job-loss rates.[5] In contrast, job-finding spells may be longer for workers with high AI exposure and adoption, consistent with lower matching efficiency and, consequently, lower job-finding probabilities.

Second, the job-finding probabilities of workers with the highest AI exposure and adoption declined more than those of other groups following the introduction of LLMs. This pattern suggests a disproportionately larger decline in labor demand for these workers after the introduction of LLMs. The separation rate of top exposure group also rose a little faster than the other groups, but the difference is not noticeable. All told, reflecting this weakness in the hiring margin, the unemployment rates of top exposure and adoption group rose a bit faster than the other groups, although the difference is small.

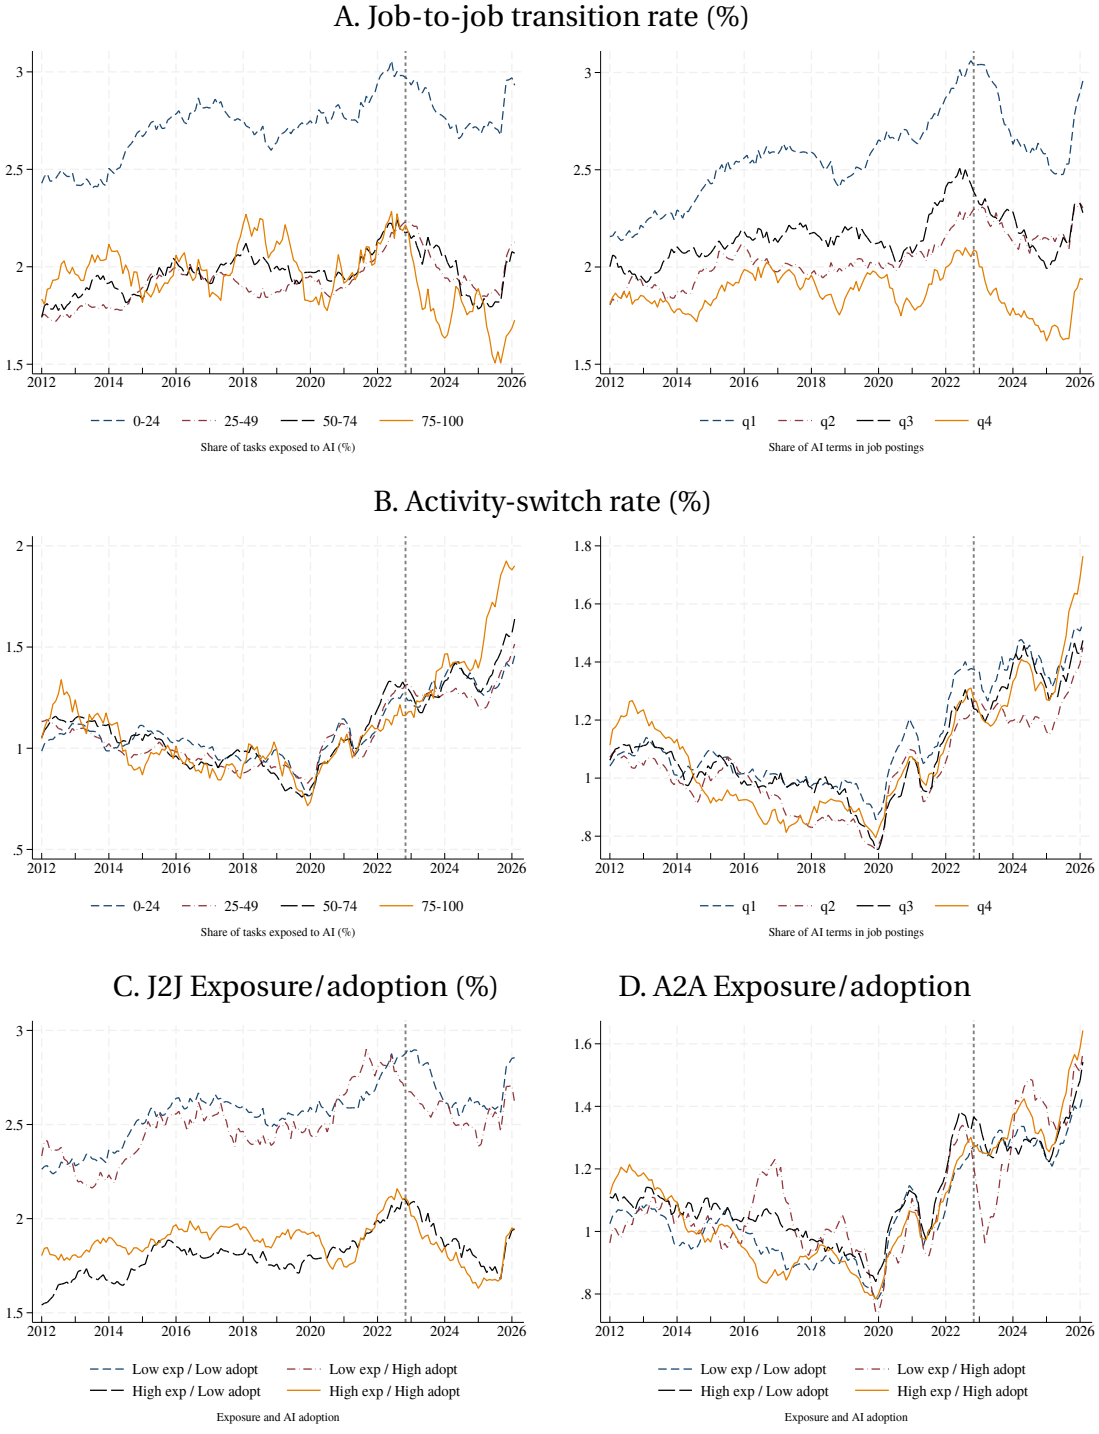
Switches of employers and job activities Figure 3 presents the job-switch rate (Panel A) and the activity-switch rate (Panel B) by AI exposure and adoption. The CPS asks respondents whether they, if employed in two consecutive months, changed employer since last month, which is used to compute the job-to-job transition rate. The CPS also asks respondents who have not changed 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 workers than for other groups. Job-switch rates fell across all exposure and adoption groups, likely reflecting weaker demand for them. Although overall job-switch rates began to recover in the second half of 2025, they continued to decline for the high-exposure group. Moreover, the job-switch rate of the top quartile of AI adoption is well below the pre-pandemic level, while those of the other groups are above their pre-pandemic levels.[6]

[5] Consistently, job creation and destruction rates are higher in the industries with low exposure.

[6] The job-to-job transition rate is about 0.5 percentage point higher in occupations with an AI-exposed task share of 25 percent or less than in other occupations. Meanwhile, the job-switch rates of the other groups are broadly

Figure 3: SWITCHES OF JOBS AND ACTIVITIES BY AI EXPOSURE



Note to figure: Panels A and B display the job-switch rate and activity-switch rate by AI exposure. For AI exposure, we use the share of tasks exposed to AI based on OpenAI's occupation-level estimates.

Source: Authors' calculation

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 rates than the other groups. A similar pattern is observed among the top quartile of AI adoption. This pattern suggests that LLM-driven reallocation has increasingly taken place within establishments and within current jobs.

All told, the evidence suggests that AI has not increased reallocation through job-to-job transitions; instead, it has dampened job switching, likely reflecting particularly weaker demand for these workers in a low-hiring environment. Rather, adjustment has occurred within jobs: while job-switch rates indicate reduced reallocation for the highest-exposure and adoption group—activity-switch rates rise disproportionately for that group. This pattern points to a buildup of unresolved reallocation pressure that is being absorbed through within-job task reorganization.

3.2 Market tightness by AI exposure and adoption

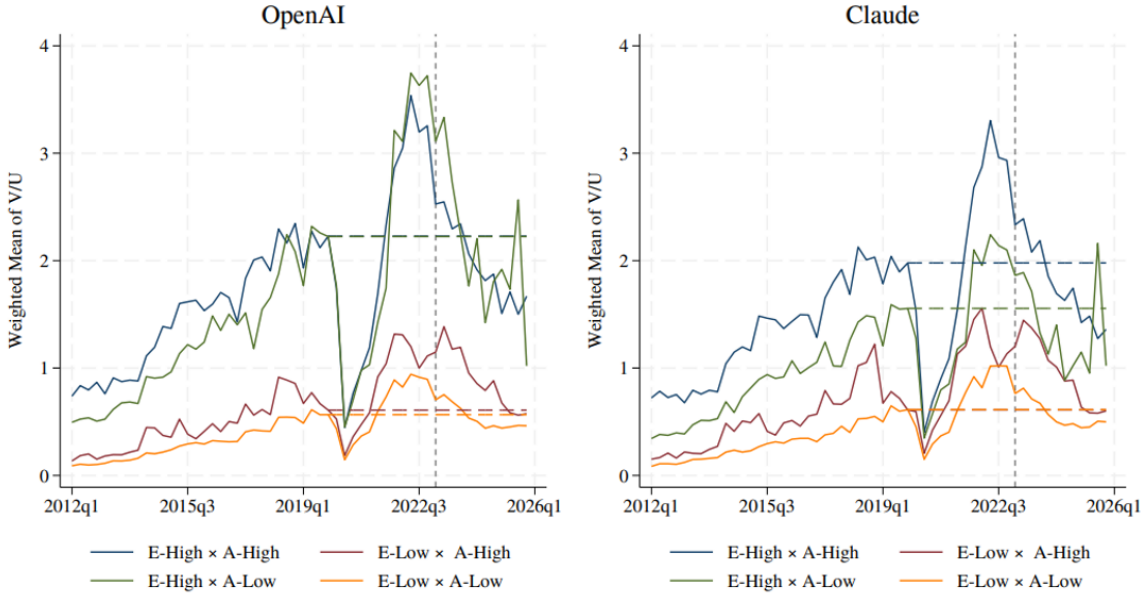
We compute market tightness for four groups of workers defined by above- and below-mean AI exposure and above- and below-median AI adoption shown in Figure 4. Panel A reports estimates based on OpenAI's AI exposure measure, while Panel B uses Anthropic's measure.

Overall, workers with high AI exposure exhibit higher market tightness, consistent with their lower unemployment rates and lower EU transition rates. In addition, workers in occupations with higher AI exposure and adoption tend to experience higher market tightness. Despite this, their job-finding rates are lower than those of other groups, reflecting lower matching efficiency and generally longer job-search spells among high-skilled workers.

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. that the evolution of tightness across groups cannot be fully explained by aggregate cyclical conditions.

similar. Given that low AI exposure occupations tend to be lower skilled, this pattern is consistent with the evidence that job-to-job transitions are more frequent among workers in low-skilled or less stable jobs (Ahn et al., 2023).

Figure 4: MARKET TIGHTNESS BY AI-EXPOSURE AND ADOPTION



Note: E-High (E-Low) denotes high (low) AI exposure, and A-High (A-Low) denotes high (low) AI adoption.
Source: Authors' calculation.

Instead, the market tightness among high-exposure workers points to structural forces associated with AI adoption. 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. This is consistent with firms adjusting labor demand through task reorganization and partial substitution of labor with AI technologies.

Evidence from market tightness is consistent with patterns in unemployment flows, job-switching, and activity-switching rates, pointing to uneven 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.

3.3 Wage inflation and hours worked

We discuss the effects of AI and AI-driven reallocation on wages and hours worked.

