FEDERAL RESERVE BANK of CLEVELAND

Economic Commentary

11.09.2026 | ISSN 2163-3738 | Number EC 2026-26

Job Finding and Unemployment in the Age of AI

This *Economic Commentary* shows that since late 2022, unemployment has risen more in occupations whose tasks are more exposed to capabilities of artificial intelligence, especially among young workers. Worker movements in and out of unemployment reveal this is not because workers in at-risk occupations are separating from their jobs at a faster pace, but because they are taking longer to find new ones.

Alexander Cline, Barış Kaymak

Topics Labor economics, Workforce development

DOI 10.26509/frbc-ec-202626

The views authors express in Economic Commentary *are theirs and not necessarily those of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System. The series editor is Tasia Hane. This paper and its data are subject to revision; please visit clevelandfed.org for updates.*

Introduction

Debates about artificial intelligence's (AI) impact on the future of work hinge on whether AI will replace workers, augment their productivity, transform the kinds of jobs that already exist, create entirely new jobs, or some combination of these. This *Economic Commentary* looks at how the rapid deployment of large language models (LLM), beginning with the launch of ChatGPT in late 2022, and followed by subsequent advances in generative AI, are reshaping the US labor market. [1] We estimate that since generative AI became widely available, the unemployment rate for workers aged 23–54 has risen by roughly 1 percentage point more in occupations most exposed to AI compared to those in the least exposed, relative to the "pre-AI" baseline period of 2015–2019. This rise in relative unemployment outcomes for

occupational groups more exposed to AI's capabilities is substantive, and the disparity is even larger for workers aged 23–30.

To understand what is driving this shift in labor market dynamics, we look beyond the unemployment rate at labor market flows, that is, how often workers separate from their jobs and how quickly unemployed workers find new ones. Our primary finding is that the rise in relative unemployment is driven primarily by meaningfully lower job-finding rates for workers in highly exposed occupations in the post-AI era compared to rates in the pre-AI baseline period. We do not find consistent evidence of elevated job separation rates among workers in occupations highly exposed to AI. Taken together, these patterns suggest that generative AI may be making it harder, to date, to enter AI-exposed occupations rather than displacing existing workers in these occupations.

Measuring AI Exposure

From recent academic literature, we use three distinct measures of how much a given occupation's tasks overlap with what AI can do, each reflecting a different approach to this problem. Two of the measures originate from Eloundou et al. (2023), who systematically evaluate more than 19,000 individual tasks catalogued in the US Department of Labor's O*NET database, a rich compendium of workers' tasks across 923 occupations. Each task was rated according to a simple criterion, whether an LLM or LLM-powered software system could reduce the time required to complete this task by at least 50 percent while maintaining output quality. This evaluation yields three exposure categories. Tasks with no such potential are classified as "unexposed"; tasks for which a basic LLM interface alone would suffice are classified as "directly exposed"; and tasks requiring additional specialized software built atop an LLM, such as domain-specific retrieval systems or image generation tools, are classified as "indirectly exposed." From these task-level ratings, the authors construct occupation-level exposure scores. We use their beta score, which credits direct exposure fully and indirect exposure at half weight, reflecting the additional investment required for indirect applications to materialize. Critically, the authors produced two variants of this measure, one based on ratings by human experts familiar with LLM capabilities and one based on ratings by GPT-4 itself. We interpret these scores as measuring the potential for AI to substitute for tasks currently performed by humans. Both are included in our analysis and yield similar results.

A conceptually distinct approach is taken by Handa et al. (2025), who measure the extent to which AI is actually used for tasks associated with a particular job. Analyzing more than four million real-world conversations on the Claude.ai platform, they map observed interactions to O*NET task descriptions and construct occupation-level usage rates. Handa et al. (2025) find

that AI use is highly concentrated; software development and writing tasks account for nearly half of all observed interactions, and actual use in very-high-wage occupations (such as for physicians) is notably lower than Eloundou et al.'s (2023) framework would predict, likely because of regulatory barriers and the physical demands of these roles. This use score is a fundamentally different object from what Eloundou et al. (2023) measure. An occupation might be highly exposed in theory yet show low actual AI use because of regulatory constraints, employer inertia, high implementation costs, or slow adoption. Conversely, occupations in which adoption has moved quickly may show high use even if theoretical exposure is moderate. The divergence between these two methodologies can be informative about the extent to which labor market flows are responding to AI's anticipated effects versus its actual use. We return to this point in the discussion of our analysis results.

