Artificial Intelligence and Labor Market Reallocation[*]

Hie Joo Ahn[†] Nicholas A. Carollo[‡]

September 16, 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 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 structural unemployment—is estimated to have risen by about 0.1-0.2 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 a new disease of which some readers may not yet have heard the name, but of which they will hear a great deal in the years to come—namely, technological unemployment.

— John Maynard 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 broader 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. Specifically, we draw on occupation-level exposure scores from Eloundou et al. (2023), 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 job 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. AI adoption is measured at the detailed industry level in job postings; unlike our time-invariant exposure measures, it varies over time. This adoption measure is strongly correlated with the Census Business Trends and Outlook Survey (BTOS, henceforth) at the industry level. 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, 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, on average, 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 (whose outcomes are typically more cyclically sensitive). This differential weakness is driven by the margin of hiring or job finding 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 remained at 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 (Hosseini and Lichtinger, 2025).

Third, LLMs have reshaped labor allocation within firms. Since the introduction of LLMs, the within-job activity-switching rate rose disproportionately for high-exposure workers, pointing to increased task reorganization. We also find that AI's positive effects on nominal wage growth and hours worked attenuated markedly after LLMs were introduced. In the pre-LLM period, both AI exposure and adoption are associated with higher nominal wage growth, and AI exposure with longer hours. These effects weakened 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 signal a potential buildup of reallocation pressure, wherein workers experience difficulties moving to sectors with better reemployment prospects due to mismatch or other frictions. Reallocation pressure is traditionally measured by the cross-sectional dispersion of market tightness and employment growth (e.g., Lilien, 1982; Jackman and Roper, 1987; Şahin et al., 2014). To assess AI-driven reallocation pressure, we construct individual-level *hirability* and *separability* indices measuring an individual's propensity for job finding from unemployment and for job loss, respectively. In the estimation, in order to isolate the effects of AI on labor-market transitions from broader cyclical conditions, we control for industry-by-month fixed effects along with occupation-level lagged labor-market tightness that absorb industry-specific shocks as well as heterogeneous responses to monetary policy shocks; we further address the potential endogeneity of AI adoption using a shift-share-style instrument whose shares are determined by occupation-level AI exposure, capturing industries' pre-existing potential to adopt LLMs. We show that the cross-sectional dispersion in the hirability index increased noticeably from the end of 2022, with AI factors accounting for most of the rise. By contrast, dispersion in 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.[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 propensity reflects increased frictions in reallocating workers to sectors with better

[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).

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 about 0.1 – 0.2 percentage point, depending on parameter calibration.

Moreover, we statistically estimate the natural rate of unemployment for worker groups with different AI exposure and adoption by feeding each group's trend job-loss and job-finding rates into the two-state flows model of unemployment rate (Shimer, 2012).[2] We employ a Bayesian trend–cycle decomposition to extract group-specific trend components of unemployment flows rates. One merit of this approach is that it allows us to control pre-LLM trends and business cycles across groups. Reflecting the persistent weakening of job-finding rates for workers with high AI exposure and adoption relative to those with low AI exposure, the natural rate is estimated to have risen by about 0.1 – 0.2 percentage point since the introduction of LLMs through mid-2026, albeit with considerable uncertainty. This statistical estimate aligns with the theoretical model's prediction.

Our empirical findings have policy implications. The persistently weaker job-finding rates of high-exposure/high-adoption workers since the introduction of LLMs may reflect 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 improved market tightness among them, may 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 AI 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 the natural rate of unemployment. Our focus on aggregate labor market functioning distinguishes this paper from prior work, which has primarily examined AI's effects on specific occupations, demographic groups, or selected outcomes such as employment and productivity (Brynjolfsson et al., 2018;

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

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 well-suited to the measurement of market tightness and reveals differential effects of AI adoption.

Our findings also depart from prior research that documents small or negligible effects of AI on the unemployment rate (e.g., The Budget Lab at Yale, 2025a), in three ways. First, we focus on workers at the right tail of the AI exposure and adoption distribution — the top 5th and 10th percentiles — whose unemployment is less cyclically sensitive than that of other workers. We find that their job-finding rate from unemployment is lower and has declined more than that of their peers, raising their unemployment rate disproportionately. Studies that rely on coarser groupings, such as quartiles, are unlikely to detect effects concentrated at the tail. Second, by considering AI exposure and adoption jointly, we uncover an interaction between the two: workers facing high levels of both experience more negative effects on job-finding than would be implied by either dimension alone, further elevating their unemployment rate relative to workers with comparable levels of exposure or adoption in isolation. Third, rather than focusing on unemployment stocks alone, we analyse labour market flows in a manner consistent with the steady-state unemployment rate, allowing us to explore the implications of AI for aggregate labour market slack — a dimension that prior studies have left largely unexamined.

Filling this gap requires new conceptual and empirical tools. To that end, this paper is the first to develop a unified framework for assessing the effects of LLMs on the natural rate of unemployment, introducing 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 AI's effects 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 cross-sectional dispersion of hirability, allowing us to quantify changes in structural unemployment due to LLMs. Third, we adopt a Bayesian trend–cycle model of unemployment flows, building on the two-state flows model of the unemployment rate, to estimate the natural rate 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 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 AI. 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 Census' BTOS asks firms whether they are using AI to substitute for tasks, supplement existing tasks, or introduce 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 high AI-exposure workers.

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 BTOS 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,

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

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 constructed in this paper extend this idea to individual workers, allowing us to characterize the extent to which AI has affected reallocation pressure in the labor market through the hiring and separation margins.

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 of employment and wages is whether AI substitutes for expert or inexpert tasks within an occupation. They 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 in particular (Crane and Soto, 2026).

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 also reached this conclusion (Eckhardt and Goldschlag, 2025). 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 (Ozkan and Sullivan, 2025).

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

[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.

compete for previously inaccessible jobs. Our findings are broadly consistent: job-finding rates are lower and declined more 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.

In this context, our 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 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.[5] Jackman and Roper (1987) argue that structural unemployment rate is determined by mismatch of job seekers and job vacancies. Building on this framework, Şahin et al. (2014) construct a mismatch index based on the dispersion of market tightness across sectors. 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

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

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 shows that firms are more likely to retrain workers in AI-exposed occupations than to lay them off (Abel et al., 2025). 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 providing the first systematic analysis of AI's influence on unemployment flows and stocks, with particular emphasis on its implications for the natural rate of unemployment—the functioning of the labor market.

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

Figure 1: RELATIONSHIP BETWEEN AI OCCUPATIONAL EXPOSURE INDICES

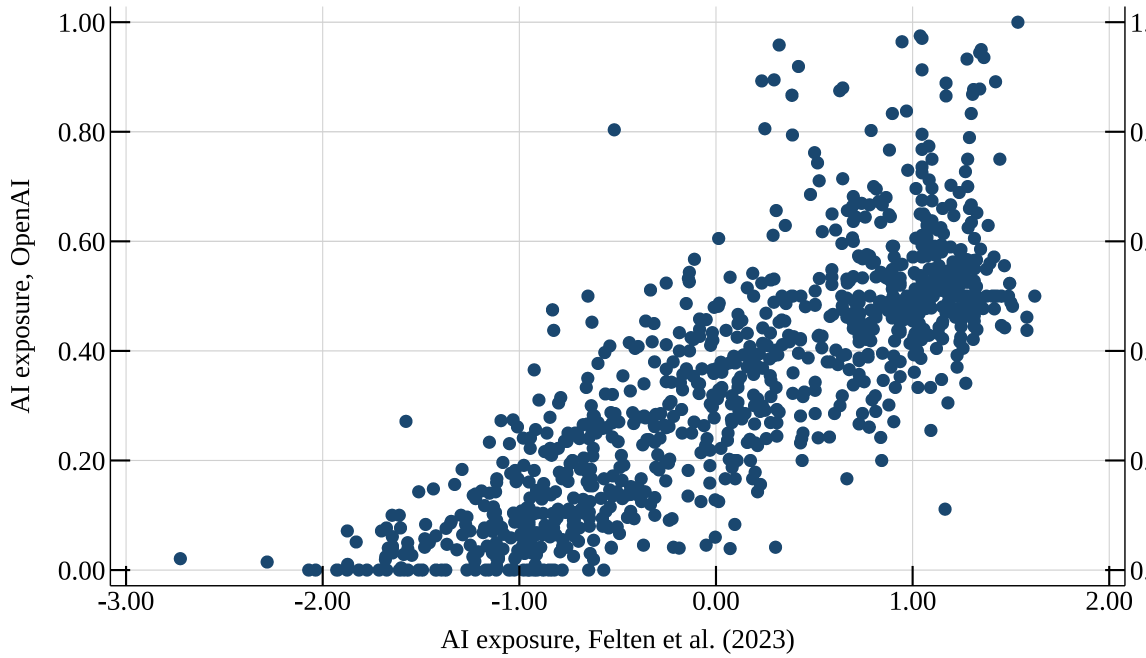

Notes: This figure displays the relationship between AI occupational exposure from Eloundou et al. (2023) (OpenAI) and Felten et al. (2023).

the use of complementary tools or software on top of LLMs. Exposure scores for 8-digit O*NET occupations are weighted averages over task-level exposure, where supplemental O*NET tasks are assigned half the weight of core tasks.

In robustness checks, we compare our AI exposure measure to the occupation-level AI exposure index of Felten et al. (2023). Both measures are constructed at the occupation level and aim to capture the degree to which workers in a given occupation are exposed to AI capabilities, but they differ in methodology: our measure follows Eloundou et al. (2023) and is based on GPT-4 assessments of the fraction of tasks within an occupation that can be performed by large language models, whereas Felten et al. (2023) link O*NET occupational ability requirements to progress benchmarks across AI applications. Despite these methodological differences, Figure 1 reveals a robust positive correlation between the two exposure measures across occupations. This is reassuring, as it suggests that our measure captures meaningful variation in the susceptibility of occupational tasks to automation by LLMs rather than idiosyncratic features of the underlying construction approach.[6]

[6] In addition, we use an alternative measure of AI exposure based on data from the Anthropic Economic Index

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 (Massenkoff and McCrory, 2026). We therefore supplement our AI exposure scores with a measure of AI utilization constructed from online job postings. 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 2 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" Using the Lightcast data, we compute the share of job postings mentioning an AI related skill requirement for each 8-digit O*NET occupation.

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 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 adoption measure, we compare it to AI adoption rates from Census' BTOS. The BTOS survey asks firms whether they have used AI in any business function over the past two weeks (Bonney et al., 2026). Unlike the Lightcast data, BTOS is only reported at

(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 8-digit O*NET occupations as described above.

[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: EXAMPLE OF JOB POSTINGS REQUIRING AI AND ML SKILLS

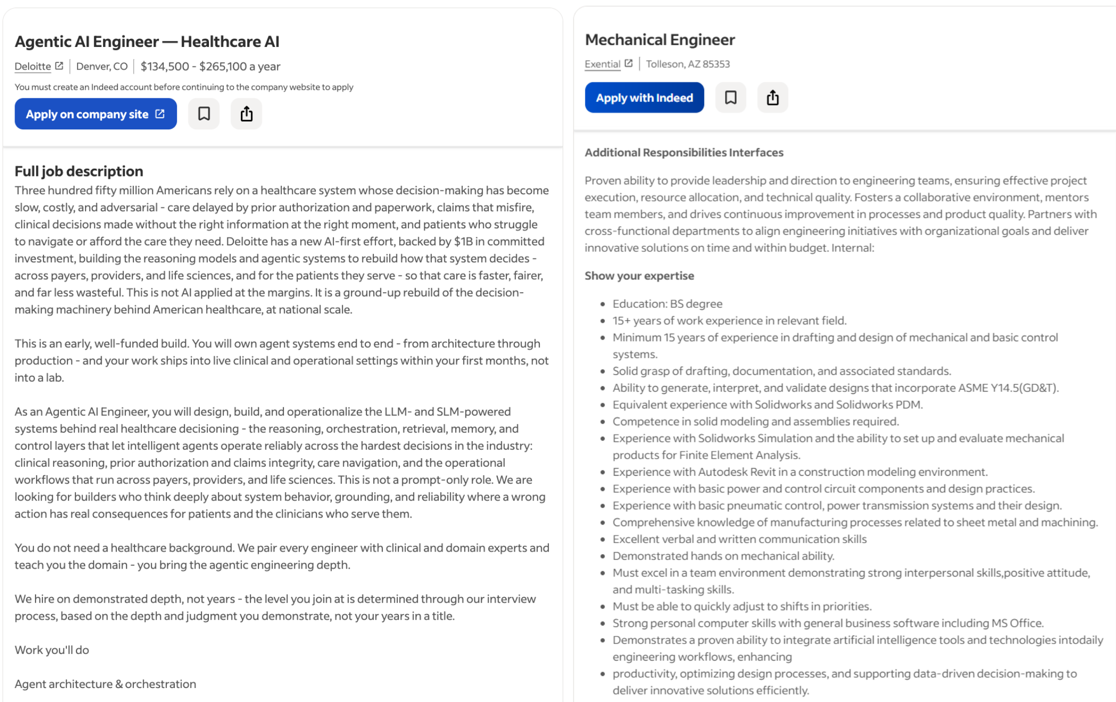

Agentic AI Engineer — Healthcare AI

Deloitte | Denver, CO | $134,500 - $265,100 a year

You must create an Indeed account before continuing to the company website to apply

Apply on company site

Full job description

Three hundred fifty million Americans rely on a healthcare system whose decision-making has become slow, costly, and adversarial - care delayed by prior authorization and paperwork, claims that misfire, clinical decisions made without the right information at the right moment, and patients who struggle to navigate or afford the care they need. Deloitte has a new AI-first effort, backed by $1B in committed investment, building the reasoning models and agentic systems to rebuild how that system decides - across payers, providers, and life sciences, and for the patients they serve - so that care is faster, fairer, and far less wasteful. This is not AI applied at the margins. It is a ground-up rebuild of the decision-making machinery behind American healthcare, at national scale.

This is an early, well-funded build. You will own agent systems end to end - from architecture through production - and your work ships into live clinical and operational settings within your first months, not into a lab.

As an Agentic AI Engineer, you will design, build, and operationalize the LLM- and SLM-powered systems behind real healthcare decisioning - the reasoning, orchestration, retrieval, memory, and control layers that let intelligent agents operate reliably across the hardest decisions in the industry: clinical reasoning, prior authorization and claims integrity, care navigation, and the operational workflows that run across payers, providers, and life sciences. This is not a prompt-only role. We are looking for builders who think deeply about system behavior, grounding, and reliability where a wrong action has real consequences for patients and the clinicians who serve them.

You do not need a healthcare background. We pair every engineer with clinical and domain experts and teach you the domain - you bring the agentic engineering depth.

We hire on demonstrated depth, not years - the level you join at is determined through our interview process, based on the depth and judgment you demonstrate, not your years in a title.

Work you'll do

Agent architecture & orchestration

Mechanical Engineer

Exential | Tolleson, AZ 85353

Apply with Indeed

Additional Responsibilities Interfaces

Proven ability to provide leadership and direction to engineering teams, ensuring effective project execution, resource allocation, and technical quality. Fosters a collaborative environment, mentors team members, and drives continuous improvement in processes and product quality. Partners with cross-functional departments to align engineering initiatives with organizational goals and deliver innovative solutions on time and within budget. Internal:

Show your expertise

- Education: BS degree
- 15+ years of work experience in relevant field.
- Minimum 15 years of experience in drafting and design of mechanical and basic control systems.
- Solid grasp of drafting, documentation, and associated standards.
- Ability to generate, interpret, and validate designs that incorporate ASME Y14.5(GD&T).
- Equivalent experience with Solidworks and Solidworks PDM.
- Competence in solid modeling and assemblies required.
- Experience with Solidworks Simulation and the ability to set up and evaluate mechanical products for Finite Element Analysis.
- Experience with Autodesk Revit in a construction modeling environment.
- Experience with basic power and control circuit components and design practices.
- Experience with basic pneumatic control, power transmission systems and their design.
- Comprehensive knowledge of manufacturing processes related to sheet metal and machining.
- Excellent verbal and written communication skills
- Demonstrated hands on mechanical ability.
- Must excel in a team environment demonstrating strong interpersonal skills,positive attitude, and multi-tasking skills.
- Must be able to quickly adjust to shifts in priorities.
- Strong personal computer skills with general business software including MS Office.
- Demonstrates a proven ability to integrate artificial intelligence tools and technologies intodaily engineering workflows, enhancing
- productivity, optimizing design processes, and supporting data-driven decision-making to deliver innovative solutions efficiently.

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.

Figure 3: CORRELATION BETWEEN LIGHTCAST AND BTOS AI ADOPTION

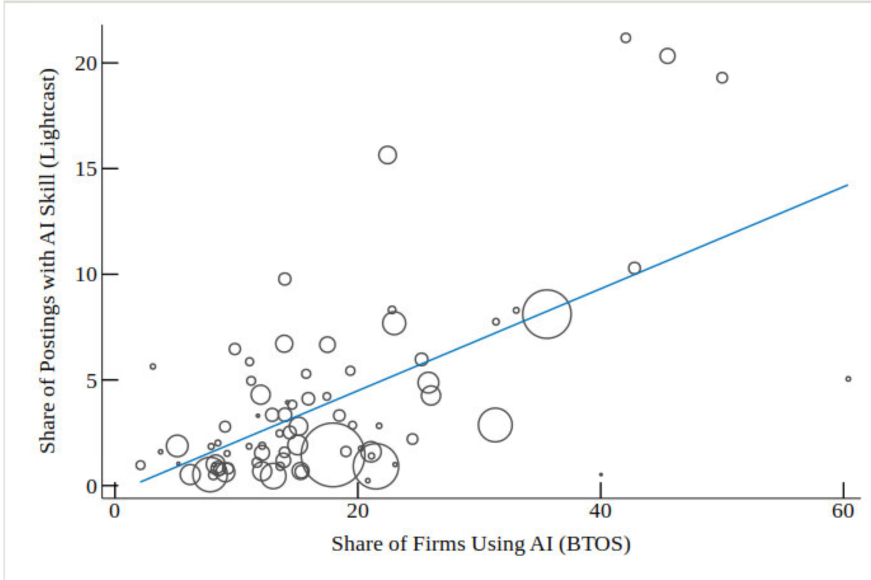

Note: This figure plots correlations between Lightcast- and BTOS-based AI adoption measures at the 3-digit NAICS level. The Lightcast measure is the share of job postings tagged with an AI skill from 2025:M1 to 2025:M12. The BTOS measure is the share of firms reporting having used AI in any business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

the 3-digit NAICS level and has only become available recently. To compare the two datasets, we therefore aggregate our Lightcast adoption data to the industry definitions used in the BTOS. Figure 3 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 AI exposure and adoption scores based on their reported occupation and industry. Exposure is based on the individual's reported occupation and the indices defined in Section 3.1. When CPS occupation categories combine more than one detailed O*NET occupation, we use the average exposure score of the CPS (4-digit) occupation's O*NET (8-digit) components. We additionally harmonize CPS occupations—revised to reflect the 2018 Standard Occupational Classification starting in 2020:M1—to definitions comparable from 2011:M1 to 2026:M7. Exposure scores therefore vary by occupation, by not by industry or time.

We assign AI adoption based on the individual's three-digit industry in the month of the interview. To do this, we first aggregate the industry-level (3-digit industries, `ind1990`) adoption scores at each point in time from Lightcast to CPS occupations as described above. The adoption scores we use therefore vary by industry and time. We consider variation across industry rather than by occupation and industry jointly for two reasons. First, the workers implementing AI applications within a firm may differ from the workers who are substituted by AI. For example, a firm may hire software engineers to implement an AI ChatBot to substitute for customer service workers. We therefore see demand for AI as better captured by employers' characteristics (which industry represents). Second, a joint categorization results in roughly 50,000 total occupation-industry cells so the adoption measure becomes quite noisy and sparse.

3.3 Imputation of Missing AI Exposure and Adoption

AI exposure is defined at the occupation level and AI adoption at the industry level; these measures can therefore be assigned only to individuals with available occupation and industry information.[8] Occupation and industry are observed for all employed workers. Unemployed individuals report their previous industry and occupation, but about 10% either lack such information or report "military occupations" without AI exposure or adoption scores. For example, some individuals transitioning from employment to unemployment lack occupation information in the current month, likely due to measurement error; others without work histories lack occupation information altogether. This issue is more severe among nonparticipants and requires more involved imputation of missing AI scores; we therefore focus on employment and unemployment and flows between these two states.

We use AI exposure and adoption measures for the current month when available, but impute the scores of AI exposure and adoption for unemployed individuals without the scores. To maximize the use of each individual's existing AI scores, we implement the following procedure. Figure 4 summarizes the imputation procedure.

1. **Step 1**: Use the individual's lagged values. 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 the exposure and adoption

[8] Individuals with occupation information also have industry information.