Consider a regression model for the wage growth:

$$w_{it} = \beta_1 openAI_{it} + \beta_2 \left(openAI_{it} \times llm_dum_t\right)$$

$$
\begin{aligned}
&+ \beta_3\, ai_skill_{it} + \beta_4\, (ai_skill_{it} \times llm_dum_t) \\
&+ \beta_5 \theta_{it} + \beta_6\, age_{it} + \beta_7\, age_{it}^2 \\
&+ \gamma' X_{it} + \delta_t + \eta_{j(i)} + \varepsilon_{it},
\end{aligned}
$$

where w_{it} is the 12-month nominal hourly wage growth of individual i at time t, $openAI_{it}$ denotes the individual's AI exposure, ai_skill_{it} captures AI adoption at the individual's industry, llm_dum_t is a dummy variable equal to one for 2023 onward, θ_{it} denotes market tightness, age_{it} and age_{it}^2 are the individual's age and its square, X_{it} includes gender and education controls, δ_t represents time fixed effects, and $\eta_{j(i)}$ denotes industry fixed effects.

Table 1 reports the regression results. AI-related factors are important determinants of wage inflation: Column (1) reports the result without industry fixed effects (`ind1990`) and column (2) reports the result with industry fixed effects. Both AI exposure and AI skill share are positively and statistically significantly associated with wage inflation, suggesting that workers in more AI-exposed occupations and industries with higher AI adoption experience faster wage growth. However, the interaction terms with the LLM period are negative and statistically significant, implying that these positive effects are attenuated after the introduction of LLMs. Quantitatively, the coefficients indicate a substantial reduction in the marginal impact of both AI exposure and AI skill share on wage inflation in the post-LLM period. This pattern is consistent with the interpretation that 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 hirings for these workers. The inclusion of time fixed effects, industry fixed effects, and industry-level market tightness suggest that the negative effects of LLM's are unlikely to reflect cyclical weakness in labor demand.

We replace the wage growth with hours worked to assess AI's effects on hours worked. Table 2 reports the regression results of hours worked. Overall, AI exposure is positively and significantly associated with hours worked, indicating that workers in more AI-exposed occupations tend to supply more labor and work longer hours. However, the interaction with the LLM period (2023 onward) is negative and highly significant, suggesting that this positive relationship weakens following the introduction of LLM technologies. In contrast, AI adoption—as captured by the AI skill share—has a negative baseline effect on hours worked, consistent with automation

reducing labor input. Importantly, the interaction with the LLM period is positive and statistically significant, indicating that this negative effect is attenuated, and partially reversed, in the post-LLM period. These findings suggest a shift in the role of AI from labor-substituting to more labor-complementary along the adoption margin, while the labor-supply effect associated with AI exposure diminishes over time. The similarity of these patterns in the post-2023 sample reinforces the interpretation that LLM technologies have materially altered how AI exposure and adoption translate into labor input decisions.

All told, AI's normal upside effects on nominal wage growth and hours worked attenuated from 2023, all suggesting that 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.

Table 1: DETERMINANTS OF WAGE INFLATION

	(1)	(2)
AI exposure	3.283***	3.097***
	(0.401)	(0.432)
AI exposure × LLM period	-1.518*	-1.538*
	(0.855)	(0.854)
AI skill share	51.14***	37.43***
	(13.55)	(12.76)
AI skill share × LLM period	-33.42**	-24.23
	(15.84)	(15.08)
Market tightness	0.0157***	0.0135**
	(0.00547)	(0.00549)
Age	-0.219***	-0.364***
	(0.0271)	(0.0285)
Age squared	0.00131***	0.00274***
	(0.000308)	(0.000321)
Female	-1.686***	-1.106***
	(0.149)	(0.169)
High-school graduation	1.583***	1.264***
	(0.237)	(0.241)
Some college and associate degree	1.258***	0.840***
	(0.241)	(0.247)
College graduation	2.258***	1.739***
	(0.287)	(0.297)
Observations	297788	297788
R-squared	0.00451	0.00684

Standard errors in parentheses
Standard errors are clustered at the individual level (cpsidp).
,
Month fixed effects are included in all columns.
,
Industry fixed effects are included in column (2).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2: DETERMINANTS OF HOURS WORKED

	(1)	(2)
AI exposure	3.625***	3.211***
	(0.0446)	(0.0478)
AI exposure × LLM period	-0.709***	-0.721***
	(0.0947)	(0.0924)
AI skill share	-9.641***	-9.674***
	(0.500)	(0.502)
AI skill share × LLM period	5.139***	5.409***
	(0.656)	(0.652)
Market tightness	0.00661***	0.00469***
	(0.000330)	(0.000324)
Age	0.990***	0.894***
	(0.00416)	(0.00411)
Age squared	-0.0108***	-0.00984***
	(0.0000481)	(0.0000474)
Female	-4.701***	-3.797***
	(0.0169)	(0.0186)
High-school graduation	2.782***	2.549***
	(0.0357)	(0.0350)
Some college and associate degree	2.491***	2.388***
	(0.0369)	(0.0364)
College graduation	3.577***	3.772***
	(0.0369)	(0.0374)
Observations	7836194	7836194
R-squared	0.0845	0.108

Standard errors in parentheses
Standard errors are clustered at the individual level (cpsidp).
,
Month fixed effects are included in all columns.
,
Industry fixed effects are included in column (2).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4 Indices of hirability and separability and the latent reallocation pressure

This section formally analyzes the effects of AI exposure and adoption on job finding and job loss. While the preceding analysis documents differences in labor market outcomes by AI exposure and adoption, these differences may reflect factors correlated with AI measures, such as education and industry composition. The purpose of this section is twofold. First, we examine the effects of individual characteristics, along with AI exposure and adoption, on an individual's likelihood of job finding and job loss. Second, we assess the extent to which AI exposure and adoption account for the dispersion in job-finding and job-loss prospects across individuals.

To address these issues, we estimate individual-level regressions of job-finding incidences on AI exposure and AI adoption, controlling for a range of worker 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. Using these indices, we examine the implications of AI-driven disparities on the build-up of reallocation pressure.

4.1 Models

Regression models Consider the baseline pooled regression:

$$\begin{aligned} UE_{it} &= \alpha + \beta_1 AIX_{it} + \beta_2 AIA_{it} + \beta_3 AIA_{it}^2 + \beta_4 \theta_{i,t-1} && (4.1) \\ &+ \beta_5 age_{it} + \beta_6 age_{it}^2 + \beta_7 sex_i + \beta_8 edu_i + \beta_9 mis_{it} + \varepsilon_{it}, && (4.2) \end{aligned}$$

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, while AIA_{it} and AIA_{it}^2 denote the AI adoption rate and its squared term, respectively; $\theta_{i,t-1}$ is lagged market tightness (vacancies-to-unemployment ratio) in the individual's industry; 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 restrict our sample period from 2023 onward, because the observations of AI adoption in the job postings are large enough during this period.[7]

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:

$$\widehat{UE}_{it}^{wo} = \widehat{UE}_{it} - \widehat{\beta}_1 AIX_{it} - \widehat{\beta}_2 AIA_{it} - \widehat{\beta}_3 AIA_{it}^2. \tag{4.3}$$

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 (4.2), we construct the *separability* index, denoted $\widehat{EU}_{it}$. We then apply equation (4.3) 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. Higher dispersion implies that a larger subset of individuals would have difficulty in moving to sectors with more favorable demand for them and higher market tightness. In the next section, we formally examine the implications of increased dispersion in job-finding probabilities for structural unemployment.

The weighted cross-sectional standard deviation is

$$sd_t^w(UE) = \sqrt{\frac{\sum_i w_{it}\left(\widehat{UE}_{it} - \mu_t^w\right)^2}{\sum_i w_{it}}}, \tag{4.4}$$

[7]We consider the squared AI exposure, but the coefficient is not statistically significant.

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

$$\mu_t^w(UE) = \frac{\sum_i w_{it}\widehat{UE}_{it}}{\sum_i w_{it}}. \tag{4.5}$$

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

$$sd_t^{wo}(UE) = \sqrt{\frac{\sum_i w_{it}\left(\widehat{UE}_{it}^{wo} - \mu_{t,wo}^w\right)^2}{\sum_i w_{it}}}, \tag{4.6}$$

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

$$\mu_t^{wo}(UE) = \frac{\sum_i w_{it}\widehat{UE}_{it}^{wo}}{\sum_i w_{it}}. \tag{4.7}$$

4.2 Estimated coefficients

Table 3 reports estimates of the determinants of job-finding (UE) and job-separation (EU) probabilities from 2023. A central result is that AI exposure is negatively associated with both transitions. The coefficient is economically and statistically significant for job finding, indicating that workers in more AI-exposed occupations are less likely to transition from unemployment to employment. A smaller but still significant negative effect is also present for separations, suggesting that AI exposure is associated with slightly lower job loss probabilities. Taken together, these patterns suggest that AI exposure tends to reduce labor market fluidity for highly exposed workers, with particularly pronounced effects on job-finding rates.

AI adoption exhibits a nonlinear relationship with both hirability and separability aligning with the evidence from the UE and EU transition rates by AI adoption quartile. The linear term enters negatively in the UE regression but positively in the EU regression, while the squared term reverses these effects. This implies a convex relationship for job-finding probabilities and a concave relationship for separations. At low levels of adoption, greater AI use is associated with lower job-finding rates and slightly higher separations, consistent with displacement effects. However, at higher levels of adoption, the positive quadratic term in the UE equation suggests some recovery in job-finding prospects, potentially reflecting complementarities or productivity gains that raise labor demand. Conversely, the negative quadratic term in the EU regression

implies that separation rates decline at higher adoption levels.

Market tightness plays a limited role in explaining job-finding probabilities in this specification, with a small and statistically insignificant coefficient in the UE regression. In contrast, tighter labor markets are associated with significantly lower separation probabilities.[8] Note that industry-level heterogeneity and supply-demand conditions are captured by θ_{it}. In addition, since AI exposure varies at the occupation level, the model also accounts for an individual's occupation as well.

Overall, the results suggest that the diffusion of AI is associated with a reconfiguration of labor market dynamics, primarily through reduced job-finding rates for highly exposed workers and nonlinear effects of adoption. These findings suggest a decline in matching efficiency for high-exposure groups and provide micro-level evidence supporting a modest increase in structural unemployment in recent years.

[8] Demographic controls display expected patterns: job-finding probabilities decline modestly with age, while separations decrease more sharply with age but at a diminishing rate, as captured by the positive quadratic term. Education gradients are more pronounced in separations than in hirability, with more educated workers experiencing significantly lower job loss probabilities. Gender differences are modest, with women exhibiting slightly higher job-finding probabilities and no statistically significant difference in separations.

Table 3: DETERMINANTS OF HIRABILITY AND SEPARABILITY

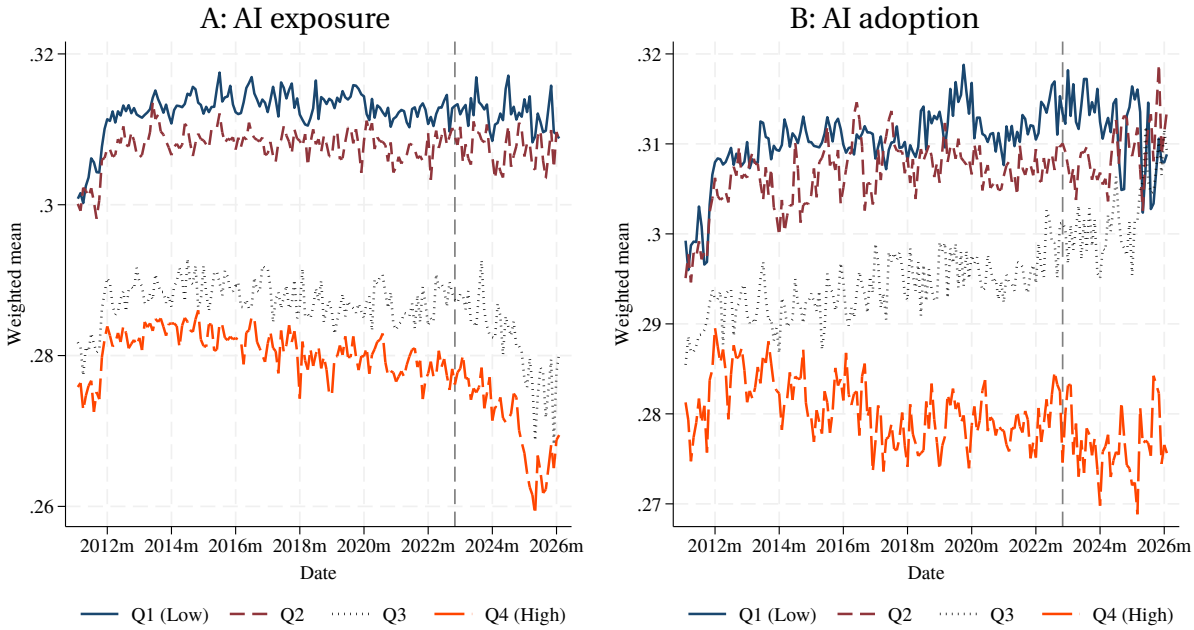
	(1) UE	(2) EU
AI Exposure	-0.0565***	-0.0114***
	(0.0137)	(0.000534)
AI Adoption	-0.696***	0.0178***
	(0.135)	(0.00408)
AI Adoption^2	1.167***	-0.0224**
	(0.331)	(0.00880)
Market Tightness	-0.0000686	-0.0000155***
	(0.000105)	(0.00000172)
Age	-0.00121	-0.000728***
	(0.000799)	(0.0000406)
Age^2	-0.00000257	0.00000631***
	(0.00000905)	(0.000000422)
Female	0.00892*	-0.000185
	(0.00478)	(0.000184)
High School	-0.0109	-0.00442***
	(0.00743)	(0.000540)
Some College / Associate	0.0153*	-0.00661***
	(0.00814)	(0.000533)
College Graduate	0.0103	-0.00842***
	(0.00859)	(0.000521)
Observations	42729	1197389
R-squared	0.00574	0.00294

Standard errors in parentheses

Constant and coefficients on MIS are not reported.

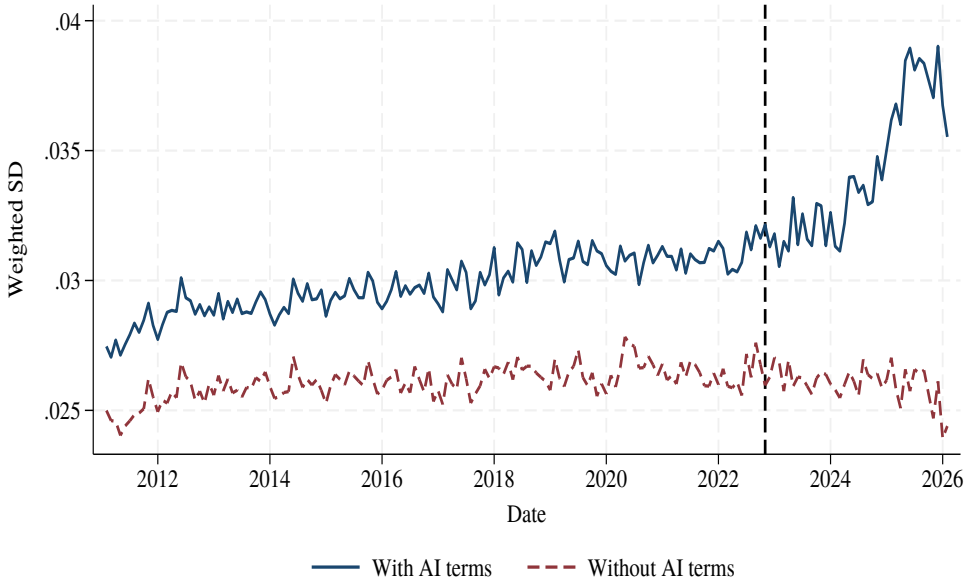
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 5: HIRABILITY INDEX BY THE QUARTILE OF AI EXPOSURE AND AI ADOPTION



Note: Panel A displays the mean hirability index by the quartile of AI exposure. Panel B displays the mean hirability index by the quartile of AI adoption.
Source: Authors' calculation.

Figure 6: DISPERSION OF HIRABILITY INDEX WITH AND WITHOUT AI FACTORS



Note: This figure plots the dispersion of hirability index with AI exposure and adoption (solid blue line) and that without AI exposure and adoption (dashed red line).
Source: Authors' calculation.

4.3 Estimated indices and their dispersion

In Figure 5, Panel A plots the mean hirability index by quartile of AI exposure. The mean hirability index is lower for individuals with higher AI exposure. Moreover, the top two exposure quartiles exhibit a more pronounced decline after the introduction of LLMs (vertical dashed line), with a particularly sharp drop in 2024, while the lower quartiles remain relatively flat.

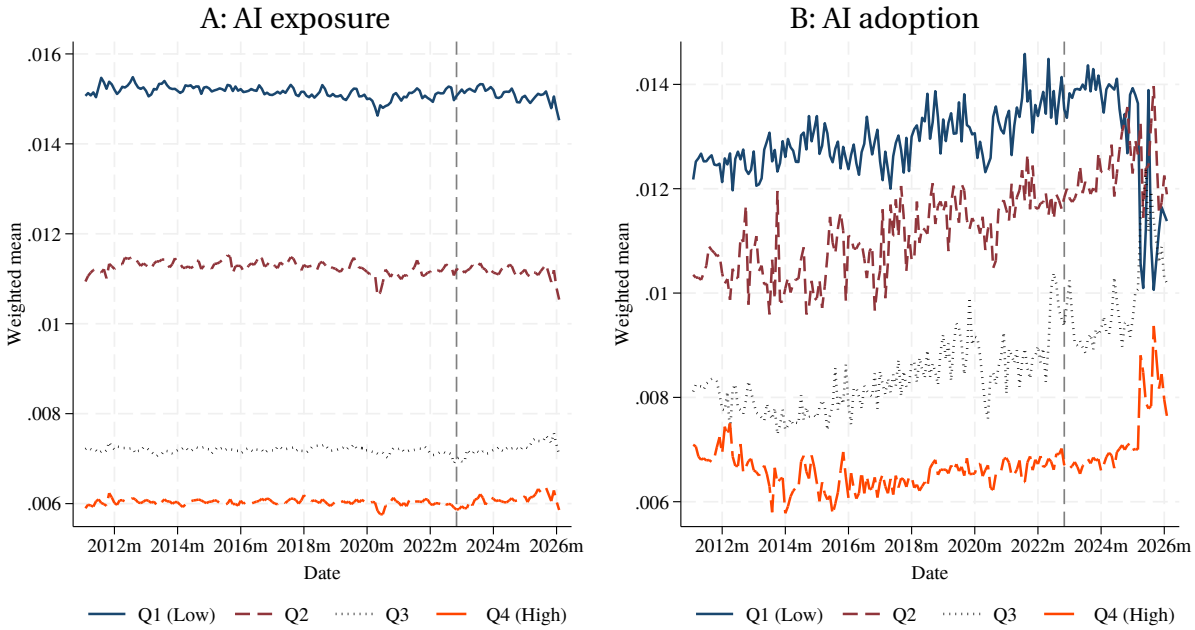
Panel B of Figure 5 plots the mean hirability index by quartile of AI adoption. The mean hirability index is lower for individuals in higher AI adoption quartiles. Notably, while the top AI adoption quartile exhibits a decline in hirability, the third quartile has risen steadily since 2016. This divergence becomes more pronounced after 2023 (vertical dashed line): hirability in the top adoption group declines further, whereas the third quartile continues to increase, reaching levels comparable to the bottom quartile. These patterns suggest that workers employed in firms with high AI adoption have experienced a deterioration in reemployment prospects.

Figure 6 displays the standard deviation of the hirability index ($sd_t^w(UE)$). The dispersion remains largely flat prior to 2023, then rises steadily and accelerates from 2024 onward, indicating a marked increase in reallocation pressure following the introduction of large language models. In contrast, the dispersion of the AI-purged hirability index ($sd_t^{wo}(UE)$) remains essentially unchanged over the entire period. This pattern suggests that AI is the primary driver of the increase in hirability dispersion and the associated rise in reallocation pressure.

Figure 7 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). However, beginning in 2025, the separability index rises for workers in high–AI-adoption occupations, while it declines for those in low–AI-adoption occupations. This observation suggests that high-adoption firms may be reallocating tasks, downsizing certain roles, or restructuring workflows, increasing separations.

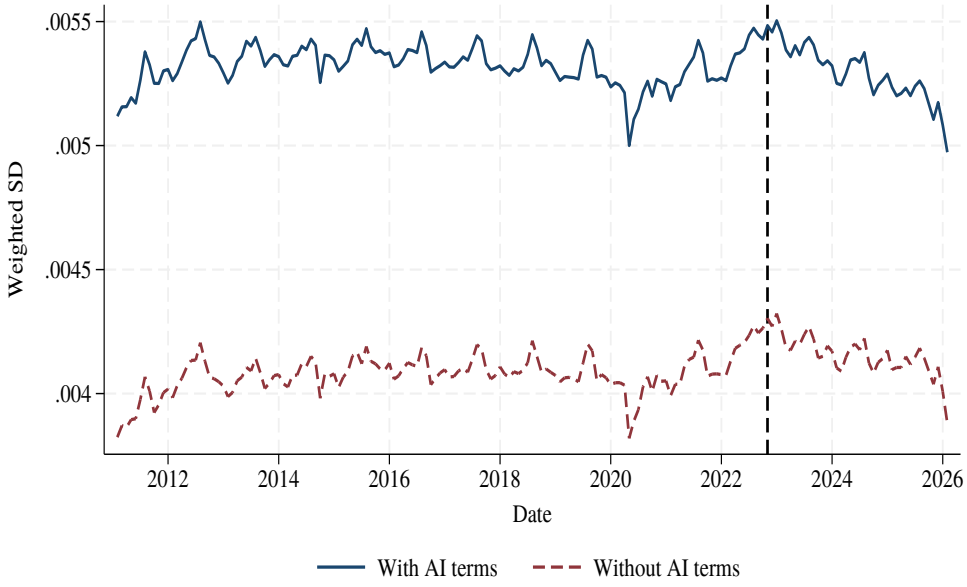
Figure 8 plots the dispersion of the separation index. Relative to the dispersion of the hirability index, the dispersion of the separation index shows no noticeable change after 2024:Q4 and remains broadly similar to that of the separation index constructed without AI-related factors. These results suggest that AI-driven reallocation pressures operate primarily through the hiring margin rather than through separations, at least through late 2025, consistent with the findings

Figure 7: SEPARABILITY INDEX BY THE QUARTILE OF AI EXPOSURE AND AI ADOPTION



Note: Panel A displays the mean separability index by the quartile of AI exposure. Panel B displays the mean separability index by the quartile of AI adoption.
Source: Authors' calculation.

Figure 8: DISPERSION OF SEPARABILITY INDEX WITH AND WITHOUT AI FACTORS



Note: This figure plots the dispersion of hirability index with AI exposure and adoption (solid blue line) and that without AI exposure and adoption (dashed red line).
Source: Authors' calculation.

of Hosseini Maasoum and Lichtinger [2026].