Figure 4: IMPUTATION PROCEDURE

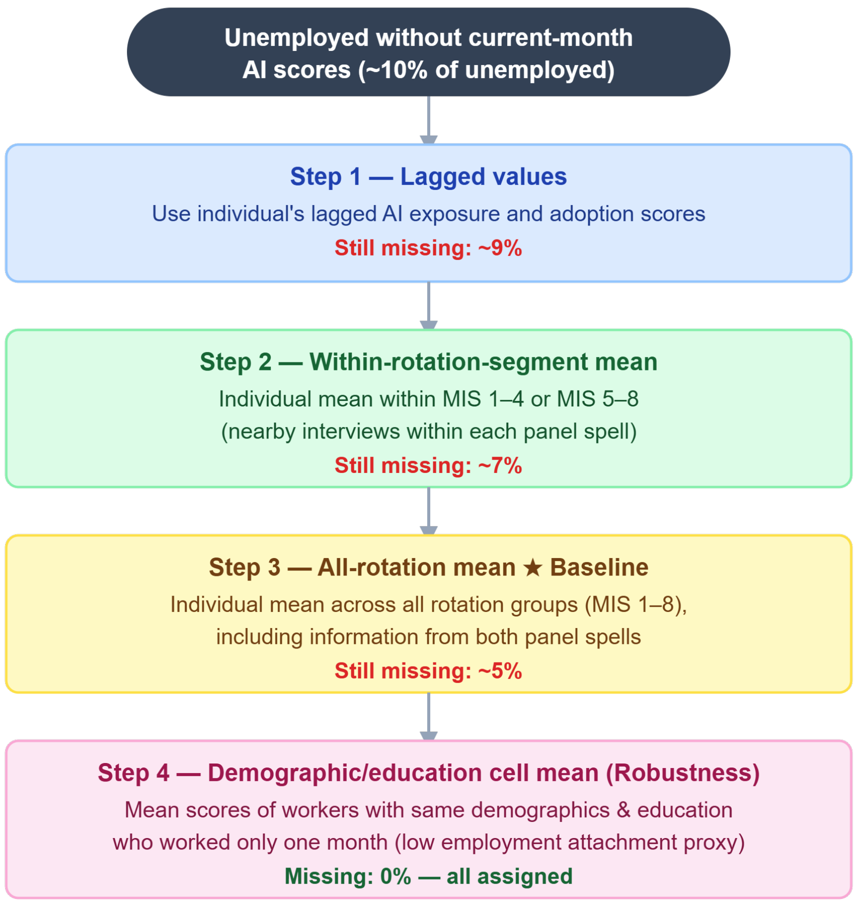

scores using the individual-specific mean within the same CPS rotation-panel segment, defined separately for MIS 1–4 and MIS 5–8, thereby relying on information from nearby interviews within each panel spell. After this step, about 7% of the individuals do not have AI information.

3. **Step 3**: Any remaining missing values are assigned the individual's average exposure and adoption scores across all the observed rotation groups, 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 who still lack exposure and adoption scores — likely those without recent work histories — we assign the mean scores of individuals with the same demographic and education characteristics who worked only one month during the survey, capturing their low attachment to employment. After this step, all unemployed individuals are assigned an AI exposure.

This procedure yields unemployment stocks and flows that are largely internally consistent. We consider two cases—with and without Step 4.[9] Omitting the final step and excluding unemployed individuals without occupation information does not materially affect the conclusion.[10]

3.4 Vacancies and Labor Market Tightness

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,

[9] 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:M7, reflecting the increase in new entrants' unemployment rate over this period. 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.

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

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

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

Next, we compute the stock of unemployment by occupation using data from the CPS, U_{ot}. To do this, we multiply the stock of unemployed workers at time t, U_t by share of workers who transitioned from unemployment to employment in occupation o, averaged over the previous 12 months. The sum of U_{ot} over all occupations therefore equals aggregate unemployment U_t.[11] To compute market tightness, we tag occupations as belonging to high and low exposure/adoption groups, then aggregate vacancies and unemployment within these groups. For each adoption/exposure (AE) group, market tightness is defined as

$$\theta_{AE,t} = \frac{\sum_{o \in AE} V_{ot}}{\sum_{o \in AI} U_{ot}}. \tag{3.2}$$

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%, 25–50%, 50–75%, 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 about 4 percent of the labor force. Instead of using employment-weighted quartiles, we adopt this approach because AI may have had particularly large effects on within-job activity changes and occupational switching among occupations in the right tail of AI exposure, such as software developers.

Next, AI adoption, measured as the share of job postings in an industry requiring AI-related skills at each point in time, ranges from near zero to approximately 15 percent across industries

[11] This approach ensures consistency with aggregate market tightness. Alternatively, one can use the number of unemployed individuals in each occupation bin constructed directly from the CPS microdata; this yields essentially the same market tightness by AI exposure and adoption group.

in our sample. Unlike AI exposure, for which we use fixed task-share thresholds, we classify industries into four adoption groups using employment-weighted percentile cutpoints: below the 25th percentile, 25th–50th, 50th–90th, and 90th percentile or above. We adopt employment-weighted quantiles here because AI adoption varies continuously across industries without natural breakpoints, and weighting by employment ensures that the groups reflect the distribution of workers rather than the distribution of industries. The bottom quartile is dominated by large, low-skill service industries — such as food services, hospitals, and trucking — which together account for a substantial share of employment despite their low adoption rates. By contrast, the top decile, while comprising only a small number of industries, captures knowledge-intensive sectors such as computer and data processing, accounting and bookkeeping, and securities brokerage, where AI skill requirements have risen sharply. As with exposure, we are particularly interested in whether AI adoption at the right tail — where LLM use is most intensive — has had outsized effects on labor market transitions.

Last, we consider a joint categorization of four groups based on whether AI exposure is above or below its 2025 mean and whether AI adoption is above or below its period-t mean.[12] The purpose of this approach is to examine the interaction effects of AI exposure and adoption on labor market outcomes.[13]

Figure 5 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 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.

[12] For AI exposure, the choice of reference year is inconsequential, since AI exposure is time-invariant at the occupation level; using alternative years yields identical classifications. For AI adoption, the median of AI adoption observed since 2012 also yields similar results. We also consider alternative cutoffs — specifically a higher threshold for AI adoption — in this two-by-two classification to examine potential nonlinear effects on job-finding and job-loss rates. The results are qualitatively similar overall.

[13] We also consider quartiles of AI adoption; the patterns are similar to those by AI exposure, though less pronounced.

Figure 5: Representative Occupations by AI Exposure

Group 1 — Lowest Exposure (34%)

- ◆ Cashiers
- ◆ Janitors and building cleaners
- ◆ Laborers and freight, stock, and material movers, hand
- ◆ Nursing, psychiatric, and home health aides
- ◆ Cooks
- ◆ Waiters and waitresses
- ◆ Construction laborers
- ◆ Stockers and order fillers
- ◆ Maids and housekeeping cleaners
- ◆ Personal care aides

Group 2 — Mid-Low Exposure (43%)

- ◆ Managers, all other
- ◆ Driver/sales workers and truck drivers
- ◆ Elementary and middle school teachers
- ◆ First-line supervisors of retail sales workers
- ◆ Registered nurses
- ◆ Retail salespersons
- ◆ Chief executives
- ◆ First-line supervisors, office and admin. support
- ◆ Financial managers
- ◆ Postsecondary teachers

Group 3 — Mid-High Exposure (19%)

- ◆ Secretaries and administrative assistants
- ◆ Customer service representatives
- ◆ Accountants and auditors
- ◆ Office clerks, general
- ◆ Receptionists and information clerks
- ◆ Sales representatives, wholesale and manufacturing
- ◆ First-line supervisors, non-retail sales
- ◆ Computer occupations, all other
- ◆ Computer and information systems managers
- ◆ Insurance sales agents

Group 4 — Highest Exposure (4%)

- ◆ Software developers, applications and systems
- ◆ Bookkeeping, accounting and auditing clerks
- ◆ Billing and posting clerks
- ◆ Computer programmers
- ◆ Insurance claims and policy processing clerks
- ◆ Data entry keyers
- ◆ Writers and authors
- ◆ Other mathematical science occupations
- ◆ File clerks
- ◆ Web developers

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

Figure 6 presents the top ten industries by AI adoption score within each group, ranked by employment weight. The bottom quartile consists of industries where tasks are predominantly physical, interpersonal, or require in-person presence — eating and drinking places, hospitals, trucking, department stores, and building services. The second quartile is more heterogeneous, mixing public-sector and community-oriented industries (elementary and secondary schools, public safety and justice) with consumer-facing services (grocery stores, health services, entertainment and recreation). The 50th-90th percentile group shifts toward more knowledge-intensive industries – construction, colleges and universities, insurance, and real estate – where professional and analytical tasks are more prevalent but AI skill requirements remain moderate. The top decile is dominated by industries at the frontier of AI adoption, where demand for AI-related skills in job postings is highest: computer and data processing services, management and public relations consulting, securities and investment companies, and accounting and

Figure 6: Representative Industries by AI Adoption Group

Bottom 25%
Lowest AI Adoption

- Eating And Drinking Places
- Hospitals
- Trucking Service
- Department Stores
- Services To Dwellings And Other Buildings
- Child Day Care Services
- Nursing And Personal Care Facilities
- Automotive Repair And Related Services
- Hotels And Motels
- Lumber And Building Material Retailing

25th–50th Percentile
Low AI Adoption

- Elementary And Secondary Schools
- Health Services, N.E.C
- Miscellaneous Entertainment And Recreation Services
- Grocery Stores
- Justice, Public Order, And Safety
- Landscape And Horticultural Services
- Offices And Clinics Of Physicians
- Executive And Legislative Offices
- Agricultural Production, Crops
- Motor Vehicle Dealers

50th–90th Percentile
Mid AI Adoption

- All Construction
- Colleges And Universities
- Insurance
- Real Estate, Including Real Estate-Insurance Offices
- Social Services, N.E.C
- Business Services, N.E.C
- Banking
- Engineering, Architectural, And Surveying Services
- Legal Services
- Administration Of Human Resources Programs

Top 10%
High AI Adoption

- Computer And Data Processing Services
- Management And Public Relations Services
- Security, Commodity Brokerage, And Investment Companies
- Accounting, Auditing, And Bookkeeping Services
- Motor Vehicles And Motor Vehicle Equipment
- Credit Agencies, N.E.C
- Electrical Machinery, Equipment, And Supplies, N.E.C
- Research, Development, And Testing Services
- Drugs

Notes: This figure reports the ten most representative industries in each of the four AI adoption groups. Numbers in parentheses are the group's share in the labor force. **Source**: Authors' calculation.

bookkeeping services. Notably, Group 4 includes not only technology-intensive industries but also professional services sectors —- such as accounting and investment — whose core tasks of processing financial data and documents make them natural early adopters of large language

Figure 7: REPRESENTATIVE OCCUPATIONS BY AI EXPOSURE/ADOPTION GROUP

Low-Exposure / Low-Adoption (41%)

- Elementary and middle school teachers
  Elementary And Secondary Schools
- Construction laborers
  All Construction
- Waiters and waitresses
  Eating And Drinking Places
- Driver/sales workers and truck drivers
  Trucking Service
- Cooks
  Eating And Drinking Places
- Grounds maintenance workers
  Landscape And Horticultural Services
- Carpenters
  All Construction
- Couriers and messengers
  Trucking Service
- Nursing, psychiatric, and home health aides
  Health Services, N.E.C
- Fast food and counter workers
  Eating And Drinking Places

Low-Exposure / High-Adoption (6%)

- Taxi drivers and chauffeurs
  Taxicab Service
- Other assemblers and fabricators
  Motor Vehicles And Motor Vehicle Equipment
- Retail salespersons
  Radio, Tv, And Computer Stores
- Aircraft pilots and flight engineers
  Air Transportation
- Flight attendants
  Air Transportation
- Counselors
  Colleges And Universities
- Health diagnosing and treating practitioner support technicians
  Drug Stores
- Printing press operators
  Printing, Publishing, And Allied Industries, Except Newspap
- Janitors and building cleaners
  Real Estate, Including Real Estate-Insurance Offices
- Electrical power-line installers and repairers
  Electric Light And Power

High-Exposure / Low-Adoption (29%)

- Registered nurses
  Hospitals
- Construction managers
  All Construction
- Food service managers
  Eating And Drinking Places
- Teacher assistants
  Elementary And Secondary Schools
- Managers, all other
  All Construction
- First-line supervisors of construction trades and extraction workers
  All Construction
- Physicians and surgeons
  Hospitals
- Registered nurses
  Health Services, N.E.C
- First-line supervisors of retail sales workers
  Grocery Stores
- Education and childcare administrators
  Elementary And Secondary Schools

High-Exposure / High-Adoption (23%)

- Software developers, applications and systems software
  Computer And Data Processing Services
- Real estate brokers and sales agents
  Real Estate, Including Real Estate-Insurance Offices
- Lawyers, and judges, magistrates and other judicial workers
  Legal Services
- Postsecondary teachers
  Colleges And Universities
- Management analysts
  Management And Public Relations Services
- Insurance sales agents
  Insurance
- Accountants and auditors
  Accounting, Auditing, And Bookkeeping Services
- Property, real estate, and community association managers
  Real Estate, Including Real Estate-Insurance Offices
- Managers, all other
  Computer And Data Processing Services
- Financial managers
  Banking

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

models.

Figure 7 shows the largest occupation–industry cells within each exposure–adoption quadrant. The low-exposure/low-adoption (LL) group, which accounts for 41 percent of the labor force, includes elementary school teachers in K–12 education, construction laborers, and waiters and waitresses in food services, reflecting the physically grounded or interpersonally intensive nature of the work. The high-exposure/low-adoption (HL) group (30 percent) is anchored by registered nurses and other healthcare workers in hospitals — occupations whose tasks are susceptible to AI assistance but whose employing industries have yet to translate this potential into widespread AI adoption. The high-exposure/high-adoption (HH) group (23 percent) includes software developers in computer and data processing services, lawyers and judges in legal services, and financial managers in banking — workers whose tasks are both AI-exposed and employed in industries that have actively integrated AI tools. The small low-exposure/high-adoption (LH) group (6 percent) contains taxi drivers and chauffeurs, assemblers in motor vehicle manufacturing, and retail salespersons in electronics stores, where industry-level AI investment has outpaced the direct task-level exposure of the workforce. Since HH and LL workers together account for roughly two-thirds of the labor force, task exposure to AI and actual AI use largely move together — but the substantial HL share points to a meaningful gap between AI potential and adoption in healthcare.

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.[14]

[14] Observable attributes of workers in each group are available upon request.

4 Stocks and Flows by AI Exposure and Adoption

4.1 Unemployment Rate and Job Finding and Loss Rates

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 intensity. First, as shown in Panel A, higher AI exposure is generally associated with lower unemployment rates, driven mainly by lower separation rates for the most exposed workers, though the relationship is not strictly monotonic across all four groups. 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.

Panel B shows a broadly similar pattern when grouping by AI adoption alone: unemployment, EU, and UE rates are generally lower in higher-adoption groups, though the middle two adoption groups show little difference from each other. The joint classification (Panel C) reinforces this picture. The low-exposure/low-adoption group has the highest EU rate, UE rate, and job-switch rate, while the high-exposure/high-adoption group has the lowest UE rate and is among the lowest in EU rate and job-switch rate (tied with the high-exposure/low-adoption group on both). Within the low-exposure group, higher AI adoption is associated with a lower EU rate as well as a lower UE rate, perhaps reflecting less churn in high-adoption jobs on average; within the high-exposure group, however, the EU rate is essentially unchanged with adoption, while the UE rate still falls.

Figure 8 displays the time series by AI exposure.[15] Several patterns emerge. First, the unemployment rate of the top AI-exposure group rose noticeably following the introduction of LLMs

[15] We focus on UE and EU flows since our primary focus is changes in the unemployment rate. Reallocation is also measured with a flows per capita concept—the monthly rate of labor force transition, which includes flows associated with nonparticipation (N). Since N-to-E, U-to-N, and N-to-U flows require substantial imputation for individuals without occupation information or work histories, we do not consider them in our baseline analysis. We provide these flows in Appendix Section E and the flows per capita in Figure A.6.

Table 1: AVERAGES OF UNEMPLOYMENT RATE AND SELECTED FLOWS RATES (2012:M1 – 2026:M7)

Group	EU rate	UE rate	Unemployment rate	Job Switch rate	Task Switch rate
Panel A: AI Exposure					
Exposure 1	2.0	27.4	7.8	2.7	1.1
Exposure 2	0.9	27.0	3.5	1.9	1.1
Exposure 3	0.8	23.8	4.0	2.0	1.1
Exposure 4	0.8	24.3	3.8	1.9	1.1
Panel B: AI Adoption					
Adoption 1 (0–25)	1.4	26.7	6.0	2.4	1.1
Adoption 2 (25–50)	1.3	27.1	5.1	2.2	1.1
Adoption 3 (50–90)	1.2	27.2	4.9	2.1	1.1
Adoption 4 (90+)	0.9	23.3	4.1	2.0	1.1
Panel C: Exposure × Adoption					
Low–Low	1.8	28.1	6.9	2.6	1.1
Low–High	1.7	26.5	6.7	2.5	1.1
High–Low	0.8	25.7	3.3	1.9	1.1
High–High	0.8	22.8	3.5	1.9	1.1

Notes: The exposure groups are defined by the share of tasks in a worker's current occupation that are exposed to AI: below 25%, 25–50%, 50–75%, and 75% or higher (fixed cutpoints). The adoption groups are defined by the labor-force-weighted monthly 25th, 50th, and 90th percentiles of AI adoption (recomputed every month): below the 25th percentile, between the 25th and 50th, between the 50th and 90th, and the top decile (90th percentile or higher). The exposure × adoption groups are defined jointly by AI exposure and adoption—above versus below the mean of AI exposure in 2025 (fixed cutoff) and above versus below the labor-force-weighted monthly mean of AI adoption (recomputed every month).

Source: Authors' calculation.

Figure 8: UNEMPLOYMENT RATE, UE AND EU RATES AND LABOR FORCE SHARE (AI EXPOSURE)

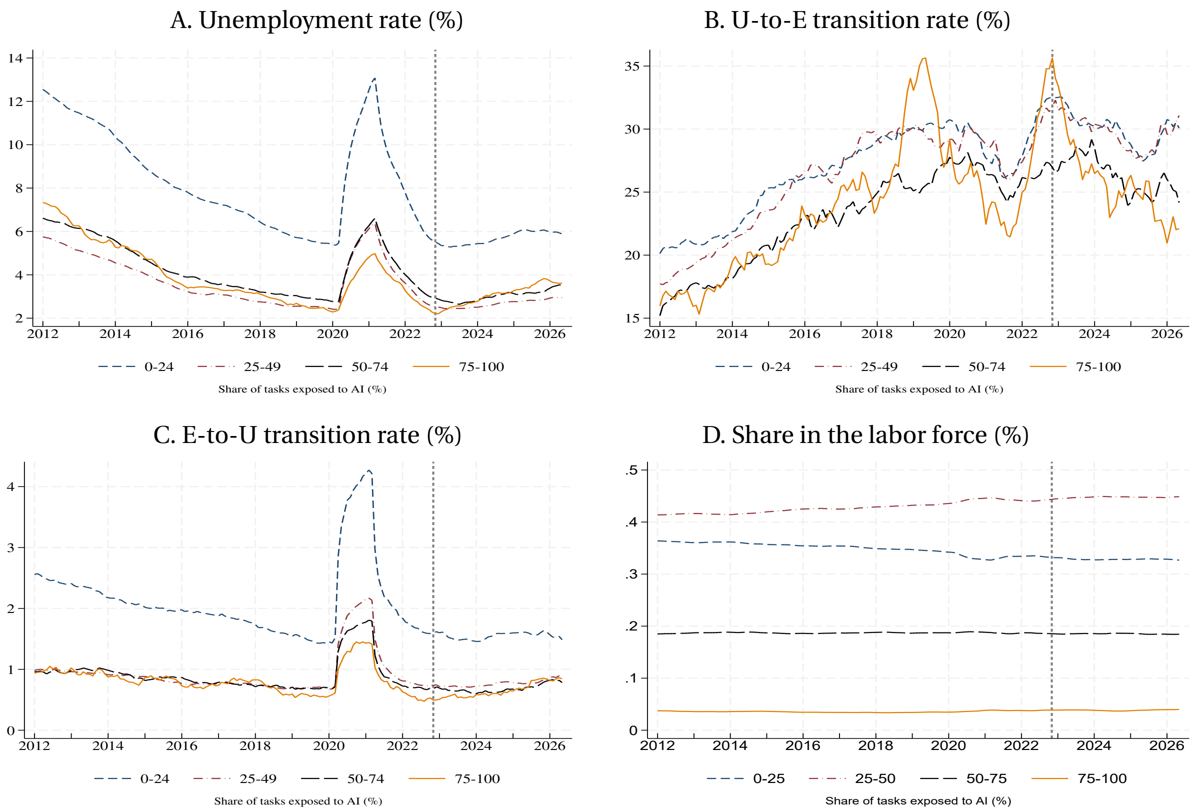

Notes: This figure displays the unemployment rate (Panel A), UE rate (Panel B), EU rate (Panel C), and labor force shares (Panel D) for four AI exposure groups (12-month moving average). Blue dashed line: 0–25; red dashed-dot line: 25–50; black long-dashed line: 50–75; solid yellow line: 75–100. **Source**: Authors' calculation (dotted vertical line, end-2022): by 1.3 pp in 2026:M7 and 1.7 pp in 2025:Q4 (yellow line, Panel A). The corresponding increases for the bottom exposure group were 0.4 pp and 0.7 pp, and +0.6 pp in 2026:M7 for the 50–75% group.[16] This divergence is driven mainly by the hiring margin: the job-finding probabilities of the highest-exposure workers declined persistently and by more than those of other groups (yellow line, Panel B). The separation rate of the top exposure group also edged up more than others, though the difference is less pronounced than in the job-finding rate (yellow line, Panel C). The labor force shares of the four exposure groups are 34%, 43%, 19%, and 4%, from bottom to top. The shares of the top two groups are little changed over time, while that of the bottom group has declined steadily and that of the second-lowest group has risen gradually (Panel D).