5 Implications for the natural rate of unemployment

The overall empirical evidence points to upward pressure on the natural rate of unemployment in the post-LLM period. To assess the effects of LLMs on the NRU, we estimate group-specific NRUs for high- and low-AI-exposure workers using the trend components of their unemployment inflow and outflow rates, and then aggregate these estimates using their respective labor force shares. Specifically, we extract trend components of unemployment flows rates using a Bayesian trend–cycle decomposition and feed them into Shimer [2012]'s steady-state unemployment framework, which is adopted by Tasci [2012], Crump et al. [2019], Ahn [2023]. We interpret this empirical exercise as providing directional evidence on the risks to natural rate estimates that abstract from AI-driven reallocation effects.

Trend–cycle decomposition by AI exposure and adoption group. Let $g \in \{1,2,3,4\}$ index the four worker groups: low exposure/low adoption, low exposure/high adoption, high exposure/low adoption, and high exposure/high adoption. For each group g, let EU_{gt} denote the unemployment inflow rate and UE_{gt} the unemployment outflow rate. We decompose each series into trend, cyclical, and measurement-error components:

$$EU_{gt} = \tau_{gt}^{EU} + c_{gt}^{EU} + r_{gt}^{EU}, \tag{5.1}$$

$$UE_{gt} = \tau_{gt}^{UE} + c_{gt}^{UE} + r_{gt}^{UE}, \tag{5.2}$$

where r_{gt}^{EU} and r_{gt}^{UE} are mean-zero Gaussian measurement errors.

In the empirical implementation, the trend in the unemployment inflow rate is modeled as a local level, whereas the trend in the unemployment outflow rate is modeled as a local linear trend:

$$\tau_{gt}^{EU} = \tau_{g,t-1}^{EU} + \nu_{gt}^{EU}, \tag{5.3}$$

$$\tau_{gt}^{UE} = \tau_{g,t-1}^{UE} + \beta_{g,t-1}^{UE} + \nu_{gt}^{UE}, \tag{5.4}$$

$$\beta_{gt}^{UE} = \beta_{g,t-1}^{UE} + \kappa_{gt}^{UE}, \tag{5.5}$$

where β_{gt}^{UE} denotes the slope of the trend in UE_{gt}. The cyclical components are modeled as stationary autoregressive processes:

$$c_{gt}^{EU} = \phi_g^{EU} c_{g,t-1}^{EU} + \omega_{gt}^{EU}, \tag{5.6}$$

$$c_{gt}^{UE} = \phi_g^{UE} c_{g,t-1}^{UE} + \omega_{gt}^{UE}. \tag{5.7}$$

All innovations are assumed to be Gaussian and mutually independent.

The model is estimated separately for each group using Bayesian methods. To avoid having unprecedented pandemic movements distort the estimated latent components, observations from March 2020 through July 2020 are excluded from the estimation sample.

Group-specific natural rate. Given the posterior draws of the trend components, the group-specific natural rate of unemployment is defined draw by draw as

$$u_{gt}^{*} = 100 \frac{\tau_{gt}^{EU}}{\tau_{gt}^{EU} + \tau_{gt}^{UE}}. \tag{5.8}$$

We then report the posterior mean of u_{gt}^{*} for each group and date.

Aggregate natural rate. Let ω_{gt} denote the labor-force share of group g at time t, where

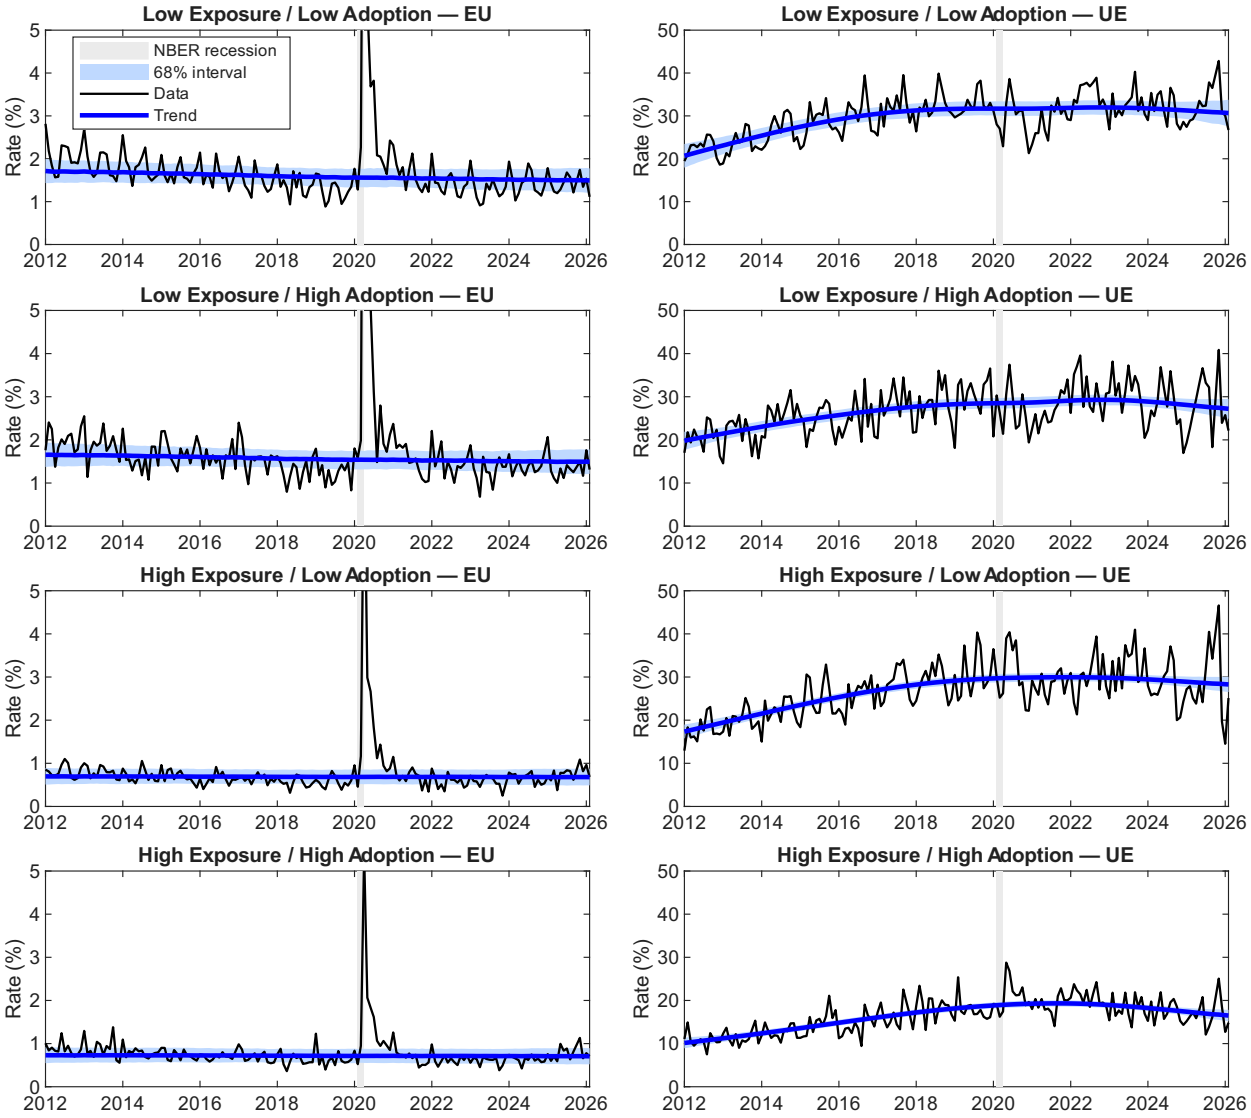
$$\sum_{g=1}^{4} \omega_{gt} = 1. \tag{5.9}$$

The aggregate natural rate is constructed as the labor-force-share-weighted average of the group-specific natural rates:

$$u_t^{*} = \sum_{g=1}^{4} \omega_{gt} u_{gt}^{*}. \tag{5.10}$$

In the estimation, this aggregation is carried out draw by draw using time-varying labor-force shares, so that posterior uncertainty in the group-specific natural rates is propagated into the posterior distribution of the aggregate natural rate. Labor force shares remain largely stable, averaging 42 percent for low-exposure, low-adoption; 9 percent for low-exposure, high-adoption; 22 percent for high-exposure, low-adoption; and 28 percent for high-exposure, high-adoption.

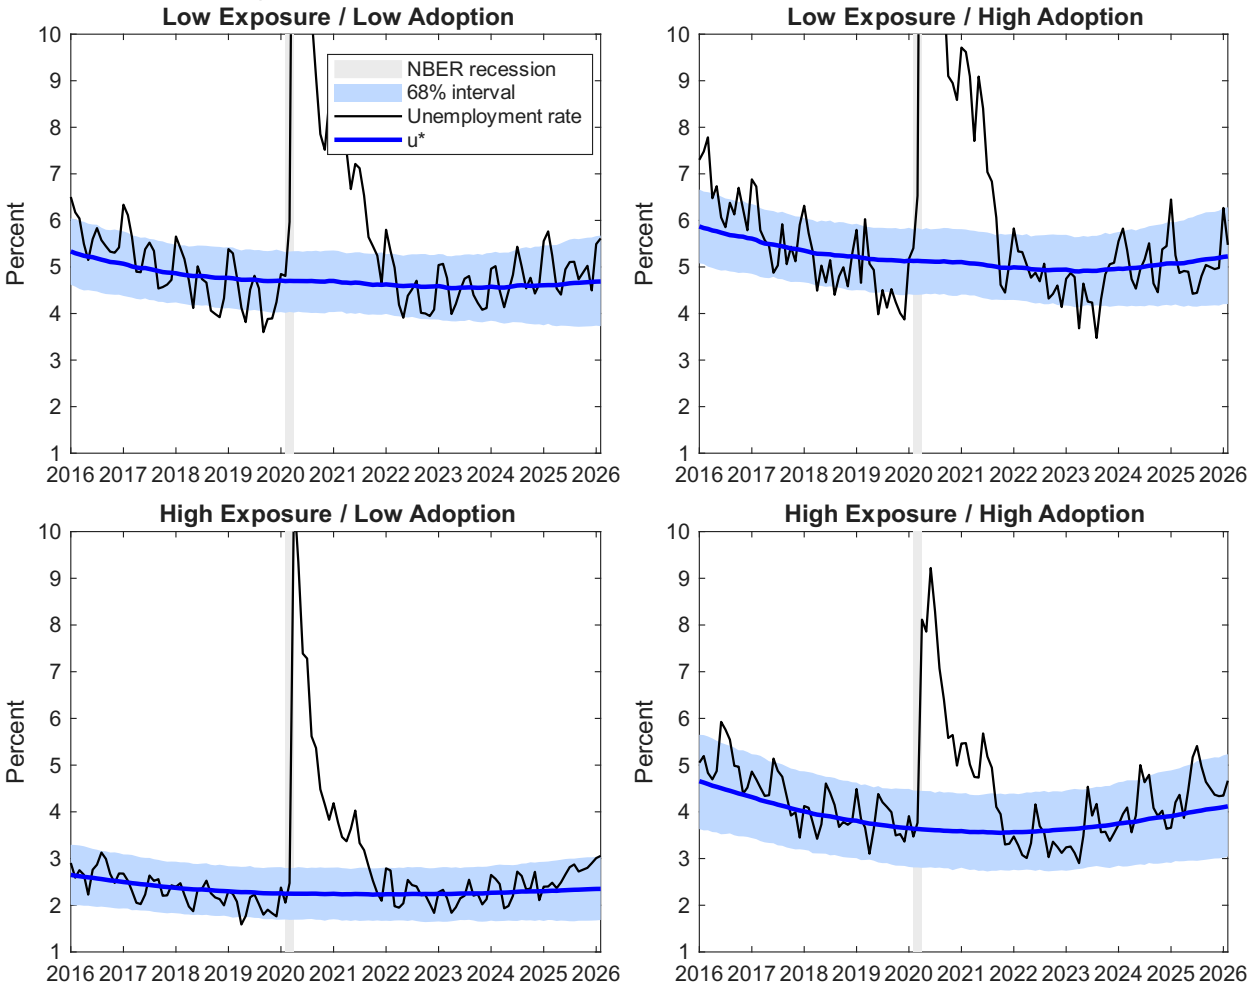
Figure 9: EU AND UE TRENDS BY AI EXPOSURE AND ADOPTION



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.