Notably, the unemployment rates of lower-exposure groups are more cyclically sensitive than those of higher-exposure groups. The faster rise in unemployment among high-exposure workers therefore likely reflects factors specific to them—above all, the introduction of LLMs. The larger increase in unemployment, the steeper decline in job-finding rates, and the somewhat faster rise in separation rates together point to a disproportionate decline and shift in labor demand for high AI-exposure workers.

Figure 9 displays labor market outcomes by industry AI adoption quantile group. Among the four groups, the top decile—industries at or above the 90th percentile of the AI adoption distribution—is the least cyclical. Its unemployment rate rose by 1.2 percentage points between December 2022 and July 2026. By contrast, the bottom quartile group is the most cyclical, with its unemployment rate rising by only 0.4 percentage point over the same period. The divergence in job-finding rates is particularly noticeable. The UE rate of the top decile group fell sharply from late 2022, reaching levels comparable to the COVID-era trough by 2023, and has remained low since. The other three groups experienced more modest or mixed trajectories—declining slowly, or declining and then partially recovering. The EU rate tells a more muted story: between-group differences are small, suggesting that adjustment in high-adoption industries has operated primarily through reduced hiring rather than elevated separations. These patterns indicate that AI has had disproportionately larger negative effects on job-finding rates and unemployment for

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

Figure 9: UNEMPLOYMENT RATE, UE AND EU RATES AND LABOR FORCE SHARE (AI ADOPTION)

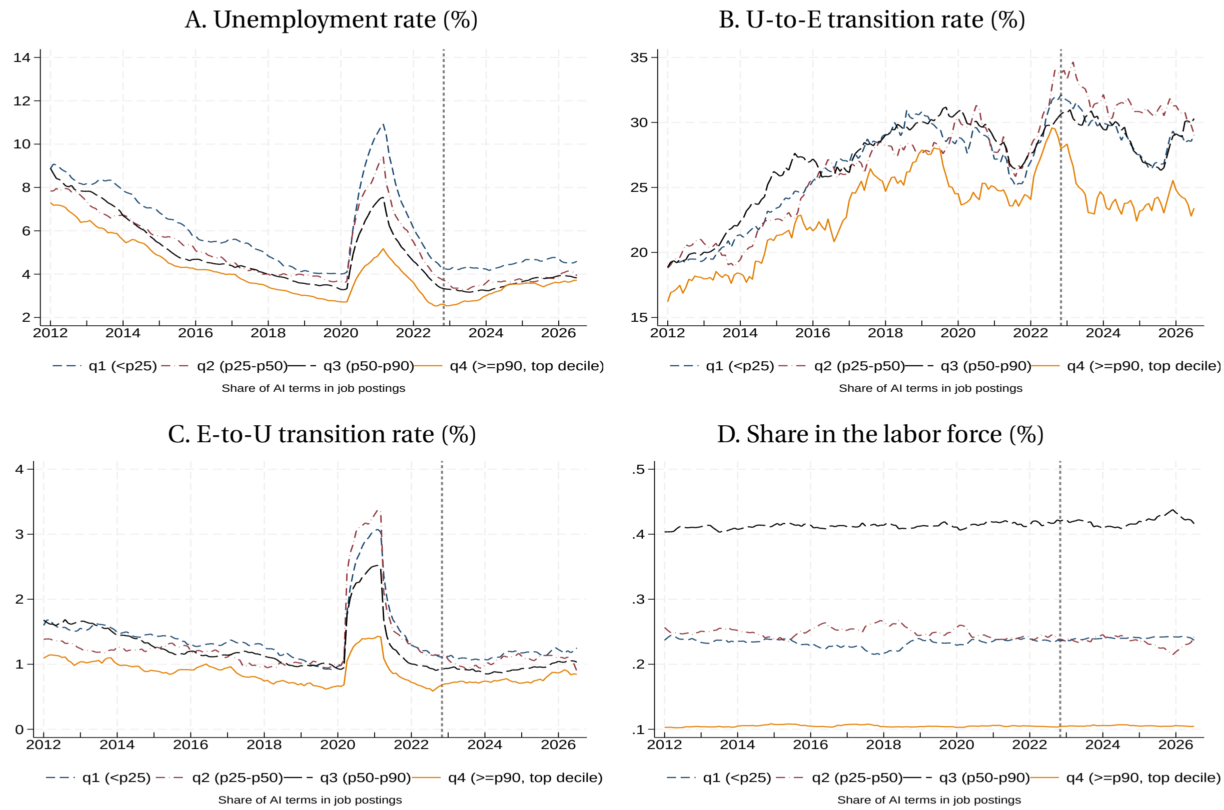

Notes: This figure displays the unemployment rate (Panel A), UE rate (Panel B), EU rate (Panel C), and labor force shares (Panel D) for four AI adoption groups (12-month moving average). Blue dashed line: bottom quartile; red dashed-dot line: second quartile; black long-dashed line: 50–90%; solid yellow line: 90+%. **Source**: Authors' calculation workers in the highest-adoption industries.[17]

Figure 10 displays the analogous data for four groups defined jointly by AI exposure and adoption. The high-exposure/high-adoption (HH) group shows patterns similar to the top exposure group. Their unemployment rate rose by slightly below 1 percentage point between end-2022 and July 2026, compared with 0.3 percentage point for the low-exposure/low-adoption (LL) group and about half a percentage point for the high-exposure/low-adoption (HL) group (Panel A). The low-exposure/high-adoption (LH) group's unemployment rate edged lower. The HH and HL groups together comprise about a half of the labor force; the LL group accounts for slightly above 40% and the LH group for less than 10%.

The rise in the HH group's unemployment rate is also driven mainly by a steady decline in the UE rate from end-2022, a pattern much less pronounced in the other groups (yellow line, Panel B). The EU rates of the high-exposure groups edged up in 2025–2026, but differences across groups remain small. The HH group's increase in unemployment rate is thus attributable primarily to a decline in the job-finding rate.

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

Switches of employers and job activities Figure 11 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-

[17] When the adoption distribution is divided into quartiles rather than the 0–25th, 25th–50th, 50th–90th, and top-decile grouping used here, this pattern is masked and the differential effects of AI adoption on unemployment and job-finding rates are not as noticeable.

Figure 10: UNEMPLOYMENT RATE AND FLOWS BY AI EXPOSURE AND ADOPTION

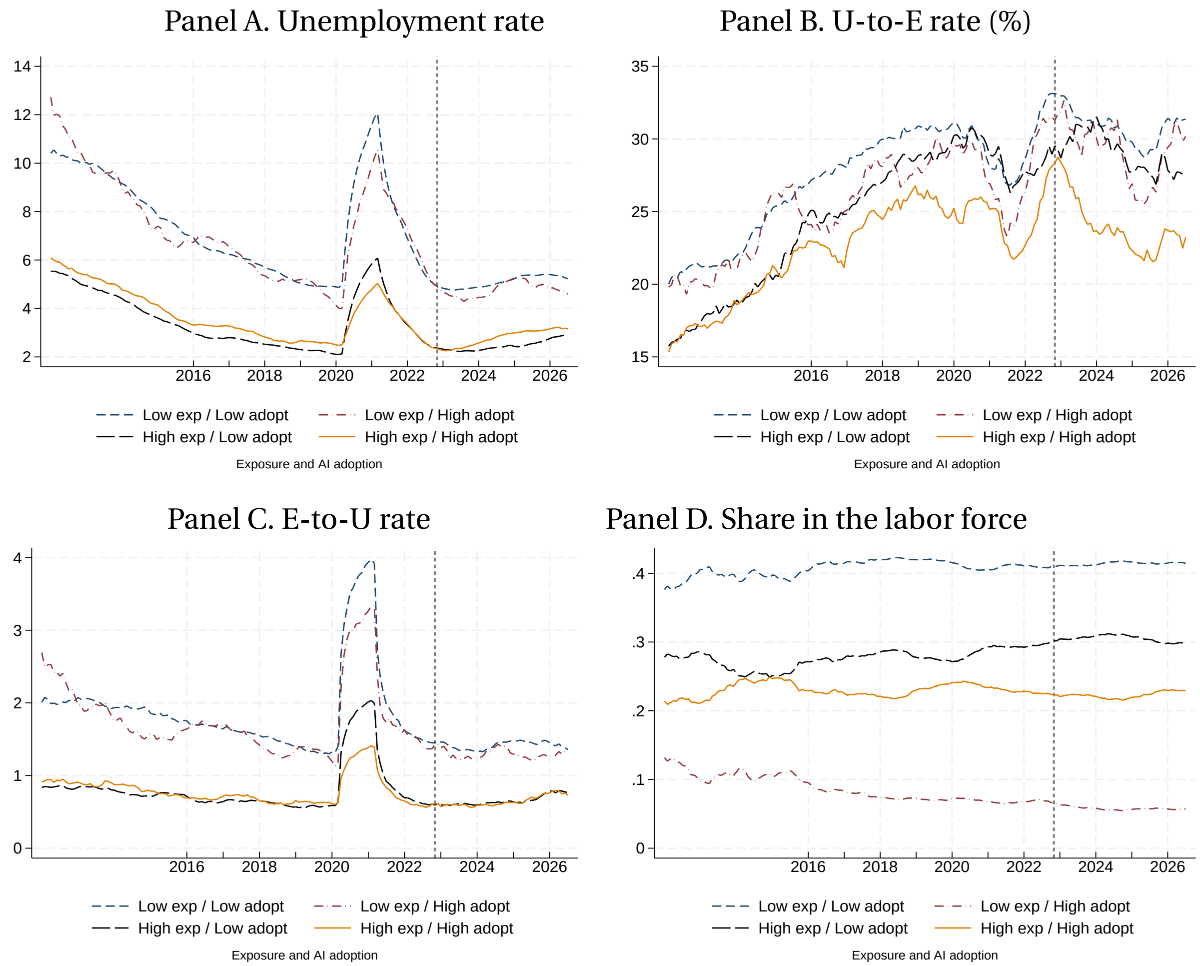

Note: This figure displays the unemployment rate (Panel A), UE rate (Panel B), EU rate (Panel C), and labor force shares (Panel D) for four AI exposure/adoption groups (12-month moving average). "Exp" stands for exposure, and "adopt" stands for adoption. Blue dashed line: Low-exp/Low-adopt; red dashed-dot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

Figure 11: SWITCHES OF JOBS AND ACTIVITIES BY AI EXPOSURE

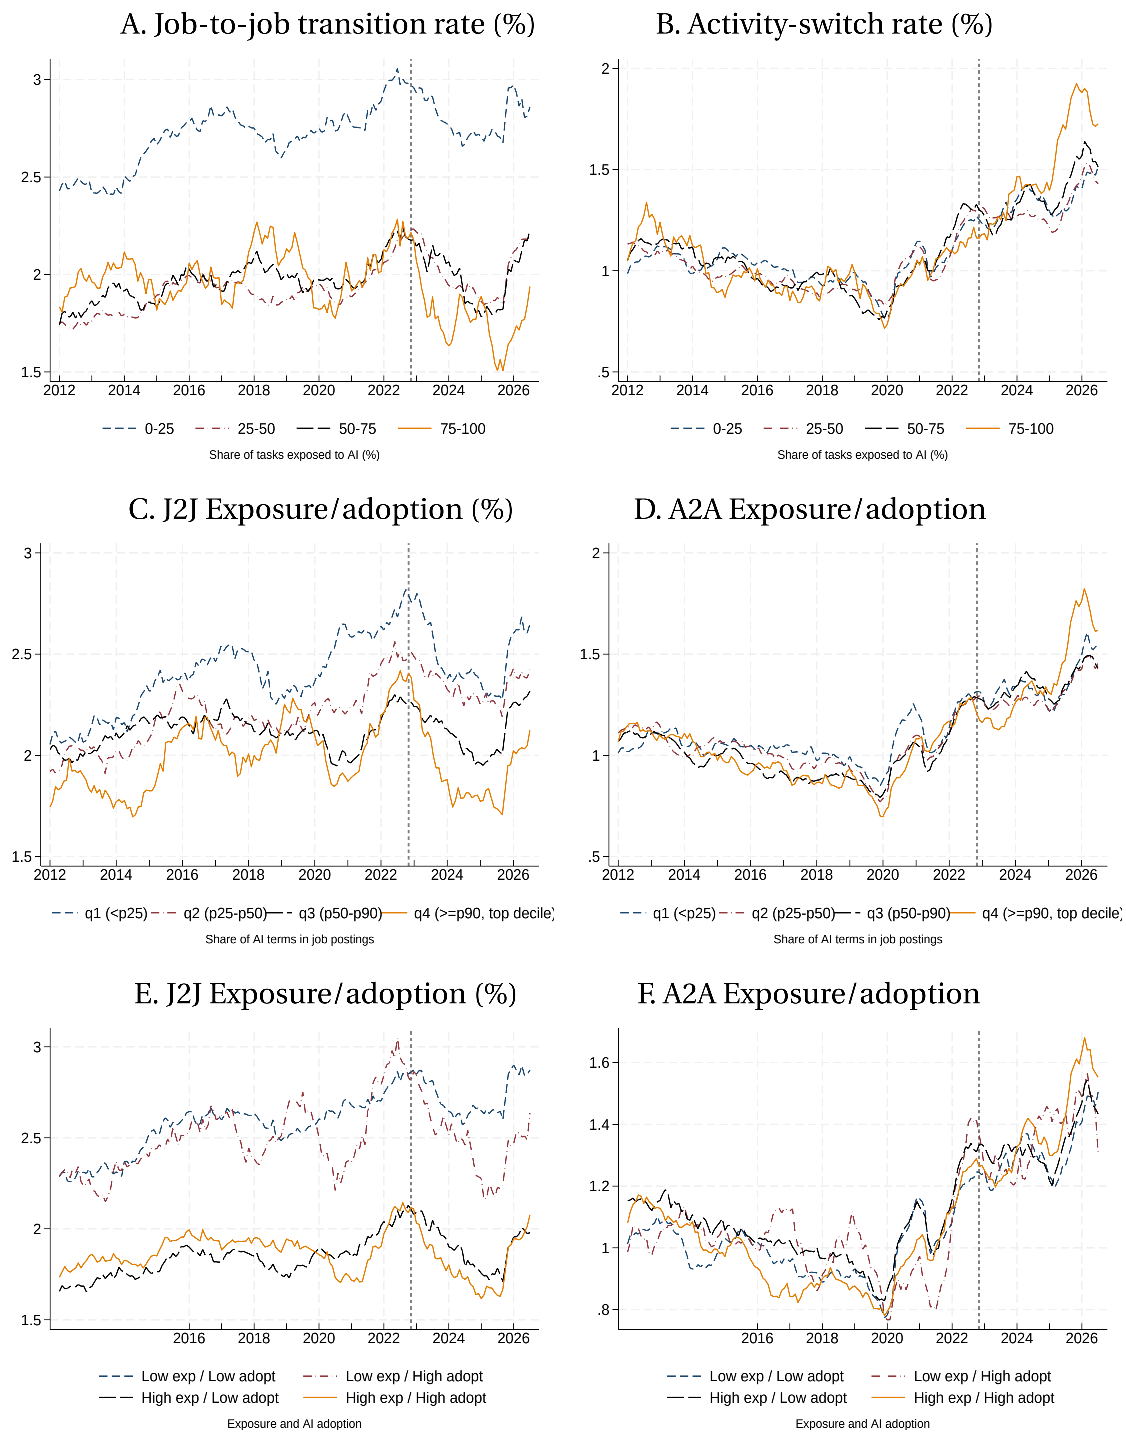

Notes: This figure displays job-to-job transition rates (Panel A) and activity-switch rates (Panel B) by AI exposure, by AI adoption (Panels C and D), and by the joint classification of exposure and adoption (Panels E and F). The 12-month moving averages are plotted. For Panels A and B: blue dashed line: 0–25 of AI exposure; 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: 0–25 percentile of AI adoption; red dashed-dot line: 25–50 percentile; black long-dashed line: 50–90 percentile; yellow solid line: 90 + percentile. For Panels E and F: 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

switch rates—another indicator of labor demand—declined more among workers with top AI exposure or adoption, and in the HH group than others. 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 uncovers a noticeable pattern in LLM-driven reallocation. While activity-switch rates were overall similar through 2024 across workers with different AI intensity, starting around 2025, the groups with the highest exposure or adoption began to show a noticeably faster increase in activity-switch rates than the other groups. This observation suggests that task reorganization is concentrated among workers in the right tail of the distribution of AI exposure or adoption, particularly in the HH group. 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 the four groups (LL, LH, HL, HH) defined by above- and below-mean AI exposure and adoption at the occupation level. 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 12, Panel A reports estimates based on OpenAI's AI exposure measure and Panel B uses the Anthropic measure. Workers with high AI exposure exhibit higher market tightness, consistent with their lower unemployment rates.[18] However, a notable divergence emerges over time. Market tightness for high-exposure workers has remained below its pre-COVID level since 2023, while that 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

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

Figure 12: MARKET TIGHTNESS BY AI-EXPOSURE AND ADOPTION

Panel A: OpenAI Panel B: Anthropic

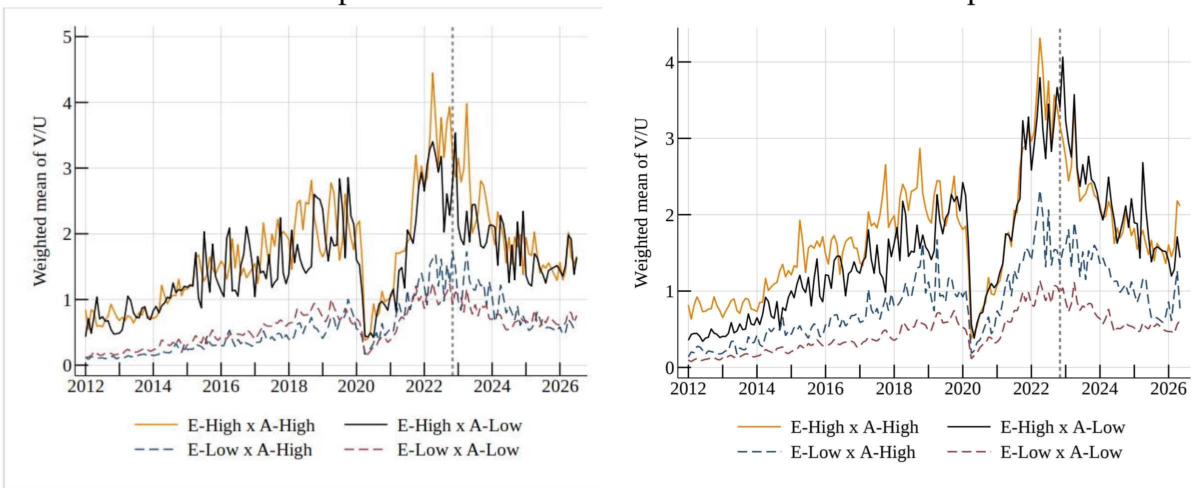

Notes: Panels A and B display market tightness by AI exposure and adoption, with the AI exposure measure constructed using the OpenAI and Anthropic data, respectively. Blue dashed line: Low-Exp/Low-Adopt; red dashed-dot line: Low-Exp/High-Adopt; black long-dashed line: High-Exp/Low-Adopt; yellow solid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

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 11, 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.

4.3 Wage Growth and Hours Worked

This section discusses effects of AI exposure and adoption on wage growth and hours worked, as these indicators also measure labor demand.[19] We focus on reduced-form evidence to document how these outcomes vary across workers with different levels of AI exposure and across industries with different levels of AI adoption.

Wage growth 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.[20]

Consider a regression model for the wage growth:

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

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 occupation at t, age_{it} and age_{it}^2 are the individual's age and its square, and X_{it} includes gender and education controls. Notation λ_t and δ_i capture time and industry fixed effects, respectively.[21]

Columns (1) and (2) of Table 2 report the regression results. Both AI exposure and adoption are associated with stronger wage growth prior to the LLM era, suggesting a wage premium for workers in high-exposure occupations and high-adoption industries. In the post-LLM period, however, these returns attenuated significantly. For AI exposure, the reversal is precisely esti-

[19] Effects of AI intensity on the level of wages are reported in Section D.3.