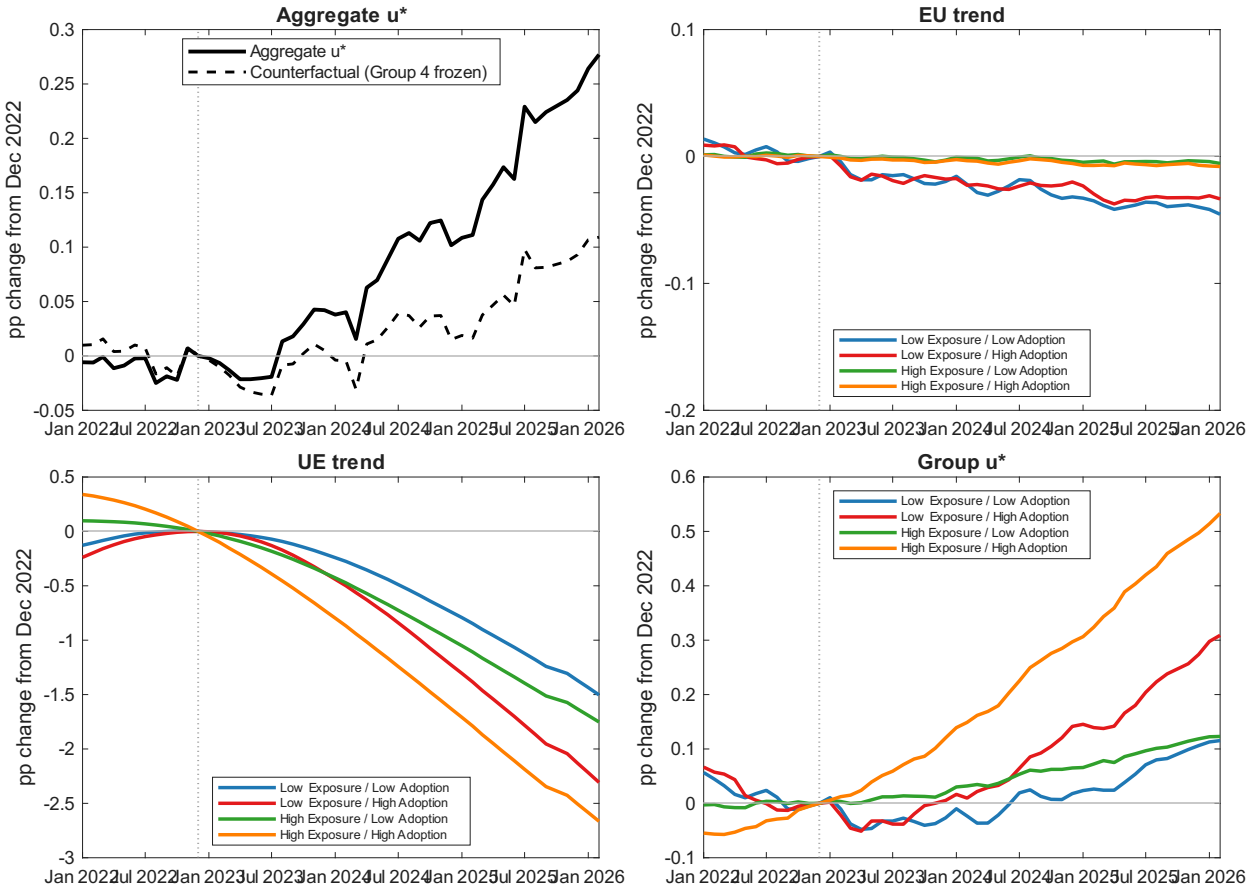
Figure 10: U-STAR ESTIMATE BY AI EXPOSURE AND ADOPTION



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.

Figure 11: TREND UNEMPLOYMENT FLOWS AND U-STAR SINCE LLMS



Note to figure: The upper-left panel displays the aggregate natural rate and a counterfactual series in which the natural rate of the high-exposure, high-adoption group is held fixed. The upper-right panel shows the EU trends for the four groups. The bottom-left panel presents the UE trend rates for the four groups, and the bottom-right panel displays the estimated natural rates for each group.

Note to figure: Authors' calculation.

Figure A.5 displays the unemployment inflow (EU) and outflow (UE) rates and their estimated trends by AI exposure and adoption. We do not seasonally adjust the EU and UE flow rates, as standard seasonal adjustment procedures applied to series with large COVID-era swings may introduce additional noise and distort the underlying dynamics.[9]

EU trends remain largely flat over the sample period. In contrast, UE trends increase across all four groups through the end of 2019, plateau until around 2022, and then begin to decline. Notably, the post-2022 decline in UE trends varies across groups, with the largest decrease observed for the high-exposure, high-adoption group relative to the others. One might argue that these differences reflect varying cyclical sensitivities in job-finding rates across the four groups. However, workers in high–AI-exposure and high–AI-adoption jobs typically have more stable employment, and thus their job-finding rates should be least responsive to adverse cyclical shocks. Instead, we observe the opposite pattern, suggesting that the larger decline in job-finding probabilities for these workers is unlikely to be driven by cyclical or aggregate uncertainty shocks.

Figure A.6 displays each group's estimated natural rate along with its 68% posterior interval. Given the relatively short sample period—about 15 years—and the large fluctuations during the COVID episode, the trend components of EU and UE are estimated with considerable uncertainty, resulting in wide posterior intervals for the natural rate. The high-exposure, high-adoption group exhibits a somewhat larger increase than the other groups.

Figure A.7 zooms in on the period following the introduction of LLMs, normalizing all estimates to zero in December 2022. The upper-left panel displays the aggregate natural rate (solid line) alongside a counterfactual in which the natural rate of high-exposure, high-adoption workers is held constant (dashed line). Between December 2022 and February 2026, the aggregate natural rate is estimated to have increased by about 0.25 percentage points. Most of this increase is attributable to a rise in the natural rate of high-exposure, high-adoption workers—captured by the gap between the solid and dashed lines and shown by the orange line in the bottom-right panel—contributing roughly 15 basis points. This increase is primarily driven by a faster decline in the trend UE rate (orange line in the bottom-left panel). The natural rate of workers in the low-exposure, high-adoption group also rises (red line in the bottom-right panel), reflecting a

[9]We consider the 12-month moving average flows rates instead of the seasonally unadjusted data, and found that the results are overall robust. See appendix for more details.

faster decline in their trend UE rate relative to low-exposure workers (red line in the bottom-left panel). However, because their labor force share is below 10 percent, their contribution to the aggregate increase is only three basis points. The remaining groups each contribute about three basis points to the aggregate natural rate.

In addition, the differing contributions to the aggregate natural rate primarily reflect variation in the pace of decline in trend UE rates rather than in trend EU rates. This pattern is consistent with previous evidence emphasizing the role of the hiring margin, rather than the separation margin, in LLM-driven reallocation. Nonetheless, the estimates are subject to considerable uncertainty, and we therefore interpret our results as indicating potential upside risks to the natural rate of unemployment associated with LLMs.

6 Conclusion

In this paper, we provide a comprehensive assessment of how the rapid diffusion of AI—particularly large language models—has reshaped aggregate labor market dynamics in the United States. By integrating traditional labor market data with novel measures of AI exposure and adoption, and by tracing their effects through stocks and flows in the labor market, we uncover a set of evidence suggesting rising reallocation pressure and structural unemployment.

Workers in highly exposed and adopting jobs continue to exhibit lower unemployment rates, largely due to lower separation rates. However, since the introduction of LLMs, these workers have experienced a disproportionate decline in job-finding rates and job-to-job mobility, alongside persistently weaker market tightness, indicating that adjustment has occurred primarily along the hiring margin. At the same time, we document clear evidence that LLMs are reshaping the allocation of labor within firms. The rise in within-job activity switching among highly exposed workers suggests that firms are reorganizing tasks rather than simply replacing workers, consistent with a reconfiguration of production processes along with reduced demand for workers with high AI exposure and high AI-adoption jobs.

Taken together, these patterns indicate a buildup of reallocation pressure in the post-LLM period. The sharp rise in the dispersion of job-finding propensities—captured by our hirability index and largely driven by AI-related factors—suggests that workers are increasingly differen-

tiated in their ability to transition into employment. The absence of a comparable increase in separation risk further underscores that the primary friction lies in matching workers to jobs, rather than in job destruction.

Taken together, these patterns suggest a buildup of reallocation pressure in the post-LLM period. The increase in the dispersion of job-finding propensities—driven largely by AI-related factors—indicates rising frictions in matching workers to jobs. These dynamics have important macroeconomic implications. By embedding AI exposure and adoption into a flow-based framework for unemployment, we show that the uneven decline in job-finding rates among highly exposed workers has likely exerted modest upward pressure on the natural rate of unemployment since late 2022. While the estimated increase is relatively small and subject to uncertainty, it is consistent with a labor market in which structural frictions—rather than cyclical weakness—play a growing role.

More broadly, our results suggest that the impact of AI on the labor market is best understood not simply in terms of substitution or complementarity, but as a process of gradual and uneven reallocation. The extent to which the labor market can efficiently reallocate workers in response to AI-driven changes in labor demand will be critical in determining the longer-run effects of LLMs on employment, productivity, and economic welfare. Future work should continue to track these dynamics as AI adoption deepens and as institutions, firms, and workers adapt to the evolving technological landscape.

References

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. The simple macroeconomics of ai. *Economic Policy*, 40(121):13–58, 2025.

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.

Hie Joo Ahn, Bart Hobijn, and Aysegul Sahin. The Dual U.S. Labor Market Uncovered. Finance and Economics Discussion Series 2023-031, Board of Governors of the Federal Reserve System (U.S.), May 2023. URL https://ideas.repec.org/p/fip/fedgfe/2023-31.html.

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, 2025a.

Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. Generative ai at work*. *The Quarterly Journal of Economics*, 140(2):889–942, 02 2025b. ISSN 0033-5533. doi: 10.1093/qje/qjae044. URL https://doi.org/10.1093/qje/qjae044.

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.

Z. Cui, M. Demirer, S. Jaffe, L. Musolff, S. Peng, and T. Salz. The effects of generative ai on high skilled work: Evidence from three field experiments with software developers. Available at SSRN 4945566, 2024.

Fabrizio Dell'Acqua, Edward McFowland III, Ethan Mollick, Katherine Kellogg Lifshitz-Assaf, Saran Rajendran, Lisa Krayer, Francois Candelon, and Karim Lakhani. Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker

productivity and quality. Working Paper No. 24-013, Harvard Business School, Cambridge, MA., 2023.

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.

Francesco Filippucci, Peter Gal, and Matthias Schief. Aggregate productivity gains from artificial intelligence: A sectoral perspective. Working Paper, 2025.

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.

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.

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.

Shakked Noy and Whitney Zhang. Experimental evidence on the productivity effects of generative artificial intelligence. Available at SSRN 4375283, 2023.

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.

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

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

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 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 and adoption are significantly affected by our imputation for missing occupation.

Figure A.1 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. AI exposure and adoption are associated with a lower likelihood of exiting the labor force overall.

Figure A.2 displays the NE and NU rates by AI exposure and AI adoption. Although workers with higher AI exposure and adoption tend to exhibit higher NE and NU rates, the relationship between AI factors and inflows into the labor force is nonlinear. In particular, both NE and NU rates are higher in the lowest AI quartiles than in the middle quartiles, while workers in the highest quartiles tend to exhibit the highest NE and NU rates overall. One exception is the NU rate by AI adoption, where the top quartile exhibits a lower NU rate than the bottom quartile.

This pattern likely reflects heterogeneity across worker types: low-skilled workers, who tend to be concentrated in low–AI-adoption jobs, are more likely to cycle between the labor force and nonparticipation, while workers in high–AI-adoption jobs exhibit stronger labor-force attachment, as the returns to job search are higher than remaining out of the labor force. However, it is notable that both NE and NU rates for workers in the highest AI-adoption quartile has declined relative to other groups, potentially suggesting that AI adoption may have discouraged job search or weakened labor force attachment among nonparticipants.

Figure A.3 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

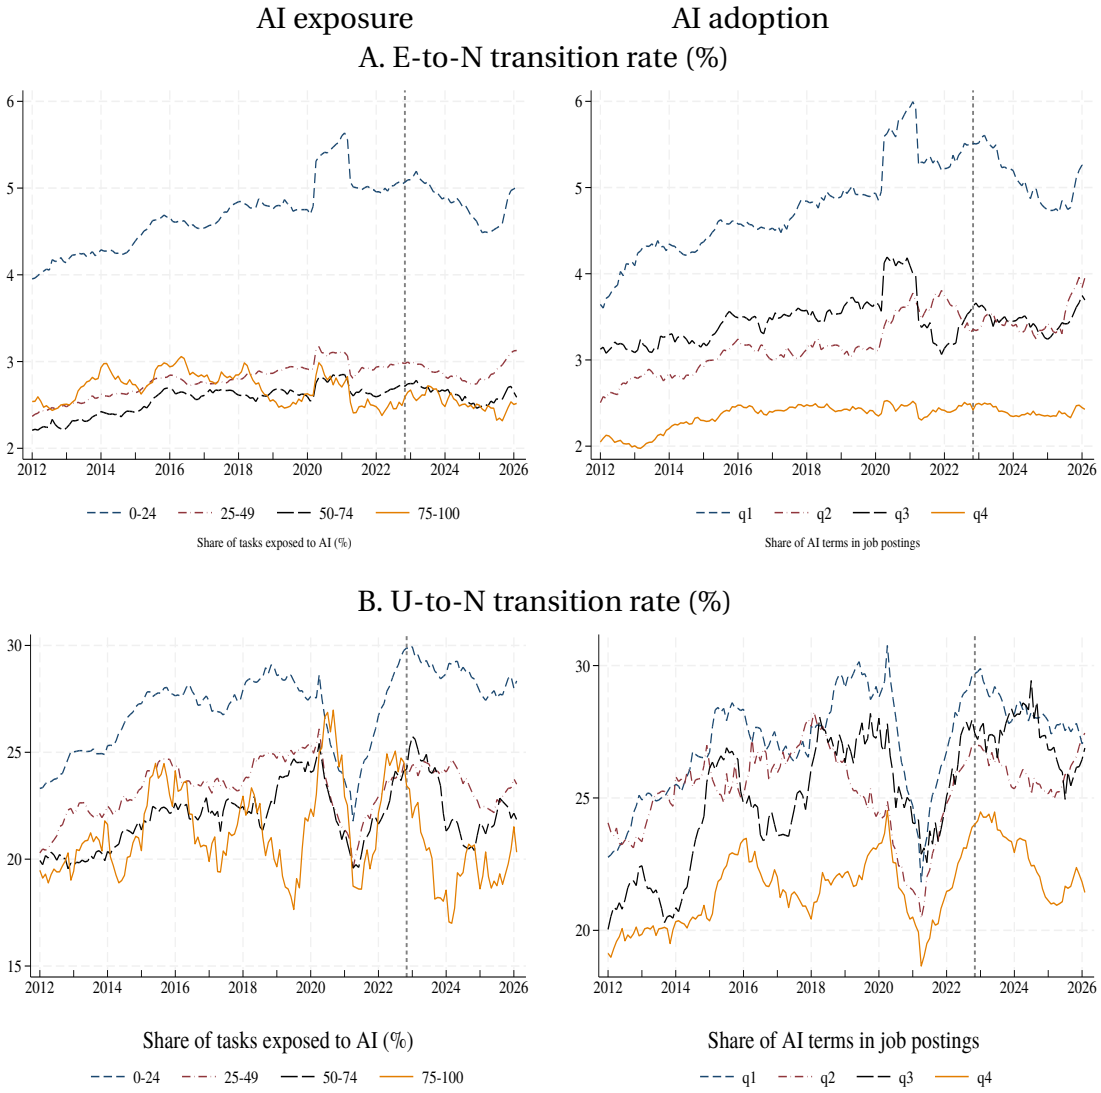
workers in the top quartiles of AI exposure and adoption. Following the COVID pandemic, LFPR and EPOP stepped down among workers in the top AI-adoption quartile, consistent with the reduced NU and NE rates for this group.

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 worker reallocation rate of the flows per capita.

Figure A.4 displays the worker reallocation rate by AI exposure. In general, there is a nonlinear relationship between worker reallocation and AI exposure and adoption. Workers with the lowest AI exposure and adoption exhibit greater reallocation, reflecting higher churn associated with job instability. Meanwhile, groups with the highest AI exposure and adoption also show greater reallocation than the middle groups—specifically, those with 25–49% AI exposure and the third quartile of AI adoption. The higher reallocation among the top groups relative to the middle groups suggests stronger job and labor force attachment, as these workers are more likely to find jobs or return to the labor force from nonparticipation.

It is notable that worker reallocation in industries with low AI adoption picked up and remained elevated relative to the pre-COVID period. As discussed earlier in the context of higher LFPR and EPOP ratios among workers in firms with low AI adoption, these firms may have continued to hire—potentially drawing from workers with higher AI exposure.

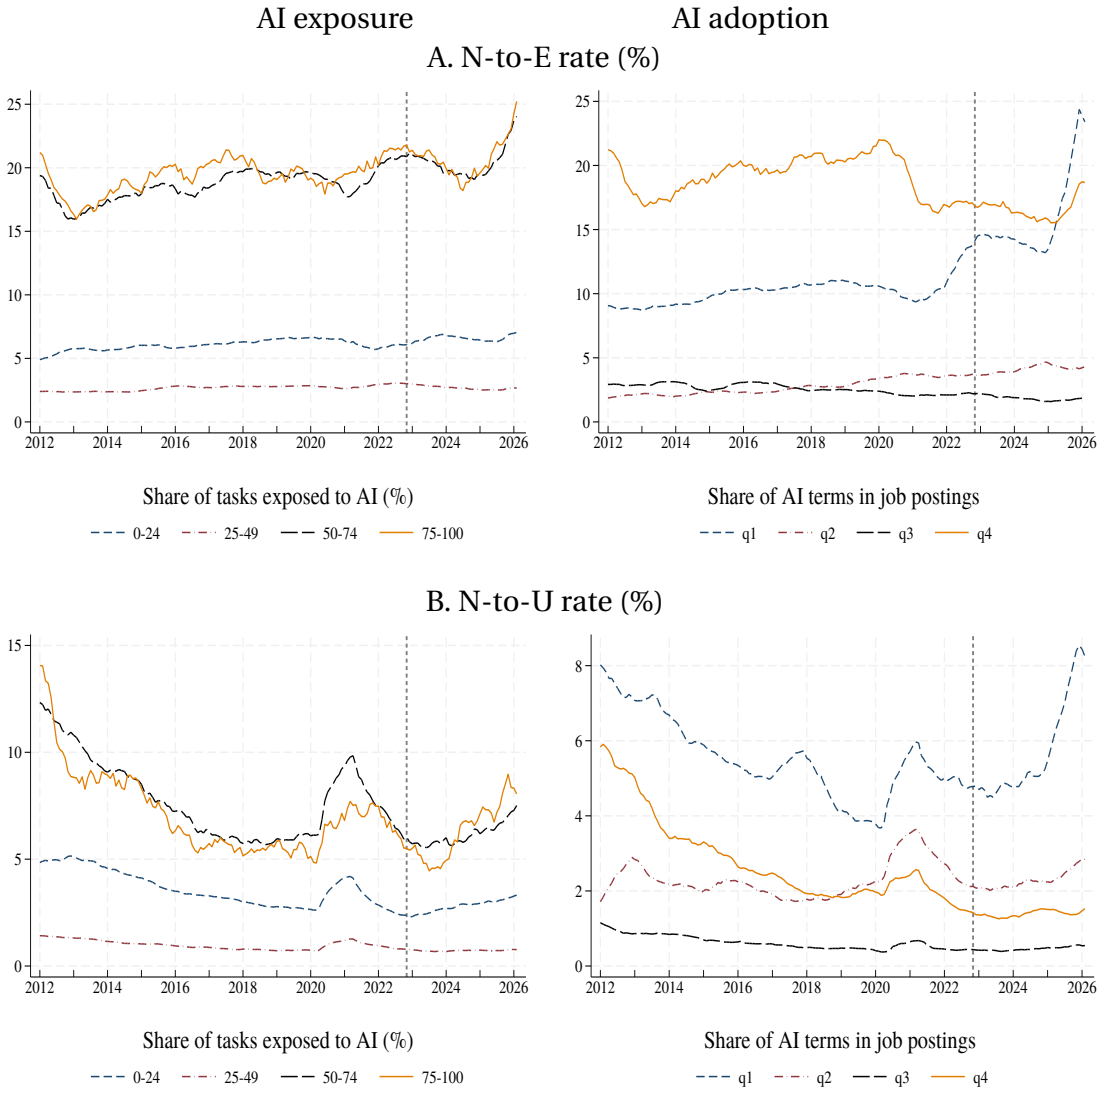
Figure A.1: EN AND UN FLOWS



Note to figure: Panels A and B display the EN and UN rates, respectively. The left panels present the breakdown by AI exposure, while the right panels present the breakdown by AI adoption.