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

[21] 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.

mated regardless of specification. Workers in more AI-exposed occupations 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. Meanwhile, for AI adoption, the post-LLM coefficient is still significantly negative without industry fixed effects, but loses significance once industry fixed effects are included, suggesting that the wage slowdown in high-adoption industries may reflect broader industry-level trends rather than the direct effect of adoption itself.

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.[22]

Overall, AI exposure is positively and statistically significantly associated with hours worked (line 1), indicating that workers in more AI-exposed occupations tend to 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

[22] 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 2: DETERMINANTS OF WAGE GROWTH

	(1) Time FE	(2) Time+Ind FE	(3) JtJ Baseline	(4) JtJ Post-LLM
AI Exposure (std.)	0.795***	0.795***	0.540***	0.662***
	(0.0969)	(0.105)	(0.0980)	(0.108)
AI Adoption (std.)	2.484***	1.977***	0.652**	1.499**
	(0.674)	(0.643)	(0.324)	(0.613)
AI Exposure × Post-LLM	-0.564***	-0.568***		-0.594***
	(0.197)	(0.197)		(0.203)
AI Adoption × Post-LLM	-1.388*	-1.106		-0.949
	(0.775)	(0.744)		(0.722)
AI Exposure × Job-to-job			2.244***	1.226
			(0.692)	(0.834)
AI Adoption × Job-to-job			1.828	16.58*
			(2.871)	(9.061)
AI Exposure × Job-to-job × Post-LLM				2.911*
				(1.719)
AI Adoption × Job-to-job × Post-LLM				-18.35*
				(9.573)
Market Tightness (θ_t)	0.0179***	0.0155**	0.0186***	0.0182***
	(0.00622)	(0.00625)	(0.00664)	(0.00664)
Age	-0.268***	-0.409***	-0.393***	-0.391***
	(0.0313)	(0.0329)	(0.0343)	(0.0343)
Age^2	0.00183***	0.00322***	0.00306***	0.00305***
	(0.000355)	(0.000370)	(0.000384)	(0.000384)
Female	-1.758***	-1.309***	-1.183***	-1.177***
	(0.172)	(0.196)	(0.202)	(0.202)
High School	1.543***	1.268***	1.126***	1.117***
	(0.270)	(0.274)	(0.286)	(0.286)
Some College / Associate	1.306***	0.933***	0.865***	0.849***
	(0.276)	(0.283)	(0.295)	(0.295)
College Graduate	2.406***	1.892***	1.762***	1.741***
	(0.330)	(0.340)	(0.351)	(0.351)
Observations	303461	303461	274342	274342
R-squared	0.00514	0.00764	0.00751	0.00762

Standard errors in parentheses

Time fixed effects included in all columns.

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

Standard errors clustered at individual level.

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

Table 3: DETERMINANTS OF HOURS WORKED

	(1) Time FE	(2) Time + Industry FE
AI Exposure (std.)	0.750***	0.656***
	(0.00997)	(0.0107)
AI Adoption (std.)	-0.362***	-0.388***
	(0.0226)	(0.0228)
AI Exposure × Post-LLM	-0.136***	-0.132***
	(0.0202)	(0.0197)
AI Adoption × Post-LLM	0.236***	0.242***
	(0.0269)	(0.0268)
Market Tightness (θ_t)	0.00613***	0.00410***
	(0.000366)	(0.000360)
Age	0.979***	0.882***
	(0.00447)	(0.00441)
Age^2	-0.0107***	-0.00968***
	(0.0000515)	(0.0000507)
Female	-4.447***	-3.614***
	(0.0183)	(0.0201)
High School	2.661***	2.440***
	(0.0379)	(0.0372)
Some College / Associate	2.326***	2.214***
	(0.0394)	(0.0390)
College Graduate	3.458***	3.585***
	(0.0393)	(0.0399)
Observations	8037399	8037399
R-squared	0.0838	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.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

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 workers with high AI exposure and adoption, 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 reduced demand for workers with high AI exposure and adoption jobs.

5 Hirability and Separability: AI-driven Reallocation Pressure

This section analyzes the effects of AI exposure and adoption on job findings and job losses. 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:

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

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_{jt} and AIA^2_{jt} denote the AI adoption rate of industry j where worker i belongs (standardized) and its squared term, respectively; age_{it} and age^2_{it} 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 8).[23] We include $\theta_{i,t-1}$, the lagged vacancy-to-unemployment ratio at the occupation level, to control for occupation-specific labor market conditions. This isolates the effects of LLMs from cyclical fluctuations that may affect occupations differently.[24] 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.[25]

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 $\widehat{UE}^{wo}_{it}$, which we define as follows:

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

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 $\widehat{EU}^{wo}_{it}$.

[23] We do not consider the squared term of AI exposure due to its high correlation with the linear term, which would introduce multicollinearity. Excluding the squared term of AI adoption does not change the result materially.

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

[25] 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.

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 more favorable job-finding prospect.

The weighted cross-sectional standard deviation is

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

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. AI adoption reduces job-finding rates and raises job-loss rates, consistent with reduced job creation and greater job destruction

Table 4: DETERMINANTS OF HIRABILITY AND SEPARABILITY

	Hirability (UE)		Separability (EU)	
	(1)	(2)	(3)	(4)
	Standardized	2×2 Group	Standardized	2×2 Group
AI Exposure (std.)	-0.00871***		-0.00231***	
	(0.00295)		(0.000125)	
AI Adoption (std.)	-0.0277***		0.000400***	
	(0.00433)		(0.000146)	
AI Adoption (std.)^2	0.00191***		-0.0000196	
	(0.000453)		(0.0000128)	
Group 2 (LH)		0.0139		-0.00279***
		(0.0238)		(0.000969)
Group 3 (HL)		-0.0158**		-0.00419***
		(0.00676)		(0.000249)
Group 4 (HH)		-0.0764***		-0.00473***
		(0.00764)		(0.000246)
Market Tightness (θ_{t-1})	-0.0000228	0.00000198	-0.0000142***	-0.0000127***
	(0.000116)	(0.000116)	(0.00000216)	(0.00000216)
Age	0.000216	0.000328	-0.000771***	-0.000748***
	(0.000840)	(0.000841)	(0.0000454)	(0.0000455)
Age^2	-0.0000182*	-0.0000187**	0.00000676***	0.00000653***
	(0.00000949)	(0.00000949)	(0.000000474)	(0.000000474)
Female	0.00198	0.00116	-0.000224	-0.000429**
	(0.00502)	(0.00500)	(0.000209)	(0.000210)
High School	-0.0154*	-0.0157**	-0.00407***	-0.00430***
	(0.00788)	(0.00787)	(0.000589)	(0.000588)
Some College / Associate	0.00918	0.00909	-0.00657***	-0.00687***
	(0.00861)	(0.00855)	(0.000579)	(0.000578)
College Graduate	-0.00198	0.00185	-0.00844***	-0.00861***
	(0.00904)	(0.00902)	(0.000564)	(0.000562)
Observations	49436	49436	1368019	1368019
R-squared	0.00626	0.00692	0.00296	0.00297

Standard errors in parentheses

Constant and MIS fixed effects not reported.

Standard errors clustered at individual level.

Sample: 2023–2026.

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

(Line 2, Columns 1 and 3). In the UE regression, the squared term, however, enters with the opposite sign and a smaller magnitude: at sufficiently high levels of adoption, the negative effect on job-finding reverses, suggesting that the job-destruction effects of LLMs diminish at the intensive margin.[26] Meanwhile, the coefficient of the squared term is not statistically significant for the EU regression.

We next examine how the interaction between exposure and adoption affects labor market transitions. The negative effect on job-finding rates is largest for the HH group and second largest for the HL group (Column 2, Table 4). On the separation side, the EU coefficients for the LH, HL, and HH groups are all negative and statistically significant, indicating lower job destruction relative to the LL group across all three. Taken together, the interaction between AI exposure and adoption exerts a statistically significant drag on hiring.[27]

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 Dispersions

Figure 13 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), while the lower exposure groups remain relatively stable.

Panel B displays the mean hirability index by AI adoption quartile. 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.

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

[27] A similar pattern emerges when the four exposure groups are interacted with quartiles of AI adoption.

Figure 13: HIRABILITY INDEX BY AI EXPOSURE AND ADOPTION

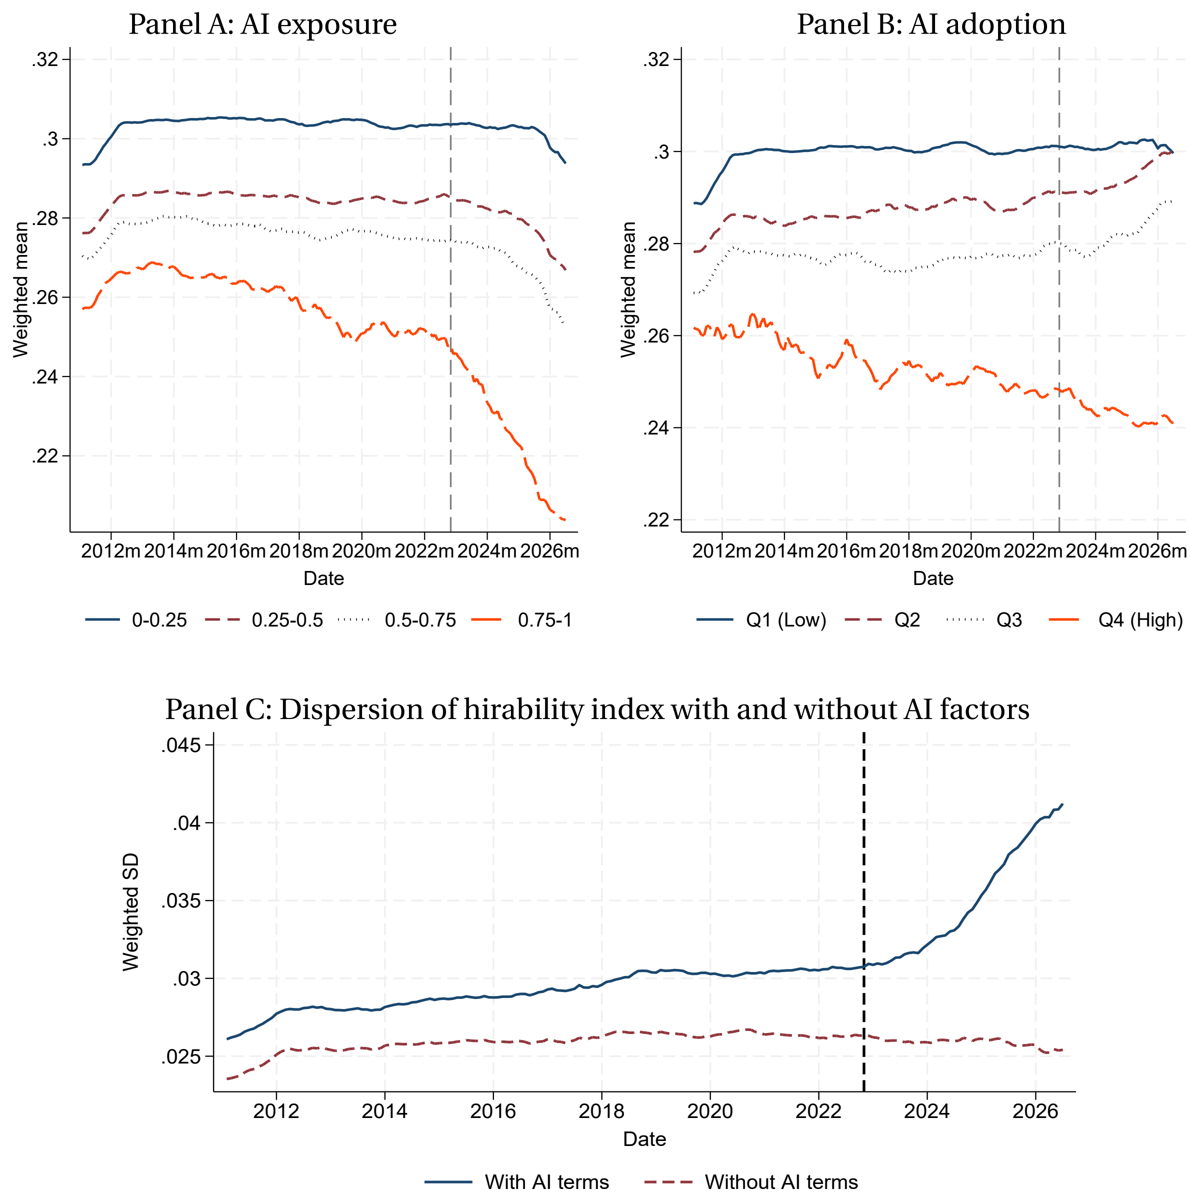

Note: Panel A displays the mean hirability index by AI exposure score. Panel B displays the mean hirability index by the quartile of AI adoption (thresholds measured in 2025). Panel C plots the dispersion of hirability index with AI exposure and adoption (solid blue line) and that without AI exposure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

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 pressure following the introduction of large language models. By contrast, the dispersion of the

Figure 14: Separability index by the quartile of AI exposure and AI adoption

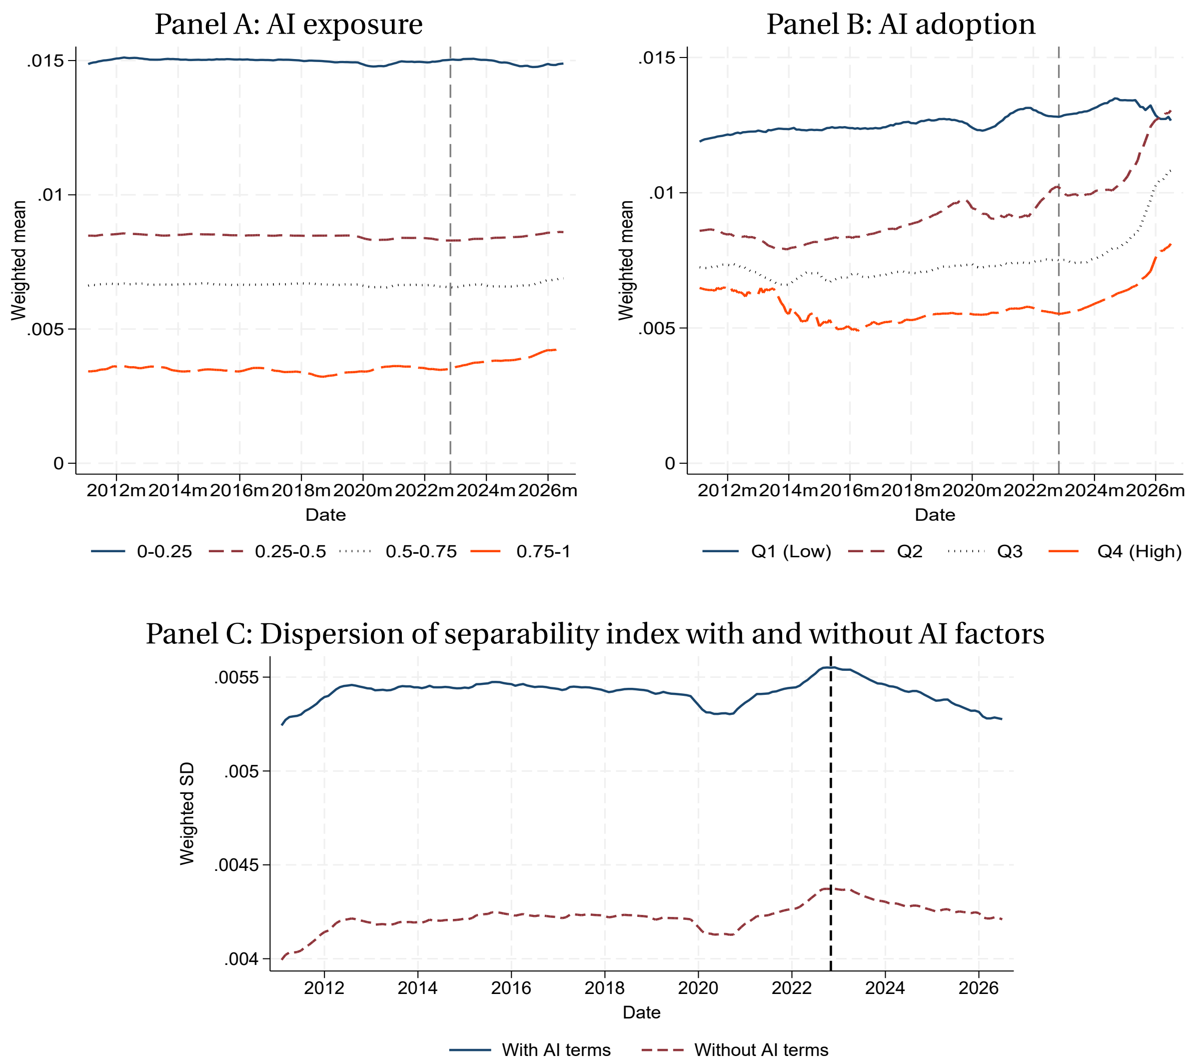

Note: Panel A displays the mean separability index by AI exposure score. Panel B displays the mean separability index by the quartile of AI adoption (thresholds measured in 2025). Panel C plots the dispersion of hirability index with AI exposure and adoption (solid blue line) and that without AI exposure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

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 14 plots the separability index ($sd_t^{w}(EU)$) by 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 14 plots the dispersion of the separability index with and without AI factors. In contrast to the hirability index, the dispersion of the separability index remains largely flat in the post-LLM era—a pattern that holds both with and without the AI factors included. In sum, AI-driven reallocation pressures operate primarily through the hiring margin rather than separations.

5.4 Identification and Confounding Factors

The analyses so far are reduced-form and do not address the causal effects of LLMs on labor market outcomes. A key concern is that the estimated effects of AI exposure and adoption may reflect confounding sector-level factors—such as differential responses to monetary tightening, sector-specific demand shocks, or the post-pandemic recovery—rather than the direct impact of AI on labor market transitions. We address this with three complementary identification strategies, each targeting a distinct threat to causal interpretation. The baseline findings prove robust across all three.[28] We also examine potential confounders, including the unwinding of the tech-sector hiring surge and the contraction of telework.

A. Alternative Identification Strategies We consider three additional analyses to address the endogeneity concerns.

The first is to include industry×month fixed effects in the regressions for UE and EU transitions (equation (C.1)), which absorb all industry-time variation — the Fed tightening cycle, sector-specific demand shocks, remote-work recovery — and identify the interaction $AIX_{it} \times AIA_{jt}$ purely off within-industry, cross-occupation variation. Note that we include the lagged occupation-level market tightness ($\theta_{i,t-1}$) to absorb occupation-level labor-market fluctuations. The interaction coefficient for the UE transitions is negative with statistical significance, and is positive with statistical significance for the EU transitions. High AI-exposure workers in higher-adoption industries are significantly less likely to return to employment when unemployed, and also face elevated separation risks. Pre-trend tests on 2019 and 2021–2022 subsamples yield interaction coefficients that are statistically insignificant, supporting the parallel-trends assumption underlying this design (Appendix Table A.2; Appendix Table A.3).

The second strategy replaces industry×month fixed effects with 2-digit industry fixed effects interacted with the federal funds rate, thereby absorbing differential monetary-policy exposure

[28] We summarize the results below; full details appear in Appendix Section C.

across broad sectors while preserving the cross-industry variation in AIA_{jt} needed for identification. The AIA_{jt} coefficient for UE transitions is negative with statistical significance, consistent with AI-adopting industries reducing labor demand at the extensive margin and slowing re-employment for unemployed workers. The AIX_{it}'s effect remains statistically insignificant, suggesting that occupational AI exposure alone does not predict job-finding in the absence of actual adoption (Appendix Table A.4).

The third analysis instruments AIA_{jt} with a leave-one-out shift-share whose shares are pre-LLM industry task content and whose shift is the national AI adoption trend, excluding each industry in turn. We call this instrument *"AI adoption-potential IV."* The IV estimate for UE transitions is more negative than the baseline OLS estimate and statistically significant, perhaps suggesting that industries facing weaker demand may be more likely to adopt LLMs, thereby compounding the reduction in hiring. Heterogeneity analysis by AIX quartile reveals that the adoption effect on job-finding is concentrated among workers in the top two exposure bins (0.5-0.75 and 0.75-1), while the lower exposure groups are largely insulated. For EU transitions, the OLS adoption coefficient is positive but becomes insignificant under instrumentation, suggesting that industries facing weaker demand shed workers and are also more likely to adopt LLMs.

To address measurement error in AIX_{it}, we instrument it with the Felten et al. AIOE score. The results remain robust (Appendix Tables A.5 and A.6).

The patterns of hirability and separability dispersion remain robust across all three analyses. In particular, the magnitudes of the AI-attributable dispersion are similar to the baseline: the weighted standard deviation of hirability rose by about 1.1 percentage points in the baseline, largely predicted by the AI factors, compared with 0.6–0.7 percentage points under the two fixed-effects specifications. The adoption-potential IV yields a larger increase—roughly four times the baseline magnitude. In this sense, we assess that our baseline estimates are robust to the endogeneity concerns and risks associated with the estimated hirability dispersion are well balanced.

B. Tech Sector's Employment Adjustment to Their Increased Hiring in the Post-COVID Period A first concern is that the estimated effect of AI adoption on hiring reflects tech-sector's unwinding of their over hiring in 2022 rather than broad-based AI adoption. To address this, we re-estimate the baseline and IV models excluding workers in Computer and Data Processing Services (`ind1990 = 732`), the industry at the center of the post-pandemic tech layoff cycle, from both the estimation sample and the instrument construction. If the baseline result were driven by a tech-specific hiring adjustment, the coefficient on AIA_{jt} should attenuate once this sector is removed. The results point in the opposite direction, robustifying our main claims. The IV estimate for UE transitions become even larger than the baseline's. The EU coefficient remains statistically insignificant. The robust effect outside the tech sector supports an interpretation based on economy-wide AI adoption rather than a confound specific to the 2022–2023 technology-sector cycle.[29]

C. Remote work Another candidate confounder is telework. Lambert and Yannick (2026) argue that remote work was already concentrated in high-AI occupations, so AI exposure may partly capture teleworkability. We confirm that workers with higher AI exposure are more likely to work remotely. However, teleworkability has opposite effects on labor market transitions relative to AI exposure and adoption: it raises job-finding probabilities, whereas AI exposure and adoption lower them (Appendix Section D.1). The adverse hiring effects we document are therefore unlikely to reflect the prevalence of remote work in the post-LLM era.

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.[30] 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 flows model of the unemployment rate (Shimer, 2012).

[29] Full results are reported in Appendix Table **??**.

[30] 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).

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 idea, Şahin et al. (2014) demonstrate 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 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

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

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

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

Equation (6.6) provides a micro-founded link between the dispersion of individual job-finding prospects and structural unemployment.

To assess the extent to which LLMs have changed the structural unemployment rate ($\triangle SU_t^{AI}$), we use the following equation derived from (6.6) after reintroducing time subscript:

$$\triangle SU_t^{AI} \approx \frac{U_{t_0}}{\alpha f_{t_0}\sqrt{2\pi}} \triangle \left(\sqrt{\mathrm{Var}_i(f_{it})} - \sqrt{\mathrm{Var}_i(f_{it}^{wo})}\right), \tag{6.7}$$

Table 5: Implied Change in Structural Unemployment Rate (pp, U_{t_0}: 4%, f_{t_0}: 0.3)

	Matching Elasticity to Vacancies (α)	
$\triangle\left(\sqrt{\mathrm{Var}_i(f_{it})}-\sqrt{\mathrm{Var}_i(f_{it}^{wo})}\right)$	$\alpha = 0.3$	$\alpha = 0.5$
LLMs' contribution (+1.1)	0.19	0.12

where $\triangle$ denotes changes from end-2022 through May 2026, and $\sqrt{\mathrm{Var}_i(f_{it})}$ and $\sqrt{\mathrm{Var}_i(f_{it}^{wo})}$ are proxied by sd_t^w(UE) and sd_t^{wo}(UE) as displayed in Figure 13. We set the baseline unemployment rate (U_{t_0}) to 4% and the baseline job-finding rate to 0.3. We consider two values of the matching elasticity to vacancies (α).

According to Figure 13, AI factors account for +0.5 pp in the cross-sectional standard deviation of hirability, with an additional widening of +1.1 pp after 2022:M11 reflecting LLMs' effects, for a total AI-driven dispersion of +1.6 pp (difference between the solid blue line and the dashed red line). The implied changes in the structural unemployment rate are 0.19 pp and 0.12 pp when $\alpha = 0.3$ and $\alpha = 0.5$, respectively.[31] We therefore evaluate that LLMs likely boosted structural unemployment rate by about 0.1–0.2 percentage point.[32]

We statistically estimate the natural rate of unemployment for the four groups defined jointly by AI exposure and adoption. We use the concept of trend U-star—unemployment rate composed of only the trend components of job-finding and job-loss rates—as a proxy for the natural rate, as identifying the natural rate with inflation is challenging in general and particularly during the sample period of our interest.[33] To do so, we recover the trend components of each group's unemployment flows rates using a Bayesian trend-cycle decomposition, then feed these trends into the two-state model of unemployment rate (e.g., Shimer, 2012) to recover each group's natural rate. We then aggregate these estimates using labor-force shares to obtain the aggregate natural rate. This framework accounts 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.

[31] Our preferred baseline is $\alpha = 0.3$ as shown in recent empirical studies (e.g., Lange and Papageorgiou, 2020; Figura and Waller, 2024). The estimated structural unemployment rates are naturally larger under lower matching elasticity, as weaker substitutability between worker types amplifies the unemployment consequences of mismatch.

[32] The fixed-effects specifications yield slightly smaller estimates—0.09 pp (3-digit industry × time) and 0.12 pp (2-digit industry × federal funds rate)—but remain close to the baseline. The adoption-potential IV implies a larger increase in the natural rate of approximately half a percentage point.

[33] This approach has been adopted by Tasci (2012); Crump et al. (2019); Ahn (2023).

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:

$$y^j_{g,t} \; = \; \tau^j_{g,t} + c^j_{g,t} + \varepsilon^j_{g,t}, \qquad \varepsilon^j_{g,t} \sim \mathcal{N}\left(0, \sigma^{2,j}_{\varepsilon,g}\right). \tag{6.8}$$

The sample period is 2013:M1-2026:M7. Observations during the COVID-19 swing (2020:M3–2021:M12) are treated as missing and excluded from likelihood evaluation.[34]

6.1.1 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:

$$\tau^{\text{EU}}_{g,t} = \tau^{\text{EU}}_{g,t-1} + \beta^{\text{EU}}_{g,t-1}, \tag{6.9}$$

$$\beta^{\text{EU}}_{g,t} = \beta^{\text{EU}}_{g,t-1} + \zeta^{\text{EU}}_{g,t}, \qquad \zeta^{\text{EU}}_{g,t} \sim \mathcal{N}\left(0, \sigma^{2,\text{EU}}_{\zeta,g}\right). \tag{6.10}$$

UE rate:

$$\tau^{\text{UE}}_{g,t} = \tau^{\text{UE}}_{g,t-1} + \beta^{\text{UE}}_{g,t-1}, \tag{6.11}$$

$$\beta^{\text{UE}}_{g,t} = \beta^{\text{UE}}_{g,t-1} + \zeta^{\text{UE}}_{g,t}, \qquad \zeta^{\text{UE}}_{g,t} \sim \mathcal{N}\left(0, \sigma^{2,\text{UE}}_{\zeta,g}\right). \tag{6.12}$$

6.1.2 Cycle

For the cyclical component, we use Harvey (1985)'s trigonometric stochastic cycle.[35]

[34] 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. We consider the sample period starting from 2013, the starting point 2012 is the period where the economy was still recovering from the Great Recession. The results are overall robust.

[35] 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

Defining the auxiliary variable $c_{g,t}^{j*}$, the pair $(c_{g,t}^{j}, c_{g,t}^{j*})$ evolves as

$$\begin{pmatrix} c_{g,t}^{j} \\ c_{g,t}^{j*} \end{pmatrix} = \rho \begin{pmatrix} \cos\lambda & \sin\lambda \\ -\sin\lambda & \cos\lambda \end{pmatrix} \begin{pmatrix} c_{g,t-1}^{j} \\ c_{g,t-1}^{j*} \end{pmatrix} + \begin{pmatrix} \kappa_{g,t}^{j} \\ \kappa_{g,t}^{j*} \end{pmatrix}, \tag{6.13}$$

where $\lambda = 2\pi/p$ is the fundamental frequency and $\kappa_{g,t}^{j}, \kappa_{g,t}^{j*} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_{\kappa,g}^{2,j})$. 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.[36]

6.2 Natural Rate of Unemployment

6.2.1 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:

$$u_{g,t}^{*} = 100 \frac{\tau_{g,t}^{\mathrm{EU}}}{\tau_{g,t}^{\mathrm{EU}} + \tau_{g,t}^{\mathrm{UE}}}, \tag{6.14}$$

where $\tau_{g,t}^{\mathrm{EU}}$ and $\tau_{g,t}^{\mathrm{UE}}$ are the trend components of the EU and UE flow rates for group g. Equation (6.14) is evaluated draw-by-draw so that posterior uncertainty in both trend components propagates into uncertainty about $u_{g,t}^{*}$.[37]

---

the dynamics of the cyclical component.

[36] The state-space representation is found in Appendix Section F.

[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:M7). 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.

6.2.2 Aggregate Estimate

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

$$u^*_{\text{agg},t} = \sum_{g=1}^{4} \omega_{g,t}\, u^*_{g,t}. \tag{6.15}$$

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

6.2.3 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 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.

6.3 Estimation Results

Figure 15 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.

[38] See Section F.4 for the background estimation results.

Figure 15: TREND UNEMPLOYMENT FLOWS AND U-STAR SINCE LLMS

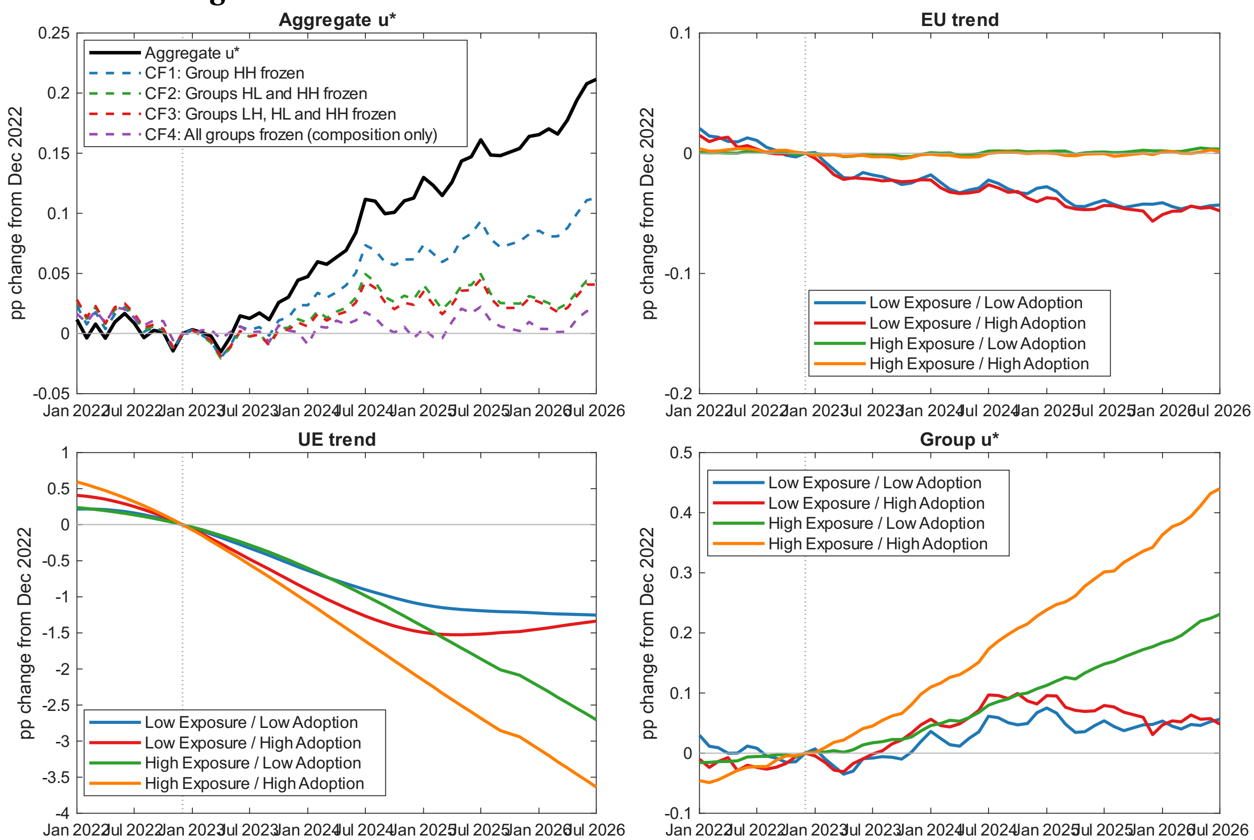

Notes: The upper-left panel displays the aggregate natural rate and the counterfactual series; the upper-right and lower-left panels show the EU and UE trend rates for the four groups, respectively; and the lower-right panel displays the estimated natural rate of each group. **Source**: Authors' calculation.

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 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 0.1 percentage point of the rise in the natural rate, while the HL group contributes approximately 5 basis points; together, high-exposure workers raise the natural rate by about 0.15 percentage point. This upward pressure stems primarily from faster declines in the trend job-finding rate among the HH group (orange line, bottom-left panel) and the HL group (green line, bottom-left panel) relative to the other groups. The LL group contributes roughly 5 basis points, likely reflecting modest upward pressure from lower matching efficiency among recent immigrants or other supply-side factors during this period.[39] The LH group's contribution is negligible, largely owing to its small labor force share. Notably, the natural rate has been on a secular downtrend driven by population aging and rising educational attainment (**??**). Against this backdrop, the concentration of upward pressure in the HH and HL groups in the post-LLM era points to unusual forces at work—forces that our empirical evidence consistently attributes to LLMs and AI-related factors.

In sum, this result underscores the importance of the hiring margin in the upward pressure on the natural rate driven by LLMs. The contribution of the higher-exposure groups aligns with the increase in structural unemployment implied by the theoretical model. Nonetheless, the estimates are subject to considerable uncertainty: the 68% posterior intervals span about 1 percentage point, well exceeding the estimated LLM effects (Figure A.8). We therefore interpret our results as indicating upside risks to the natural rate of unemployment associated with LLMs.

[39] Lower-skilled immigrants tend to face less efficient job searches during the labor market assimilation process.

7 Pervasiveness of LLMs' Effects on the Labor Market

The analyses so far suggest that the diffusion of LLMs may have pervasive effects on the aggregate labor market, raising both the structural unemployment rate and the natural rate of unemployment. While the estimated magnitudes are modest and surrounded by considerable uncertainty, the empirical patterns point to a broad reorientation of labor demand that extends across multiple margins of the labor market.

A growing body of microeconomic evidence documents that the adverse effects of LLMs have fallen disproportionately on early-career workers, young workers, and new labor force entrants (Brynjolfsson et al., 2025; Hosseini and Lichtinger, 2025; Tucker, 2026). Our findings are broadly consistent with this picture: as shown in Figure 16, the composition of unemployment among new labor force entrants shifted toward workers in high-exposure/high-adoption groups, particularly among those aged 25–54 with some college or higher education (Panels C and D). This compositional shift was modest among the youngest entrants (aged 20–24, Panel B), suggesting that the effects on this group have been partially obscured by the influx of new graduates into the labor force.

However, our evidence indicates that the disruption associated with LLMs reaches well beyond young workers and new entrants. Several findings point to broader pervasiveness. First, the geographic distribution of AI exposure and adoption reveals striking co-location: as shown in Figure 18, metro areas with high AI exposure tend to have high AI adoption, so that high-intensity markets are doubly exposed to the reallocation pressures documented above. This concentration creates both risk — if adjustment is slow, local labor markets in high-intensity metros may absorb disproportionate and persistent unemployment — and opportunity, since the same agglomeration of AI-intensive workers and firms could facilitate faster reallocation and productivity gains once uncertainty recedes.

Second, the composition of unemployment among permanent job losers shifted toward high-exposure/high-adoption workers not only among prime-age workers (aged 25–54) but also among workers aged 55–64, and the shift was more pronounced among college graduates aged 55–64 (Figure 17). This finding implies that LLMs may be affecting labor supply decisions near retirement, as elevated job-loss risk among older high-exposure workers could pull forward

Figure 16: Composition of New Entrants to the Labor Force (Blue: 2022; Red: 2025~)

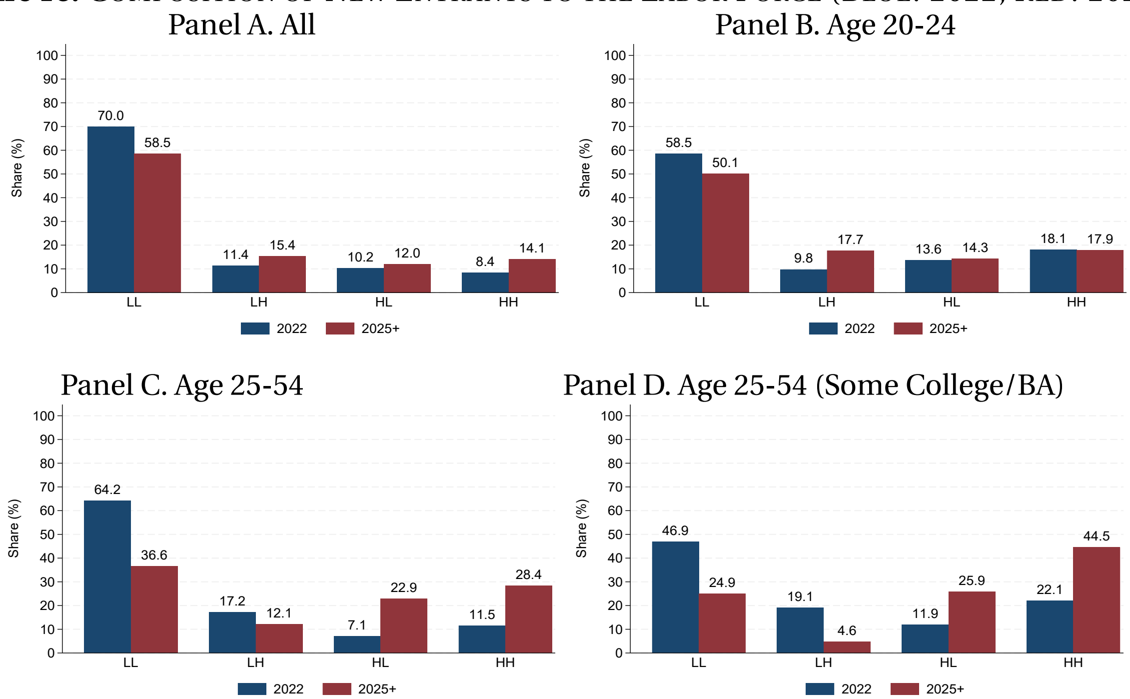

Notes: Bars display the share of each AI exposure/adoption group within new entrants (Panel A), new entrants aged 20–24 (Panel B), new entrants aged 25–54 (Panel C), and new entrants aged 25–54 with education beyond high school (Panel D). Blue bars denote 2022 shares and red bars denote 2025–2026 shares.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

retirement or discourage re-entry into the labor force, with lasting consequences for participation rates and aggregate labor supply.

Taken together, although the estimated effects of LLMs on aggregate labor market functioning remain small in magnitude and carry significant uncertainty, the empirical evidence suggests that their reach is broad, affecting hiring, separation, participation, and geographic reallocation across multiple corners of the labor market. The co-occurrence of these patterns across age groups, reasons for unemployment, and local labor markets points to a structural reconfiguration rather than a cyclical or narrowly targeted disruption.

8 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

Figure 17: COMPOSITION OF OTHER SEPARATIONS (BLUE: 2022; RED: 2025˜)

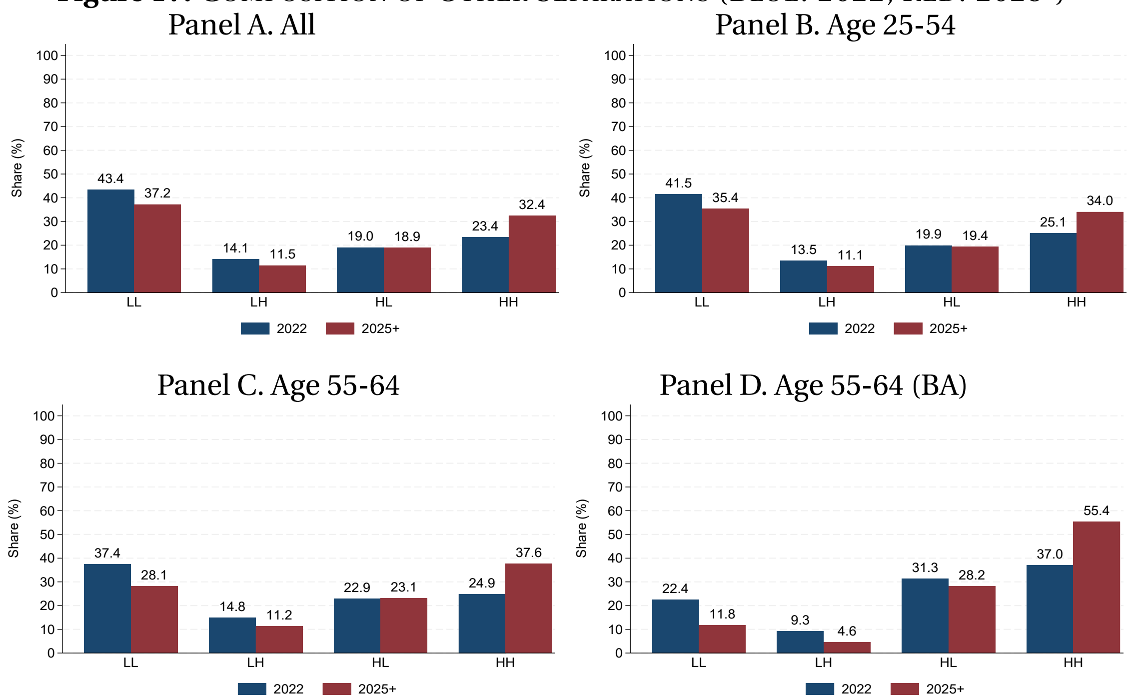

Notes: Bars display the share of each AI exposure/adoption group within other separations (Panel A), other separations aged 25–54 (Panel B), other separations aged 55–64 (Panel C), and other separations aged 55–64 with a college degree (Panel D). Blue bars denote 2022 shares and red bars denote 2025–2026 shares. **Source**: Authors' calculation.