Source: Authors' calculation

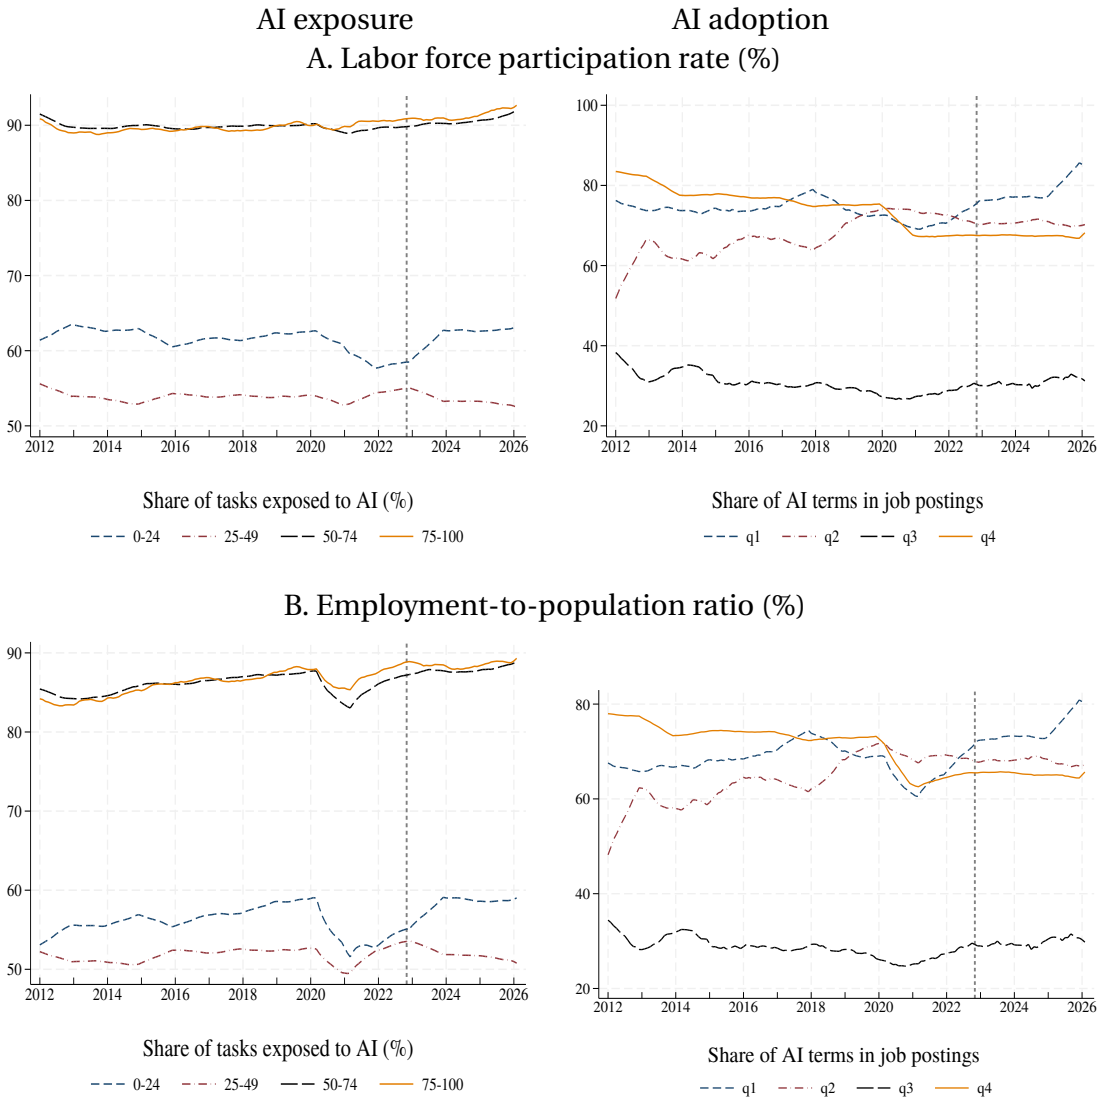
Figure A.2: NE AND NU FLOWS RATES



Note to figure: Panels A and B display the NE and NU rates, respectively. The left panels present the breakdown by AI exposure, while the right panels present the breakdown by AI adoption.

Source: Authors' calculation

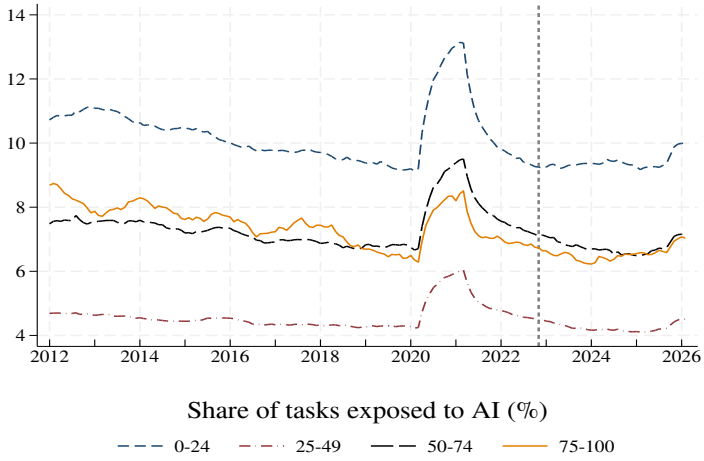
Figure A.3: LFPR AND EPOP RATIO WITH/WITHOUT AI EXPOSURE



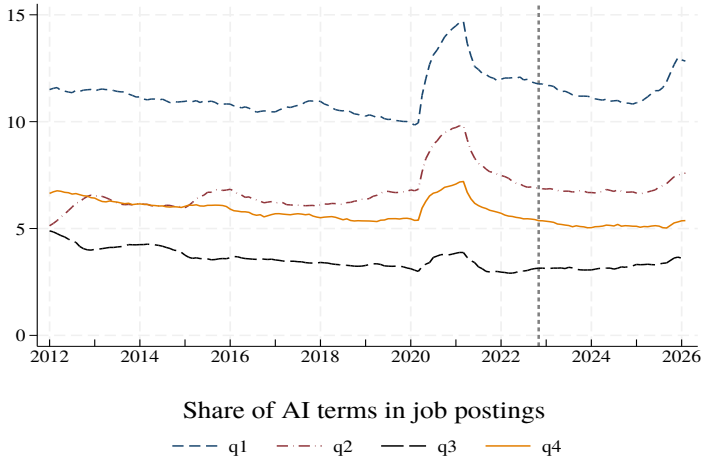
Note to figure: Panels A and B show LFPR and the EPOP ratio by AI exposure (left) and AI adoption (right).
Source: Authors' calculation

Figure A.4: WORKER REALLOCATION RATE: FLOWS PER CAPITA (12-MONTH MOVING AVERAGE)

Panel A. AI exposure



Panel B. AI adoption



Note to figure: This figure displays the worker reallocation rate by AI exposure. The solid blue line shows the average reallocation rate; the red dashed line represents the reallocation rate for workers with high AI exposure; and the dotted black line represents the reallocation rate for workers with low AI exposure.
Source: Authors' calculation

B Bayesian Trend–Cycle Decomposition

B.1 Data and setup

We partition workers into four groups $g \in \{1,2,3,4\}$, defined by combinations of AI exposure (low/high) and AI adoption (low/high). For each group, we construct monthly unemployment inflow and outflow rates, denoted EU_{gt} and UE_{gt}.

To mitigate the influence of extreme pandemic movements, observations from March 2020 through July 2020 are excluded from estimation.

B.2 Model

For each group, we decompose the observed series into trend, cycle, and measurement error:

$$EU_{gt} = \tau_{gt}^{EU} + c_{gt}^{EU} + \varepsilon_{gt}^{EU}, \tag{B.1}$$

$$UE_{gt} = \tau_{gt}^{UE} + c_{gt}^{UE} + \varepsilon_{gt}^{UE}, \tag{B.2}$$

where measurement errors are Gaussian.

Trend. The inflow rate follows a local level, and the outflow rate follows a local linear trend:

$$\tau_{gt}^{EU} = \tau_{g,t-1}^{EU} + \nu_{gt}^{EU}, \tag{B.3}$$

$$\tau_{gt}^{UE} = \tau_{g,t-1}^{UE} + \beta_{g,t-1}^{UE} + \nu_{gt}^{UE}, \tag{B.4}$$

$$\beta_{gt}^{UE} = \beta_{g,t-1}^{UE} + \kappa_{gt}^{UE}. \tag{B.5}$$

Cycle. Cyclical components follow AR(1) processes:

$$c_{gt}^{EU} = \phi_g c_{g,t-1}^{EU} + \omega_{gt}^{EU}, \tag{B.6}$$

$$c_{gt}^{UE} = \phi_g c_{g,t-1}^{UE} + \omega_{gt}^{UE}, \tag{B.7}$$

with $\phi_g = 0.90$.

All innovations are Gaussian and mutually independent.

State-space form. Let $y_{gt} = (EU_{gt}, UE_{gt})'$ and define the state vector

$$x_{gt} = \begin{bmatrix} \tau_{gt}^{EU} & \tau_{gt}^{UE} & \beta_{gt}^{UE} & c_{gt}^{EU} & c_{gt}^{UE} \end{bmatrix}'.$$

The model can be written as

$$y_{gt} = H x_{gt} + \varepsilon_{gt}, \tag{B.8}$$

$$x_{gt} = F_g x_{g,t-1} + \eta_{gt}, \tag{B.9}$$

with Gaussian errors.

B.3 Prior specification

Innovation variances follow inverse-gamma priors:

$$\sigma^2_{\nu^{EU},g} \sim \mathcal{IG}(a_\eta, b_{\eta,g}), \tag{B.10}$$

$$\sigma^2_{\omega^{EU},g} \sim \mathcal{IG}(a_\zeta, b_{\zeta,g}), \tag{B.11}$$

$$\sigma^2_{\kappa^{UE},g} \sim \mathcal{IG}(a_\kappa, b_\kappa), \tag{B.12}$$

$$\sigma^2_\varepsilon \sim \mathcal{IG}(a_\varepsilon, b_\varepsilon), \tag{B.13}$$

with

$$a_\eta = 30, \quad a_\zeta = 35, \quad a_\kappa = 2, \quad a_\varepsilon = 2.$$

To align priors with the scale of the data, the inflow-related scale parameters are adjusted across groups:

$$b_{\eta,g} = b_\eta \cdot \frac{\mathrm{Var}(EU_g)}{\mathrm{Var}(EU_1)}, \qquad b_{\zeta,g} = b_\zeta \cdot \frac{\mathrm{Var}(EU_g)}{\mathrm{Var}(EU_1)}.$$

B.4 Estimation

Conditional on parameters, the model is linear and Gaussian. The likelihood is evaluated using the Kalman filter:

$$p(y_{1:T} \mid \Theta) = \prod_{t=1}^{T} p(y_t \mid y_{1:t-1}, \Theta). \tag{B.14}$$

We estimate the model separately for each group using Bayesian methods. Posterior simulation is conducted via a Gibbs sampler that alternates between: (i) drawing latent states using a simulation smoother and (ii) drawing variance parameters from their conditional posterior distributions.

We generate 6000 draws, discard the first 1500 as burn-in, and retain every fifth draw for inference.

B.5 Natural rate of unemployment

For each posterior draw m, the group-specific natural rate is defined as

$$u_{gt}^{*,(m)} = 100 \cdot \frac{\tau_{gt}^{EU,(m)}}{\tau_{gt}^{EU,(m)} + \tau_{gt}^{UE,(m)}}. \tag{B.15}$$

Posterior means and credible intervals are computed from the retained draws.