Figure 18: AI EXPOSURE AND ADOPTION BY METRO AREA

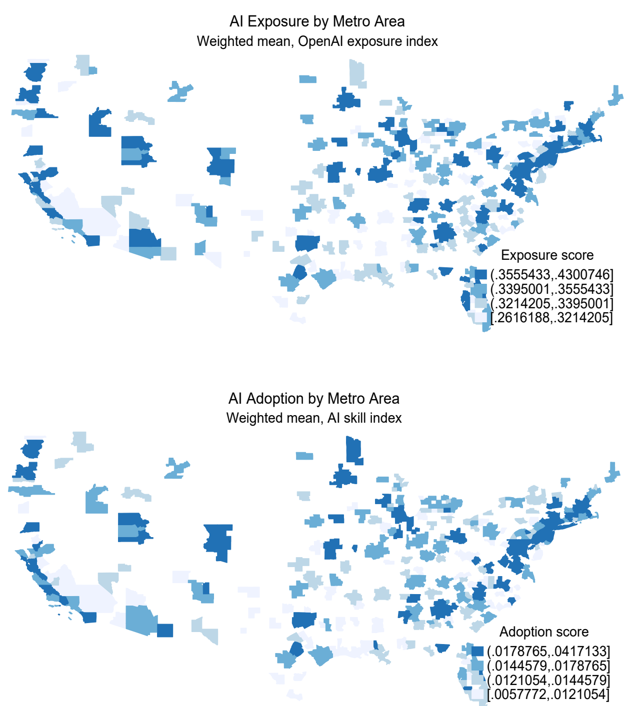

Notes: This figure displays AI exposure and adoption intensities across U.S. metropolitan areas. Darker shades indicate higher quartiles of AI intensity. **Source**: Authors' calculations.

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 AI-driven rise in the dispersion of job-finding propensities suggests growing disparities in reemployment prospects, which our empirical framework translates into modest upward pressure on the natural rate since late 2022. Since weak hiring of high-exposure/high-adoption workers may reflect uncertainty about both the macroeconomic environment and the capabilities of LLMs, policies that reduce uncertainty may help alleviate reallocation pressure and facilitate more efficient labor market adjustment. To the extent that the low-hiring, low-firing environment is itself an endogenous response to uncertainty, recent signs of recovery in job-finding and job-switching rates among HH workers, along with rising market tightness, may signal reduced uncertainty about hiring workers with high AI intensity.

References

Abel, Jaison R., Richard Deitz, and Natalia Emanuel. 2025. "Survey Evidence on Firms' Use of Artificial Intelligence." Federal Reserve Bank of New York, unpublished survey evidence.

Abraham, Katharine G., and Lawrence F. Katz. 1986. "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?" *Journal of Political Economy* 94 (3): 507–522, `http://www.jstor.org/stable/1833046`.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. "Artificial Intelligence and Jobs: Evidence from Online Vacancies." *Journal of Labor Economics* 40 (S1): 293–340.

Acemoglu, Daron, and Pascual Restrepo. 2019. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33 (2): 3–30. 10.1257/jep.33.2.3.

Ahn, Hie Joo. 2023. "Duration Structure of Unemployment Hazards and the Trend Unemployment Rate." *Journal of Economic Dynamics and Control* 151 104670. 10.1016/j.jedc.2023.104670.

Ahn, Hie Joo, and James D. Hamilton. 2022. "Measuring labor-force participation and the incidence and duration of unemployment." *Review of Economic Dynamics* 44 1–32. https://doi.org/10.1016/j.red.2021.04.005.

Althoff, Lukas, and Hugo Reichardt. 2026. "Task-Specific Technical Change and Comparative Advantage." June, `https://hugoreichardt.com/pdf/tstc_compadvantage.pdf`, Working Paper.

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

Baslandze, Salome, Erik Brynjolfsson, Daniel Rock, and Chad Syverson. 2026. "Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives."Technical Report 34984, National Bureau of Economic Research, `https://www.nber.org/papers/w34984`.

Bonney, Kathryn, Cory L Breaux, Emin Dinlersoz, Lucia S. Foster, John C Haltiwanger, and Aditya A Pande. 2026. "The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks." NBER Working Paper No. 35141.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine?: Six Facts about the Recent Employment Effects of Artificial Intelligence." Working Paper.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. "What Can Machines Learn, and What Does it Mean for Occupations and the Economy?" *AEA Papers and Proceedings* 108 43–47.

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

Crane, Leland D., and Paul E. Soto. 2026. "AI and Coder Employment: Compiling the Evidence." Finance and Economics Discussion Series (FEDS): https://doi.org/10.17016/FEDS.2026.018.

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

Daly, Mary C., Bart Hobijn, AysegÃ¼l Sahin, and Robert G. Valletta. 2012. "A Search and Matching Approach to Labor Markets: Did the Natural Rate of Unemployment Rise?" *Journal of Economic Perspectives* 26 (3): 3–26. 10.1257/jep.26.3.3.

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

Eckhardt, Sarah, and Nathan Goldschlag. 2025. "AI and Jobs: The Final Word (Until the Next One)."Technical report, Economic Innovation Group, https://eig.org/ai-and-jobs-the-final-word/.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. "GPTs and GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." arXiv:2303.10130.

Felten, Edward W., Manav Raj, and Robert Seamans. 2023. "Occupational Heterogeneity in Exposure to Generative AI." Working Paper.

Figura, Andrew, and Chris Waller. 2024. "What does the beveridge curve tell us about the likelihood of soft landings?" *Journal of Economic Dynamics and Control* 169 104957. https://doi.org/10.1016/j.jedc.2024.104957, Heterogeneity and Macroeconomics of Labor Markets.

Freund, Lukas B., and Lukas F. Mann. 2026. "Job Transformation, Specialization, and the Labor Market Effects of AI." Institute Working Paper 114, Federal Reserve Bank of Minneapolis and CESifo, https://doi.org/10.21034/iwp.114.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller. 2025. "Artificial Intelligence and the Labor Market." *NBER Working Paper No. 33509*.

Handa, Kunal, Alex Tamkin, Miles McCain et al. 2025. "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations." https://arxiv.org/abs/2503.04761.

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

Hosseini, Seyed M., and Guy Lichtinger. 2025. "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data." Working Paper.

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

Lambert, Peter J., and Schindler Yannick. 2026. "The Broken Ladder: AI, Remote Work, and Early-Career Hiring." http://dx.doi.org/10.2139/ssrn.6787638.

Lange, Fabian, and Theodore Papageorgiou. 2020. "Beyond Cobb-Douglas: Flexibly Estimating Matching Functions with Unobserved Matching Efficiency." Working Paper 26972, National Bureau of Economic Research. 10.3386/w26972.

Lilien, David M. 1982. "Sectoral Shifts and Cyclical Unemployment." *Journal of Political Economy* 90 (4): 777–793, http://www.jstor.org/stable/1831352.

Manning, Sam J., and TomÃ¡s Aguirre. 2026. "How Adaptable Are American Workers to AI-Induced Job Displacement?" NBER Working Papers 34705, National Bureau of Economic Research, Inc. None.

Massenkoff, Maxim, and Peter McCrory. 2026. "Labor Market Impacts of AI: A New Measure and Early Evidence." Anthropic.

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

Ng, Serena. 2021. "Modeling Macroeconomic Variations After COVID-19." Working Paper 29060, National Bureau of Economic Research. 10.3386/w29060.

Ozkan, Serdar, and Nicholas Sullivan. 2025. "Is AI Contributing to Rising Unemployment? Evidence from Occupational Variation."Technical report, Federal Reserve Bank of St. Louis, https://www.stlouisfed.org/on-the-economy/2025/aug/is-ai-contributing-unemployment-evidence-occupational-variation.

Rissman, Ellen R. 1993. "Wage growth and sectoral shifts : Phillips curve redux." *Journal of Monetary Economics* 31 (3): 395–416. None.

Shimer, Robert. 2012. "Reassessing the ins and outs of unemployment." *Review of Economic Dynamics* 15 (2): 127–148. https://doi.org/10.1016/j.red.2012.02.001.

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

The Budget Lab at Yale. 2025a. "Artificial Intelligence and Labor Market Outcomes." https://budgetlab.yale.edu, Report finding no relationship between AI exposure and recent employment or unemployment changes.

The Budget Lab at Yale. 2025b. "Evaluating the Impact of AI on the Labor Market: Occupational Composition and Employment Trends." https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market, Finds that the occupational mix of the U.S. labor market is changing somewhat faster than historically.

Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. 2025. "Working with AI: Measuring the Occupational Implications of Generative AI." Working Paper.

Tucker, Lee C. 2026. “You’re (not) Hired: Artificial Intelligence and Early Career Hiring in the Quarterly Workforce Indicators.” Census Bureau Working Paper CES-26-27.

Webb, Michael. 2020. “The Impact of Artificial Intelligence on the Labor Market.” Working Paper.

Yin, Michelle, Hoa Vu, and Claudia Persico. 2026. “How (un)Stable Are LLM Occupational Exposure Scores? Evidence from Multi-Model Replication.” NBER Working Paper No. 35110.

Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2014. “Mismatch Unemployment.” *American Economic Review* 104 (11): 3529–3564. 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 **??**), and joint categorization of AI exposure and adoption (Table **??**).

B Additional Data by AI Exposure and Adoption

Figures A.2 and Figures A.3 display the fully-imputed data by AI exposure and the joint classification, respectively.

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

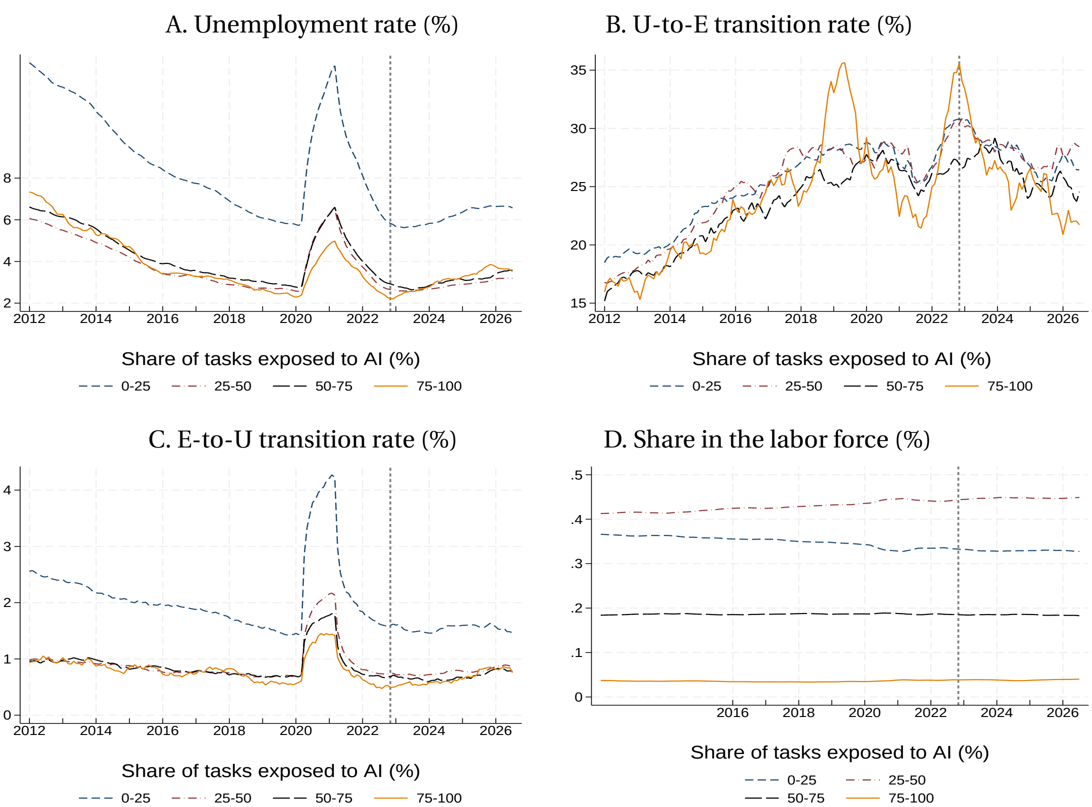

Note: This figure displays the unemployment rate (Panel A), UE rate (Panel B), EU rate (Panel C), and labor force shares (Panel D) for four AI exposure groups with full imputation (12-month moving average). Blue dashed line: 0-0.25; red dashed-dot line: 0.25-0.5; black long-dashed line: 0.5-0.75; solid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

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

Figure A.2: Unemployment rate, UE and EU rates and Labor Force Shares by AI adoption (full imputation)

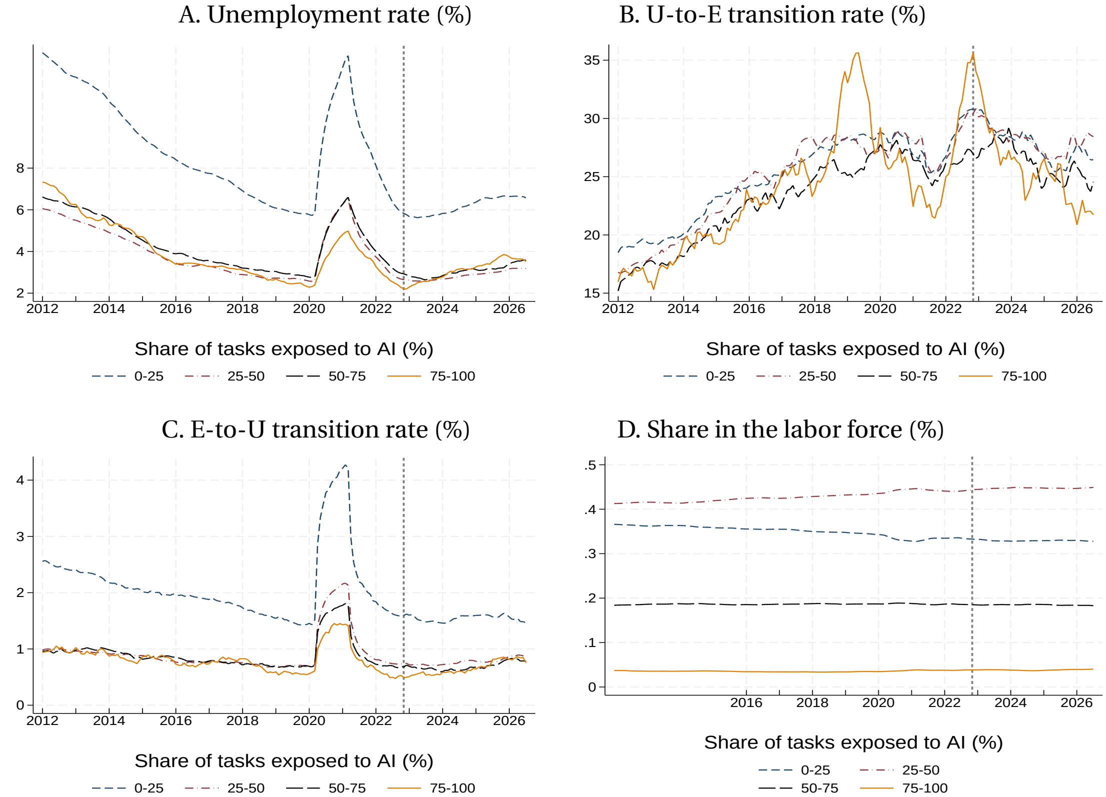

Note: This figure displays the unemployment rate (Panel A), UE rate (Panel B), EU rate (Panel C), and labor force shares (Panel D) for four AI exposure groups with full imputation (12-month moving average). Blue dashed line: 0-0.25; red dashed-dot line: 0.25-0.5; black long-dashed line: 0.5-0.75; solid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

Figure A.3: UNEMPLOYMENT RATE, UE AND EU RATES AND LABOR FORCE SHARES BY AI EXPOSURE/ADOPTION (FULL IMPUTATION)

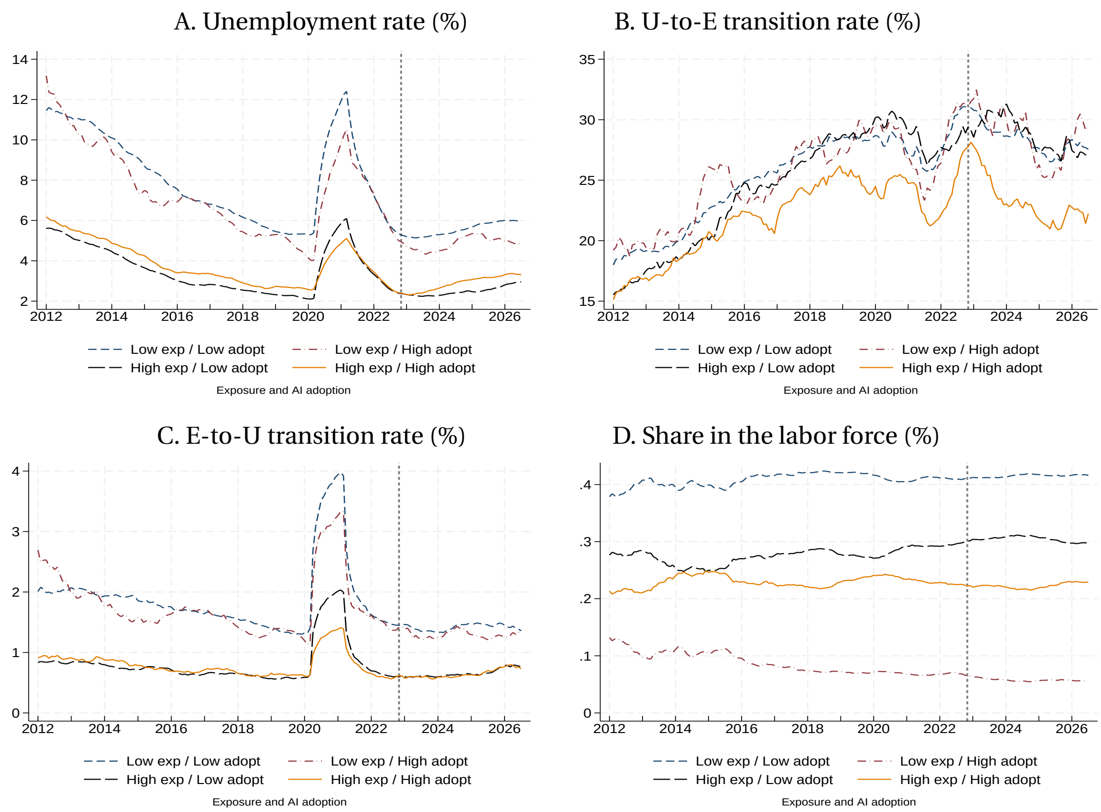

Note: This figure displays the unemployment rate (Panel A), UE rate (Panel B), EU rate (Panel C), and labor force shares (Panel D) for four AI exposure/adoption groups (12-month moving average). "Exp" stands for exposure, and "adopt" stands for adoption. Blue dashed line: Low-exp/Low-adopt; red dashed-dot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

C Identification of AI's Influence

C.1 Industry×Month Fixed Effects

Identification. The industry-by-month fixed effects specification is the most defensible approach to identifying AI's effect on labor market transitions. It directly absorbs all industry-time variation — the Fed tightening cycle, sector-specific demand shocks, remote-work recovery — leaving only within-industry, cross-occupation variation in exposure to identify the key interaction $AIX_{it} \times AIA_{jt}$. Workers in the same adopting industry and month are compared, so the coefficient on the interaction captures whether those more exposed to AI are differentially more (or less) likely to find or lose jobs, purged of any confound that operates at the industry-month level. The event-study evidence presented below is the visual counterpart of the same identification argument: it shows that high-exposure workers diverged from low-exposure workers sharply after end-2022 and not before, making a cyclical attribution implausible. The regression estimates and the event-study figures together constitute the paper's primary evidence.

Empirical model. For each worker i in industry j observed in month t, we estimate

$$\begin{aligned} y_{it} &= \alpha_{jt} + \beta_1 (AIX_{it} \times AIA_{jt}) + \beta_2 AIX_{it} \\ &+ \beta_3 \theta_{i,t-1} + \beta_4 age_{it} + \beta_5 age_{it}^2 + \beta_6 sex_i + \beta_7 edu_i + \beta_8 mis_{it} + \varepsilon_{it}, \end{aligned} \tag{C.1}$$

where $y_{it} \in \{UE_{it}, EU_{it}\}$ denotes either an indicator for an unemployment-to-employment (UE) transition or an employment-to-unemployment (EU) transition; α_{jt} is an industry×month fixed effect absorbed via `reghdfe`; AIX_{it} is the worker's AI exposure index (occupation-level, O*NET-based); AIA_{jt} is the industry-level AI adoption rate (Lightcast AI skill share); $\theta_{i,t-1}$ is lagged market tightness (vacancies-to-unemployment ratio) at the occupation level; and age_{it}, age_{it}^2, sex_i, edu_i, and mis_{it} denote age, its quadratic, sex, education categories, and the month-in-sample indicator, respectively. Because AIA_{jt} varies only at the industry×month level, its main effect is collinear with α_{jt} and is therefore absorbed — the interaction $AIX_{it} \times AIA_{jt}$ is the sole channel through which industry-level adoption enters the regression. Standard errors are clustered at the industry level. For UE transitions, j denotes the worker's *origin* industry, constructed to avoid the sample selection that arises from conditioning on current employment.

Table A.2: INDUSTRY × MONTH FE: AI EXPOSURE × ADOPTION INTERACTION (STANDARDISED)

	(1) UE	(2) EU
AIX × AIA (std.)	-0.00872**	0.000323**
	(0.00420)	(0.000131)
AI Exposure (std.)	-0.0000874	-0.00244***
	(0.00978)	(0.000458)
Observations	47,821	1,327,520
R-squared	0.150	0.016

Standard errors in parentheses
Sample: year > 2022 (post-ChatGPT). AIX standardised using post-2022 mean and SD.
AIX × AIA (std.): the raw AIX × AIA product, standardised as its own variable using its post-2022 sample mean and SD (not the product of the separately standardised AIX and AIA).
Fixed effects: industry × month (absorbed via reghdfe); AIA's own main effect is not separately identified (collinear with the FE).
Controls: lagged market tightness θ_{t-1}, age, age^2, sex, education, MIS FE.
UE: SE clustered at uind1990 (origin industry); EU: at ind1990 (current industry).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Results. Table A.2 reports the estimates. For job-finding (column 1), the interaction coefficient is -0.00872^{**} (s.e. 0.00420): within the same adopting industry and month, workers with higher AI exposure are significantly less likely to return to employment. The AIX_{it} main effect is near zero and insignificant, indicating that occupational exposure per se does not predict re-employment in the absence of actual industry adoption. For separations (column 2), the interaction coefficient is 0.000323^{**} (s.e. 0.000131): high-exposure workers in high-adoption industries face elevated separation rates relative to their low-exposure counterparts in the same industry. The negative AIX_{it} main effect (-0.00244^{***}, s.e. 0.000458) captures a structural feature of high-exposure occupations — absent adoption, these workers are *less* likely to separate, consistent with AI complementing their tasks. The positive interaction indicates that active industry adoption reverses this advantage, turning the complement into a substitute for the most exposed workers.

Table A.3: PRE-TREND TEST: INDUSTRY × MONTH FE (STANDARDISED)

	2019		2021–2022	
	(1)	(2)	(3)	(4)
	UE	EU	UE	EU
AIX × AIA (std.)	0.0285	0.000232	0.00731	0.000411
	(0.0218)	(0.000427)	(0.0126)	(0.000306)
AI Exposure (std.)	-0.00451	-0.00230***	0.000555	-0.00309***
	(0.00979)	(0.000512)	(0.00896)	(0.000578)
Observations	17,130	481,913	35,716	813,067
R-squared (within)	0.143	0.014	0.131	0.016

Standard errors in parentheses
Null hypothesis: the exposure x adoption interaction is zero pre-period.
Rejection would suggest the interaction reflects pre-existing differential trends, not a post-ChatGPT AI effect.
AIX and AIA standardised using the post-2022 sample mean and SD. AIX × AIA (std.) is the raw AIX × AIA product standardised as its own variable (not the product of the two separately standardised variables), same scale as Table A.2.
Fixed effects: industry × month, via reghdfe absorb(). Controls: θ_{t-1}, age, age^2, sex, education.
UE: SE clustered at uind1990 (origin industry); EU: at ind1990 (current industry).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Pre-trend evidence. The estimates in equation (C.1) have a causal interpretation only if high- and low-exposure workers within the same industry were on parallel trends in job-finding and separation rates prior to the LLM era. We test this by re-estimating equation (C.1) on two pre-LLM subsamples: 2019 (a clean pre-COVID baseline) and 2021–2022 (the post-COVID recovery period immediately before ChatGPT).

Table A.3 reports the results, using the same standardised specification (AIX × AIA standardised as its own variable, with θ_{t-1} as a control) as Table A.2. For job-finding, the interaction coefficient is 0.029 (s.e. 0.022) in 2019 and 0.007 (s.e. 0.013) in 2021–2022 — both small and statistically indistinguishable from zero, supporting parallel trends on this margin. For separations, the interaction coefficients are 0.00023 (s.e. 0.00043) in 2019 and 0.00041 (s.e. 0.00031) in 2021–2022 — also both insignificant, and smaller in magnitude than the post-2022 estimate (0.00032^{**} in Table A.2). The parallel-trends assumption is therefore not rejected on either margin in either pre-LLM window. The AIX_{it} main effect is small and insignificant for job-finding in both pre-periods, and negative and significant for separations in both (-0.0023^{***} in 2019, -0.0031^{***} in 2021–2022), consistent with the post-2022 estimate and indicative of a stable structural feature of occupations rather than a post-2022 treatment effect.

C.2 2-Digit Industry × FFR

Identification. The second specification includes 2-digit industry fixed effects interacted with the federal funds rate (FFR_t), providing a flexible control for differential responses to monetary policy shocks across broad sectors, in place of the industry×month FE of Spec 1. Because identification here does not rely on absorbing all industry-time variation, AIX_{it} and AIA_{jt} enter directly as main effects (rather than through an interaction), and we retain lagged market tightness $\theta_{i,t-1}$ as an additional cyclical control. The industry×FFR_t interaction absorbs the most plausible confound related to heterogeneous responses to monetary policy while preserving the cross-industry variation in AIA_{jt} needed to identify β_2.

Empirical model. We estimate

$$\begin{aligned} y_{it} &= \sum_k \delta_k \mathbf{1}[J_i = k] \cdot FFR_t + \beta_1 AIX_{it} + \beta_2 AIA_{jt} + \beta_3 \theta_{i,t-1} \\ &+ \beta_4 age_{it} + \beta_5 age_{it}^2 + \beta_6 sex_i + \beta_7 edu_i + \beta_8 mis_{it} + \varepsilon_{it}, \end{aligned} \tag{C.2}$$

where J_i denotes the worker's 2-digit industry and k indexes 2-digit industry groups, so that each broad sector absorbs its own cyclical slope with respect to monetary conditions. All other notation follows equation (C.1). Unlike Spec 1, the interaction term $AIX_{it} \times AIA_{jt}$ is excluded here, so β_1 and β_2 recover the unconditional main effects of exposure and adoption.

Table A.4: 2-DIGIT INDUSTRY × FFR (STANDARDISED)

	(1)	(2)
	UE	EU
AI Exposure (std.)	-0.00718	-0.00220***
	(0.00805)	(0.000375)
AI Adoption (std.)	-0.0115***	0.000208
	(0.00376)	(0.000127)
Observations	48,216	1,327,581
R-squared	0.013	0.003

Standard errors in parentheses

Sample: year > 2022 (post-ChatGPT). AIX and AIA standardised using post-2022 mean and SD.

Fixed effects: month-in-sample (MIS); 2-digit industry × federal funds rate replaces market tightness as the cyclical control. No exposure×adoption interaction term.

Controls: lagged market tightness θ_{t-1}, age, age^2, sex, education.

UE: SE clustered at uind1990 (origin industry); EU: at ind1990 (current industry).

FFR: monthly effective federal funds rate (FRED FEDFUNDS).

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

Results. Table A.4 reports the estimates. For job-finding (column 1), the AIA_{jt} coefficient is -0.0115^{***} (s.e. 0.00376): industries with higher AI adoption have significantly lower job-finding rates for unemployed workers, conditional on the industry's cyclical position. The AIX_{it} main effect is -0.00718 (s.e. 0.00805) and insignificant, corroborating the finding in Spec 1 that occupational exposure alone does not predict re-employment absent actual industry adoption.

For separations (column 2), the AIX_{it} main effect is -0.00220^{***} (s.e. 0.000375), replicating the structural feature documented in Spec 1: more exposed occupations see fewer separations absent adoption. The AIA_{jt} main effect is 0.000208 (s.e. 0.000127) and statistically insignificant, indicating that, unlike job-finding, industry-wide adoption does not have a detectable uniform effect on separation rates once the FFR-interacted 2-digit industry control and market tightness are included.

C.3 Adoption Potential IV

Identification. The IV specification addresses the residual endogeneity concern that industries selecting into AI adoption may already have been on different labor market trajectories. OLS estimates of β_2 are biased if, for example, demand-buoyant industries both adopt AI and reduce unemployment mechanically, or if labor-shedding industries adopt AI to substitute for workers while simultaneously contracting. In either case the OLS coefficient conflates the causal effect of adoption with selection. (The industry×FFR control in the second model absorbs aggregate cyclical patterns but cannot eliminate industry-specific selection.) In addition, AIX_{it} is measured with error, and we address the resulting potential bias using an instrumental variable for AIX_{it}.

We instrument AIA_{jt} with an IV whose shares are pre-LLM industry task content (pre-determined) and whose shift is aggregate national AI adoption (external to any single industry). Separately, we instrument AIX_{it} with the Felten et al. AIOE score, an alternative pre-determined occupational AI exposure measure constructed from an independent data source.

The two instruments support two 2SLS specifications: a single-endogenous design (only AIA_{jt} instrumented, AIX_{it} treated as exogenous) and a two-endogenous design (both instrumented; exactly identified). Comparing estimates across these two designs reveals whether endogeneity bias in AIX_{it} materially affects the adoption coefficient.

Empirical model. For worker i in industry j and month t, the estimating equation is

$$y_{it} = \alpha \text{AIX}_{it} + \beta \text{AIA}_{jt} + X_{it}'\gamma + \mu_m + \epsilon_{it}, \tag{C.3}$$

where X_{it} includes age, its quadratic, sex, education, and lagged tightness θ_{t-1}, and μ_m are MIS fixed effects. Standard errors are clustered at the industry level. The adoption instrument is

$$Z_{jt} = \widehat{AdoptPotential}_j^{\text{pre}} \times \overline{NatAdopt}_{t,-j}, \tag{C.4}$$

where $\widehat{AdoptPotential}_j^{\text{pre}}$ is the pre-LLM cognitive task intensity of industry j from O*NET (interacted with pre-period employment shares), and $\overline{NatAdopt}_{t,-j}$ is the leave-one-out national AI adoption index at month t, excluding industry j to eliminate any mechanical own-industry correlation. For UE transitions, Z_{jt}^{UE} is keyed on origin industry (`uind1990`); for EU transitions,

Z_{jt} is keyed on current industry (`ind1990`). In the two-endogenous design, both AIX_{it} and AIA_{jt} are treated as endogenous, and the first-stage system is

$$\text{AIX}_{it} = a\,\text{AIOE}_{it} + b\,Z_{jt} + X_{it}'\lambda_1 + \mu_m + u_{it}, \tag{C.5}$$

$$\text{AIA}_{jt} = c\,\text{AIOE}_{it} + d\,Z_{jt} + X_{it}'\lambda_2 + \mu_m + w_{it}. \tag{C.6}$$

The system is exactly identified (two instruments, two endogenous variables). In the single-endogenous design, AIX_{it} is instead treated as exogenous: it enters equations (C.5)–(C.6) directly as a regressor rather than being instrumented, and AIOE is dropped from both first stages. Although AIX_{it} is determined by the worker's current occupation and does not vary within an occupation over time, it varies at the worker×month level because workers may change occupations across CPS rotation months. The reported partial F-statistics are Sanderson–Windmeijer conditional statistics: the partial F for AIOE → AIX conditions out Z_{jt}'s contribution to the AIX first stage, and vice versa for Z_{jt} → AIA. In practice the AIOE → AIX partial F exceeds 1,000 in the pooled sample, so weak identification of AIX_{it} is not a concern except in the smallest exposure-bin subsamples.

Aggregate estimates. Table A.5 reports OLS and IV estimates for the pooled post-2022 sample. Three findings stand out.

First, the Adoption-potential instrument substantially amplifies the adoption coefficient for job-finding. OLS yields $\hat{\beta}_{\text{AIA}} = -0.332^{***}$ for UE transitions; the single-endogenous IV raises this to -1.326^{***} and the two-endogenous IV to -1.716^{***} (Table A.5, columns 2–3). The upward correction from IV is consistent with demand-driven selection: industries in which labor demand happens to be strong tend to adopt AI precisely because they are hiring, biasing OLS toward zero. The IV design, identifying off pre-period task content interacted with a national adoption trend, corrects for this. The first-stage partial F-statistics are 44 (single-endogenous) and 44

Table A.5: AI EXPOSURE AND ADOPTION: OLS AND ADOPTION-POTENTIAL IV ESTIMATES (AGGREGATE)

	UE Transitions (U → E)			EU Transitions (E → U)		
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	IV (AIA only)	IV (AIX+AIA)	OLS	IV (AIA only)	IV (AIX+AIA)
AI Exposure (AIX)	−0.054	0.009	0.004	−0.011***	−0.011***	−0.016***
	(0.036)	(0.045)	(0.067)	(0.002)	(0.002)	(0.003)
AI Adoption (AIA)	−0.332***	−1.326***	−1.716***	0.007*	−0.001	0.013
	(0.096)	(0.400)	(0.500)	(0.004)	(0.015)	(0.015)
Observations	49,436	46,363	36,374	1,327,581	1,274,503	1,274,503
Partial F: AIOE → AIX			1,913			1,397
Partial F: Z_{jt} → AIA		44	44		46	37

Standard errors in parentheses, clustered at industry.

Sample: year > 2022. UE: origin industry (`uind1990`); EU: current industry (`ind1990`).

IV (AIA only): AIA_{jt} instrumented by $Z_{jt} = \widehat{AdoptPotential}_j^{\text{pre}} \times \overline{NatAdopt}_{t,-j}$; AIX exogenous.

IV (AIX+AIA): AIX instrumented by AIOE (Felten et al.); AIA by Z_{jt}. Exactly identified (2 instruments, 2 endogenous).

Controls: age, age^2, sex, education, θ_{t-1}, MIS fixed effects.

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

(two-endogenous), comfortably above standard weak-instrument thresholds.

Second, the AIX exposure coefficient for UE transitions is near zero and insignificant across all three estimators (OLS: −0.054; IV: +0.004), confirming the Spec 1 finding that occupational exposure alone does not predict re-employment independently of actual industry adoption.

Third, for separation (EU) transitions the IV estimates diverge from OLS in a revealing way. The OLS AIA_{jt} coefficient is positive (+0.007*, column 4), but both IV specifications produce point estimates near zero or insignificant (−0.001 and +0.013, columns 5–6). The OLS positive sign disappears under instrumentation, suggesting it reflects reverse-selection — industries contracting for non-AI reasons have both elevated separations and lower AI adoption, generating a spurious positive correlation between adoption and separations in OLS. By contrast, the AIX_i coefficient on EU grows in magnitude when instrumented (−0.011*** under OLS, −0.016*** under two-endogenous IV, columns 4 and 6), indicating that occupational exposure genuinely protects against separation and that OLS understates this protective effect due to downward bias in the AIX measure.

Heterogeneity by exposure bin. Table A.6 disaggregates estimates by four fixed bins of AIX_i: G1 (0.00–0.24), G2 (0.25–0.49), G3 (0.50–0.74), and G4 (≥ 0.75). Each panel reports both IV designs.

For UE transitions (Panel A), the adoption effect is concentrated in the upper bins. In G1 and G2 — workers with below-median AI exposure — the AIA_{jt} coefficient is statistically indistinguishable from zero in both specifications, despite adequate first-stage power. In G3, the single-endogenous estimate is -2.587^{***} and the two-endogenous estimate is -3.121^{***}: adoption sharply curtails job-finding for workers whose occupations are substantially exposed to AI. G4 yields -1.153^{**} and -1.216^{**}. The nonlinearity is economically meaningful — below-median exposure workers are largely insulated from the labor-demand compression that AI adoption generates, while above-median workers face a strong and growing barrier to re-employment. The AIX coefficient in Panel A is consistently small and insignificant in the single-endogenous design but flips large and positive in the two-endogenous design for G3 and G4, where AIOE instruments a residualized exposure measure. This sign reversal is consistent with strong upward-sorting bias: workers who sort into high-AIX occupations have unobserved productivity advantages that lower their unemployment rates independently of AI, and the OLS (and single-endogenous IV) coefficient on AIX reflects this sorting rather than a causal effect.

For EU transitions (Panel B), the pattern is more nuanced. In G1, both AIX and AIA reduce separations under two-endogenous IV (-0.052^{***} and -0.138^{*} respectively): in low-exposure industries, AI adoption is associated with lower separation risk even for moderately exposed workers. In G3, the AIX coefficient becomes -0.118^{*} (more protective than OLS implied) while AIA is positive ($+0.042^{**}$): within high-adoption industries, workers in highly exposed occupations are less likely to separate, but adoption itself raises mean separation rates in this bin — consistent with AI substituting for some tasks while preserving or complementing others in the same occupational category. The weak first-stage partial F for AIOE→AIX in G3 ($11.8^{\dagger}$) counsels caution in that cell. G4 corroborates the G3 AIX pattern (-0.077^{*}) with a stronger first stage.

Table A.6: AI EXPOSURE AND ADOPTION BY FIXED AIX BIN: ADOPTION-POTENTIAL IV 2SLS ESTIMATES

	G1	G2	G3	G4
	0.00–0.24	0.25–0.49	0.50–0.74	≥ 0.75
Panel A: UE Transitions (U → E) — Z_{jt}^{UE} keyed on origin industry (`uind1990`)				
(i) IV (AIA only) — AIA_{jt} instrumented; AIX exogenous				
AI Exposure (AIX)	-0.082	-0.440^{***}	0.028	-0.207
	(0.112)	(0.143)	(0.123)	(0.299)
AI Adoption (AIA)	0.504	-0.843	-2.587^{***}	-1.153^{**}
	(1.511)	(0.660)	(0.686)	(0.479)
Observations	22,597	15,195	7,040	1,531
Partial F: $Z_{jt}^{UE} \to$ AIA	98.5	124.8	26.5	$18.0^{\dagger}$
(ii) IV (AIX + AIA) — AIOE → AIX; $Z_{jt}^{UE} \to$ AIA				
AI Exposure (AIX)	-0.141	-0.424	1.375	1.732
	(0.151)	(0.372)	(1.162)	(2.963)
AI Adoption (AIA)	-0.901	-1.176	-3.121^{***}	-1.216^{**}
	(1.504)	(0.827)	(0.957)	(0.479)
Observations	17,300	12,170	5,654	1,250
Partial F: AIOE → AIX	359.0	481.3	59.1	$7.9^{\dagger}$
Partial F: $Z_{jt}^{UE} \to$ AIA	86.5	141.7	25.0	$14.6^{\dagger}$
Panel B: EU Transitions (E → U) — Z_{jt} keyed on current industry (`ind1990`)				
(i) IV (AIA only) — AIA_{jt} instrumented; AIX exogenous				
AI Exposure (AIX)	-0.023^{***}	-0.022^{***}	0.007	-0.006
	(0.008)	(0.007)	(0.004)	(0.009)
AI Adoption (AIA)	-0.178^{**}	0.015	0.032^{*}	0.007
	(0.079)	(0.024)	(0.017)	(0.010)
Observations	393,620	592,994	238,631	49,258
Partial F: $Z_{jt} \to$ AIA	131.8	110.4	53.4	25.6
(ii) IV (AIX + AIA) — AIOE → AIX; $Z_{jt} \to$ AIA				
AI Exposure (AIX)	-0.052^{***}	-0.025	-0.118^{*}	-0.077^{*}
	(0.019)	(0.017)	(0.068)	(0.044)
AI Adoption (AIA)	-0.138^{*}	0.017	0.042^{**}	0.010
	(0.081)	(0.027)	(0.021)	(0.011)
Observations	393,620	592,994	238,631	49,258
Partial F: AIOE → AIX	236.4	380.3	$11.8^{\dagger}$	28.1
Partial F: $Z_{jt} \to$ AIA	131.5	132.0	48.8	26.8

Standard errors in parentheses, clustered at industry. Sample: year > 2022.