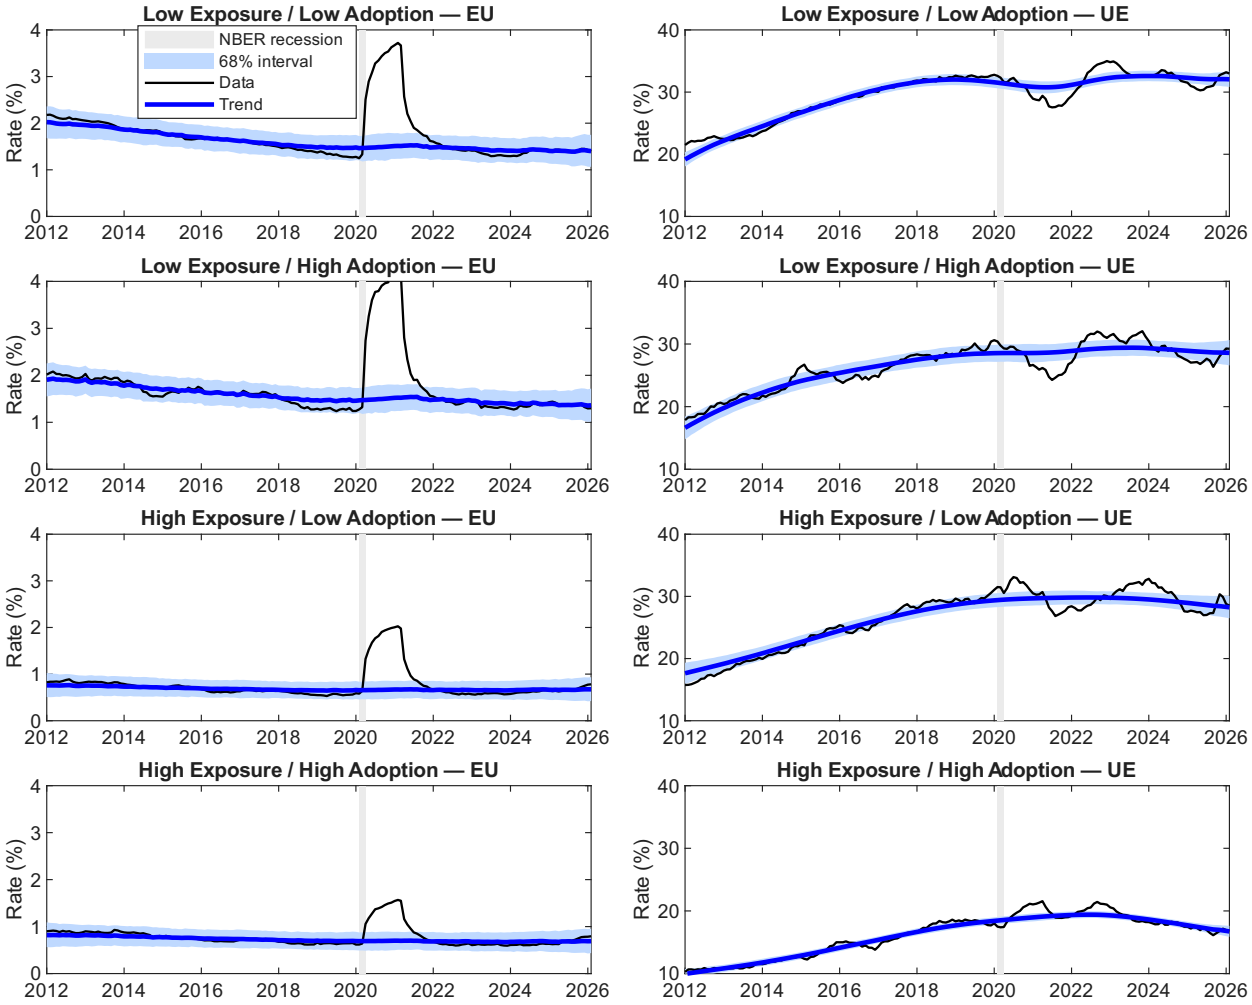
B.6 Aggregate natural rate

Let ω_{gt} denote the time-varying labor-force share of group g, with $\sum_{g=1}^{4} \omega_{gt} = 1$. The aggregate natural rate is computed draw by draw:

$$u_t^{*,(m)} = \sum_{g=1}^{4} \omega_{gt}\, u_{gt}^{*,(m)}. \tag{B.16}$$

This procedure propagates uncertainty in group-specific estimates into the aggregate natural rate.

Figure A.5: EU AND UE TRENDS BY AI EXPOSURE AND ADOPTION

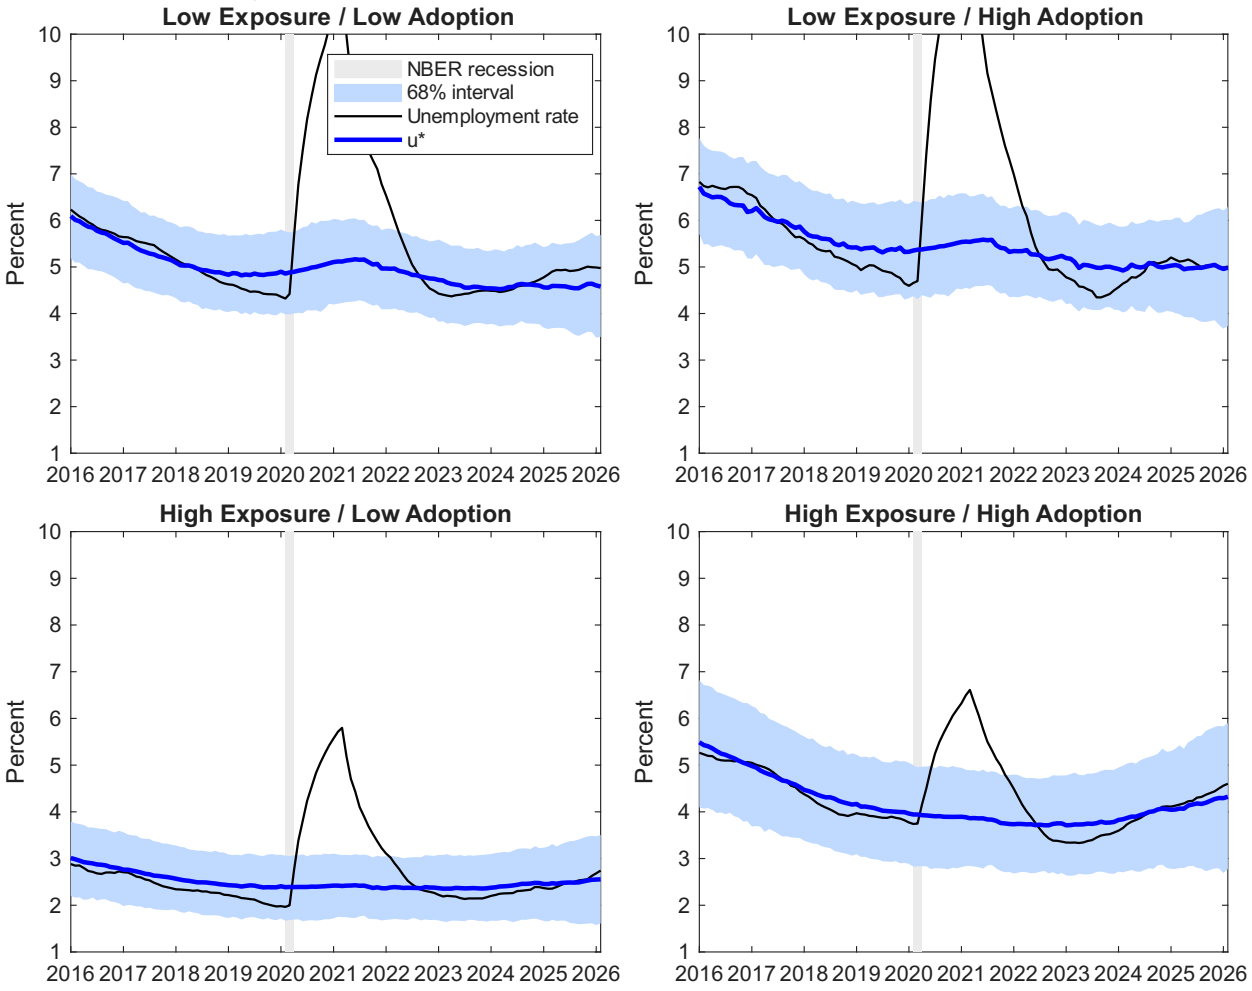


Note to figure: This figure displays unemployment inflow (EU) and outflow (UE) rates by AI exposure and adoption. The black lines represent the observed data, the blue lines denote the posterior median estimates, and the shaded blue areas indicate the 68 percent posterior intervals.

Note to figure: Authors' calculation.

Remark. Because the natural rate is a nonlinear function of latent trend components, posterior simulation provides a convenient way to obtain its full distribution without relying on linear approximations.

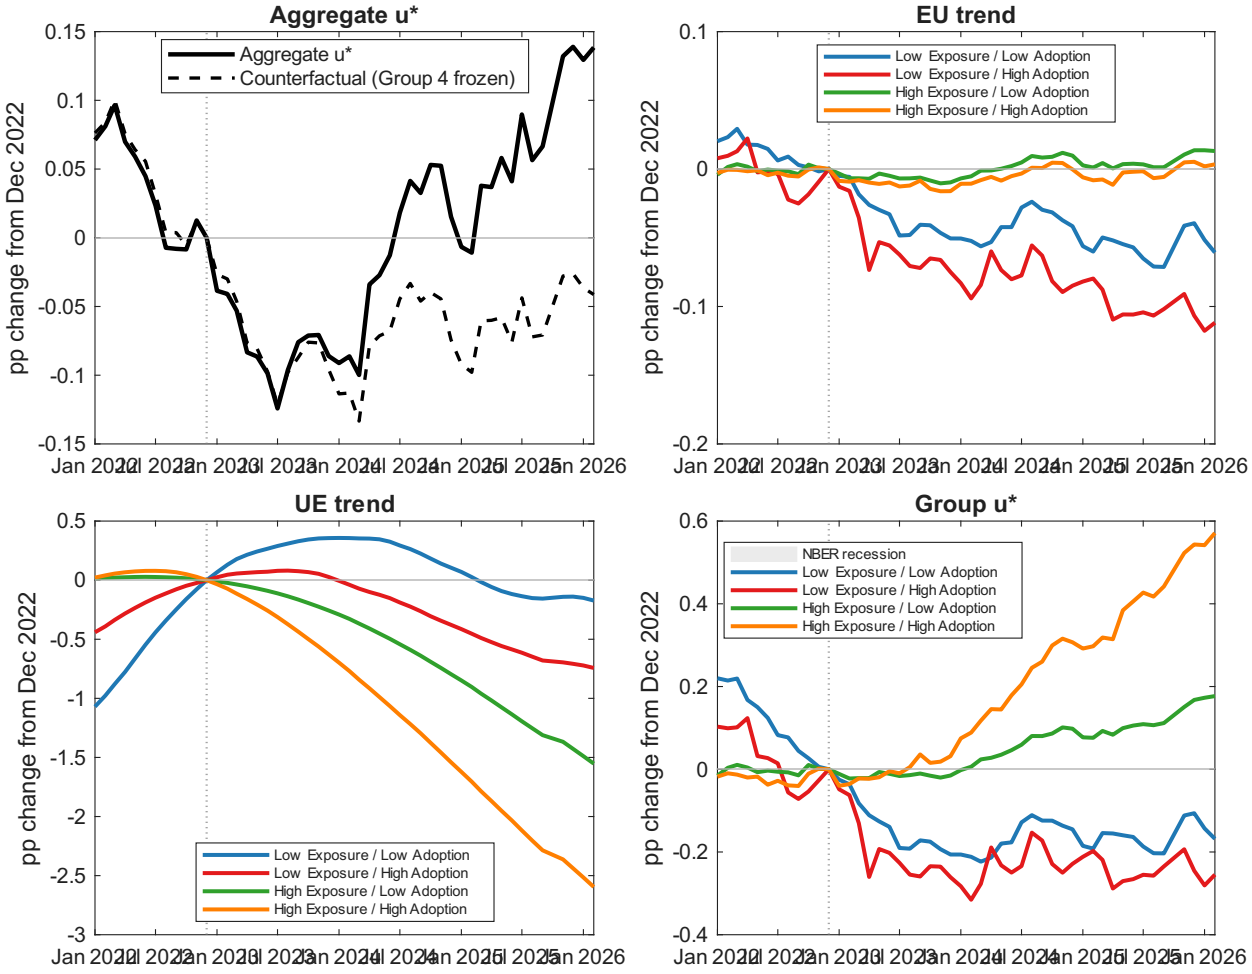
Figure A.6: U-STAR ESTIMATE BY AI EXPOSURE AND ADOPTION



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.

Figure A.7: TREND UNEMPLOYMENT FLOWS AND U-STAR SINCE LLMS



Note to figure: The upper-left panel displays the aggregate natural rate and a counterfactual series in which the natural rate of the high-exposure, high-adoption group is held fixed. The upper-right panel shows the EU trends for the four groups. The bottom-left panel presents the UE trend rates for the four groups, and the bottom-right panel displays the estimated natural rates for each group.

Note to figure: Authors' calculation.