Bins defined on `openAI_har` (AIX). Spec (i): AIA_{jt} endogenous, instrumented by Z_{jt}; AIX exogenous.

Spec (ii): AIX instrumented by AIOE (Felten et al.); AIA instrumented by Z_{jt}. System exactly identified.

Controls: age, age^2, sex, education, θ_{t-1}, MIS fixed effects.

† Partial F below 20: instrument may be weak in this subgroup.

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

D Additional Regression Results

D.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:

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

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.7 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.7: 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.

D.2 Tech-Sector Confounds

The baseline 2SLS specification regresses each labor market transition indicator on AI adoption ($\widehat{AIA}_{it}$), AI exposure (AIX_i), market tightness (θ_{t-1}), demographic controls (age, age^2, sex, education dummies), and month-in-sample fixed effects (μ_s), using survey weights. Formally, for the unemployment-to-employment (UE) transition, the second-stage equation is

$$\Pr(UE_{ist}) = \beta_0 + \beta_1\, AIX_i + \beta_2\, \widehat{AIA}_{jt} + \mathbf{X}'_{it}\boldsymbol{\gamma} + \mu_s + \varepsilon_{ist}, \tag{D.2}$$

where the endogenous regressor AIA_{jt} (the Lightcast AI skill share in industry j at time t) is instrumented by a leave-one-out adoption-potential IV shift-share: $Z_{jt} = \widehat{AdoptPotential}_j^{\text{pre}} \times \overline{NatAdopt}_{t,-j}$. The share component, $\widehat{AdoptPotential}_j^{\text{pre}}$, is the employment-weighted mean occupational AI exposure within industry j computed over employed workers observed before December 2022; the shift component, $\overline{NatAdopt}_{t,-j}$, is the leave-one-out national average AI skill share in month t, excluding industry j. Standard errors are clustered at the worker's origin industry (`uind1990`) for UE transitions and at the current industry (`ind1990`) for employment-to-unemployment (EU) transitions. The estimation sample is restricted to observations with year > 2022.

A natural concern is that the 2022–2023 tech-sector layoff cycle—rather than economy-wide AI adoption—drives the estimated effect. Two specifications address this directly. The first re-estimates the baseline IV on a sample that excludes workers employed in Computer and Data

Processing Services (`ind1990` = 732). This removes the industry most directly affected by the 2022 tech correction from both the estimation sample and the instrument construction. If the baseline finding reflected a tech-specific hiring adjustment rather than broad AI adoption, the coefficient on AIA_{jt} should attenuate when this sector is excluded. Instead, the IV estimate for UE transitions strengthens to -1.563 (SE = 0.438, $p < 0.01$), compared with -1.326 in the baseline, and the first-stage partial F-statistic rises sharply to 163.8 (EU sample: 256.8), well above conventional weak-instrument thresholds. The EU transition coefficient remains statistically indistinguishable from zero (-0.008, SE = 0.018). The persistence—indeed amplification—of the UE effect outside the tech sector indicates that the result is not an artifact of tech-sector dynamics.

The second specification (Specification 2) splits the post-COVID sample into a pre-ChatGPT window (2021–2022) and the post-ChatGPT window (year > 2022, the baseline sample), estimating the baseline 2SLS separately for each. If the estimated coefficient captures an AI adoption effect that operates through the diffusion of large language models after December 2022, the coefficient should be negligible or statistically absent in the pre-period, when AI adoption had not yet accelerated. Consistent with this prior, the IV estimate for UE transitions in the pre-period is -6.635 (SE = 4.352), which is imprecisely estimated and statistically insignificant despite its large point value; the coefficient in the post-period is -1.326 (SE = 0.400, $p < 0.01$), recovering the baseline estimate exactly. For EU transitions, the pre-period estimate is -0.677 (SE = 0.409, $p < 0.10$)—marginally significant and substantially larger in magnitude than the post-period estimate of -0.001 (SE = 0.015, insignificant)—suggesting that if anything the 2022 tech correction modestly raised EU separation risk in high-adoption industries before the ChatGPT era, while the post-2022 separation margin is not systematically related to AI adoption after conditioning on market tightness. Together, the two specifications support interpreting the baseline UE coefficient as reflecting genuine post-ChatGPT AI adoption effects operating across industries, rather than a spurious correlation with the 2022–2023 technology-sector cycle.

D.3 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:

$$\ln w_{it} = \alpha + \beta_1 \text{Educ}_i + \beta_2 \text{Exper}_{it} + \beta_3 \text{Exper}_{it}^2 + \beta_4 \text{AIX}_{it} + \beta_5 \text{AIA}_{it} \tag{D.3}$$
$$+\beta_6 (\text{AIX}_{it} \times \text{Post-LLM}_t) + \beta_7 (\text{AIA}_{it} \times \text{Post-LLM}_t) + \gamma X_{it} + \lambda_t + \delta_{j(i)} + \varepsilon_{it},$$

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.8) 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 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]

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

[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.

Table A.8: MINCER WAGE REGRESSION

	(1) Baseline	(2) AI Factors	(3) 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.

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

E 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.4: EN, UN, NE AND NU FLOWS RATES BY AI EXPOSURE

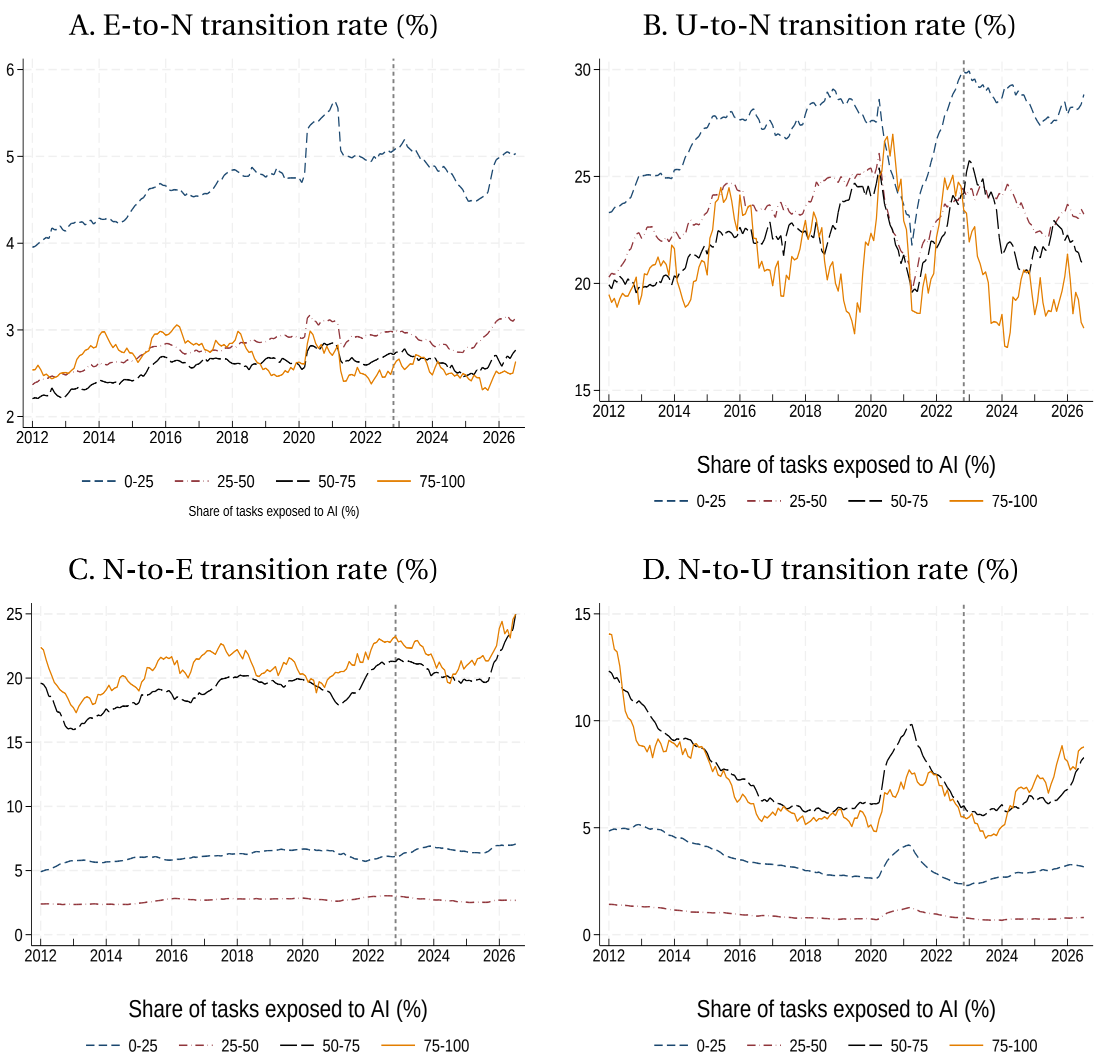

Notes: This figure displays the EN rate (Panel A), UN rate (Panel B), NE rate (Panel C), and NU rate (Panel D) for four AI exposure groups (12-month moving average). Blue dashed line: 0–25; red dashed-dot line: 25–50; black long-dashed line: 50–75; solid yellow line: 75–100. **Source**: Authors' calculation

Figure A.5: LFPR AND EPOP RATIO

A. Labor force participation rate (%) B. Employment-to-population ratio (%)

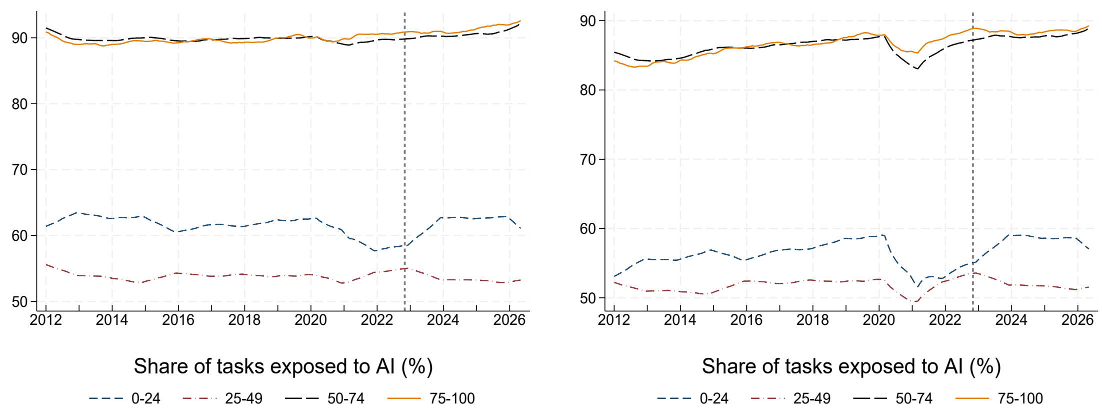

Note to figure: Panels A and B show LFPR and the EPOP ratio by AI exposure. Blue dashed line: 0–25; red dashed-dot line: 25–50; black long-dashed line: 50–75; solid yellow line: 75–100. **Source**: Authors' calculation

Figure A.4 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.4 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.5 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.6: WORKER REALLOCATION RATE: FLOWS PER CAPITA BY AI EXPOSURE (12-MONTH MOVING AVERAGE)

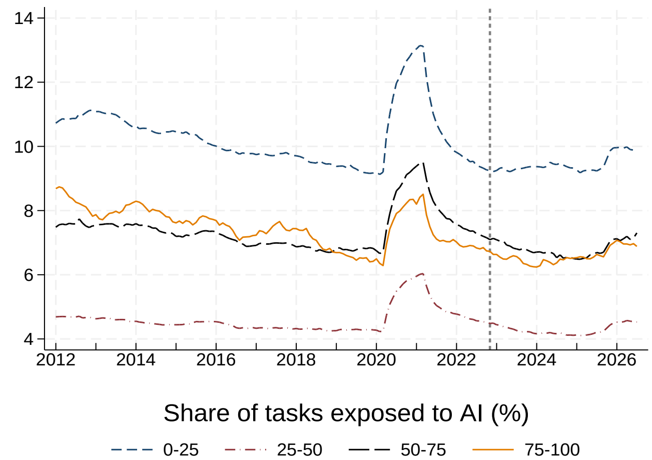

Note to figure: This figure displays the monthly worker reallocation rate by AI exposure. The unit for Y-axis is percent. **Source**: Authors' calculation

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.6 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.6).

F Bayesian Estimation of U-star

F.1 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

$$y^j_{g,t} = Z\,\alpha^j_{g,t} + \varepsilon^j_{g,t}, \tag{F.1}$$

$$\alpha^j_{g,t} = T\,\alpha^j_{g,t-1} + u^j_{g,t}, \tag{F.2}$$

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 F.2 details the estimation including priors and posterior simulations.

F.2 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:

$$\sigma^{2,j}_{\zeta,g} \sim \mathrm{IG}\big(a_\zeta,\ b_\zeta \cdot r_g\big) \qquad \text{(slope innovation)}, \tag{F.3}$$

$$\sigma^{2,j}_{\kappa,g} \sim \mathrm{IG}(a_\kappa,\ b_\kappa) \qquad \text{(cycle innovation)}, \tag{F.4}$$

$$\sigma^{2,j}_{\varepsilon,g} \sim \mathrm{IG}(a_\varepsilon,\ b_\varepsilon) \qquad \text{(observation noise)}, \tag{F.5}$$

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.9.

Table A.9: 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

F.3 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:

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

where $\mathscr{T}$ 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.

F.4 Trend Estimates

Figures A.7 and A.8 display the estimated EU and UE rates and the natural rate estimates of the four AI exposure/adoption groups, respectively.

Figure A.7: EU AND UE TRENDS BY AI EXPOSURE AND ADOPTION

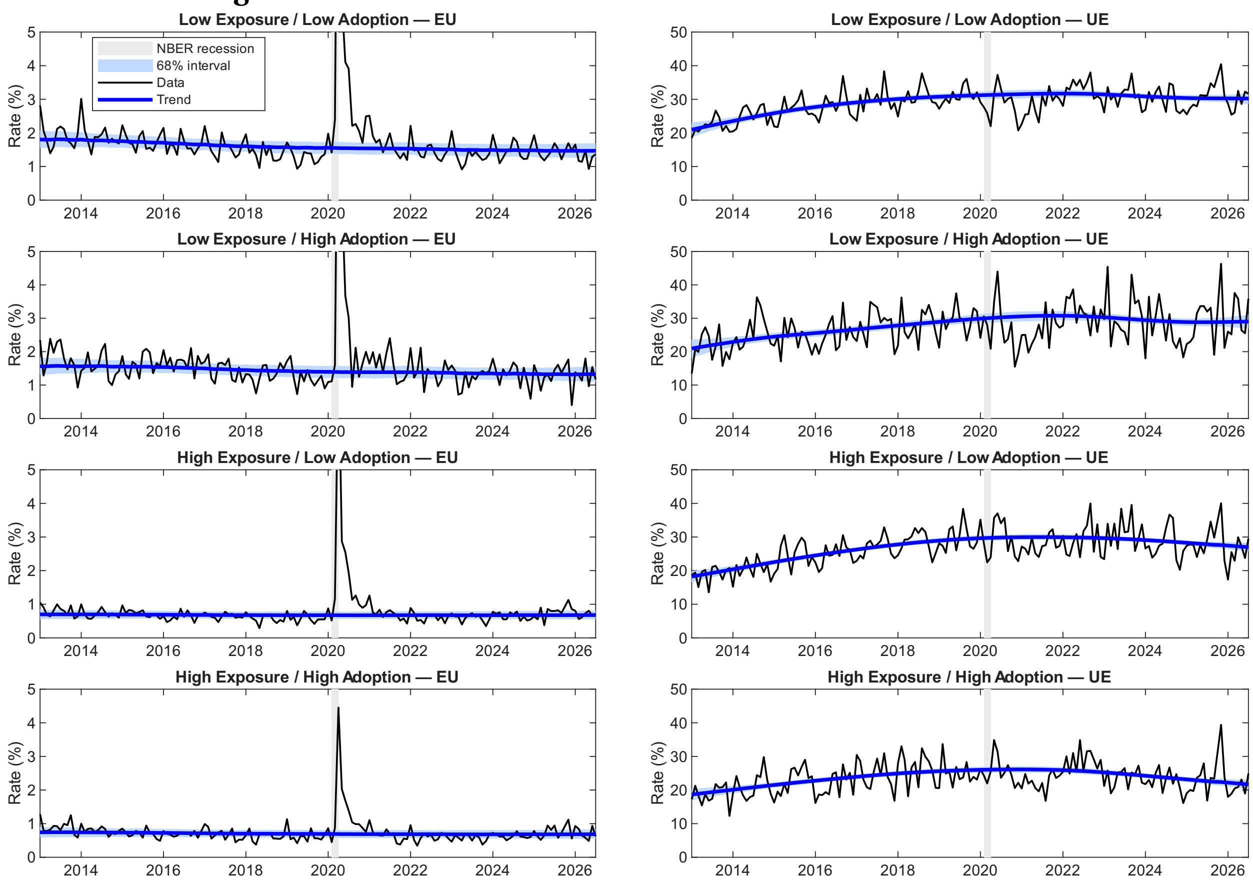

Note to figure: This figure displays unemployment inflow (EU) and outflow (UE) rates by AI exposure and adoption. The black lines represent the observed data, the blue lines denote the posterior median estimates, and the shaded blue areas indicate the 68 percent posterior intervals.

Note to figure: Authors' calculation.

Figure A.8: U-STAR ESTIMATE BY AI EXPOSURE AND ADOPTION

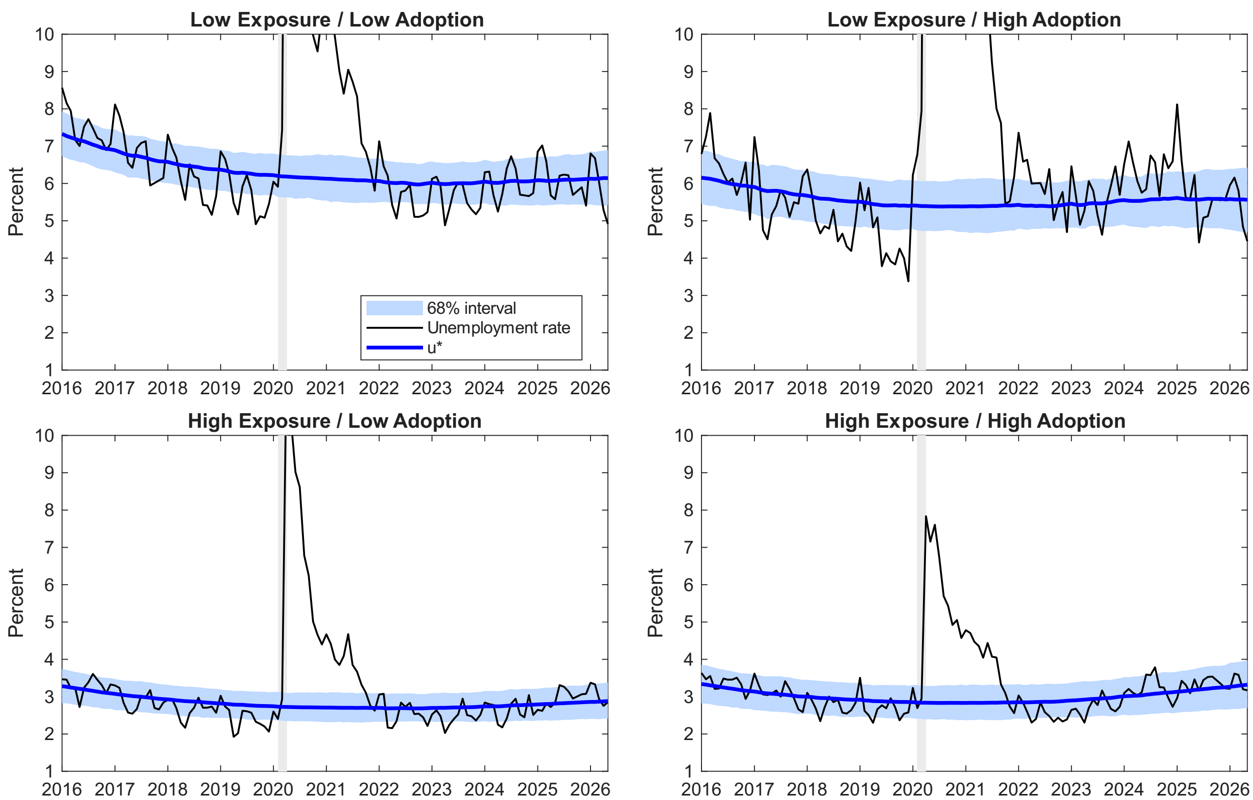

Note to figure: This figure displays the unemployment rate and the estimated natural rate for the four groups defined by AI exposure and adoption. The black lines represent the observed data, the blue lines denote the posterior median estimates, and the shaded areas indicate the 68 percent posterior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

Note to figure: Authors' calculation.