How We Track Job Finding and Job Loss across Occupations

To measure how AI exposure affects workers' odds of finding a job or losing one, we need to follow individuals over time to observe changes in their employment status. This requires repeated observations of the same worker across consecutive months linked to information about the AI exposure of their occupation so that we may compare how these transition rates evolved before and after the widespread availability of generative AI tools across occupations with varying degrees of AI exposure. Our primary data source is the Basic Monthly Current Population Survey (CPS), administered jointly by the US Census Bureau and the Bureau of Labor Statistics (IPUMS, Flood et al., 2024). [2] The CPS interviews households for four consecutive months, rotates them out for eight, and then interviews them for four more. This panel structure allows us to link individuals across adjacent months and directly observe transitions between employment states.

We restrict the sample to civilians aged 23–54, excluding most younger workers still in school and older workers near retirement whose labor market behavior may reflect lifecycle considerations unrelated to AI. We focus on the unemployment rate along with two transition outcomes. An unemployed worker is classified as a "job finder" if they are employed in the subsequent month. For example, if a person is unemployed in January and employed in February, we say that they are a job finder in February. Similarly, an employed worker is classified as "separated from their job" if they are unemployed or out of the labor force in the subsequent month. These individual flows, grouped at the occupational level (according to their AI exposure), form the dependent variables in our analysis.

Because unemployed workers do not currently hold jobs, they cannot be assigned to an occupation in the conventional sense. The CPS surveys all unemployed respondents about the occupation of their most recent job, and we use this information to classify unemployed workers into occupational groups by AI exposure level. [3] , [4]

Empirical Strategy

We group the 923 O*NET occupations into 20 equal-sized groups (ventiles) ranked by AI exposure according to each of our three measures. We then estimate, for each ventile group, the change in the average unemployment rate, job-finding rate, and job-separation rate between a pre-AI baseline period and a post-AI comparison period.

Our pre-AI baseline period is January 2015 through December 2019. This window predates the public availability of generative AI tools and is also free of COVID-19-pandemic-related distortions to the labor market, distortions which were severe and idiosyncratic during 2020–2022. Our comparison period for unemployment covers May 2025 through April 2026, a period in which LLMs had become widely available and extensively adopted by businesses and individuals (Bick, Blandin, and Deming, 2026). For job-finding and job-separation rates, we extend our comparison period to include January 2024 through April 2026. [5] Because of its rotating panel structure and attrition in the CPS, unemployed individuals with the requisite month-to-month linkages are a smaller subset of the sample used to compute unemployment rates. Extending the comparison period allows us to regain some of this lost sample size and maintain the statistical precision of the patterns documented below. [6]

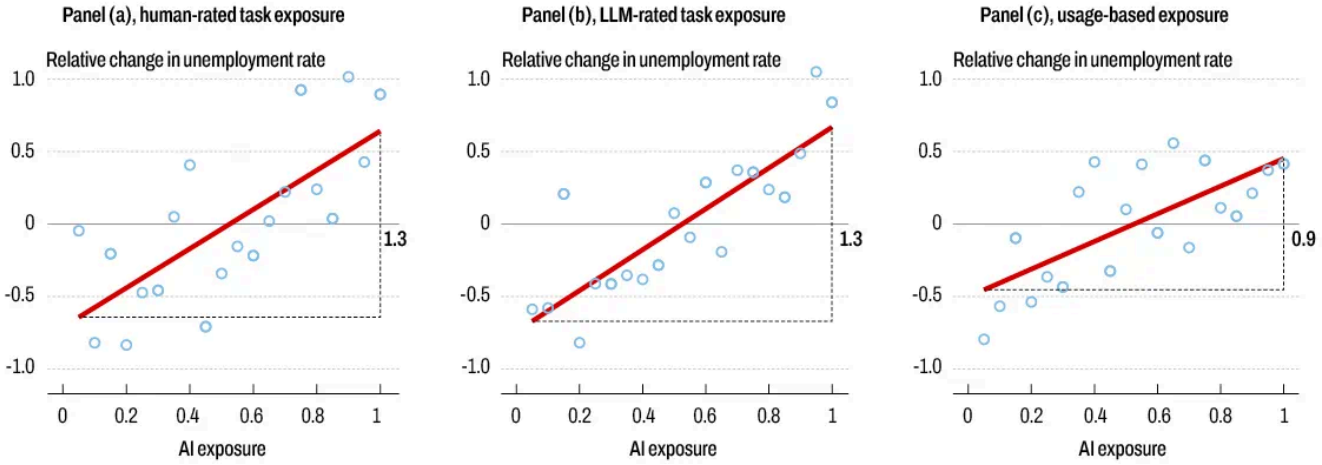
To see how AI has affected the labor market, we plot changes in the unemployment rate and the associated flow rates for each occupational group against its AI exposure ranking. We compare how much the labor market has changed before versus after the advent of generative AI across occupations with varying degrees of exposure. [7] The key question is whether unemployment or flow rates changed more in occupations in which AI can do more of the work. If AI is reshaping the labor market, we would expect the biggest changes to show up precisely in those occupations.

AI Exposure and Labor Market Outcomes

Figure 1 presents our headline result: the relationship between AI exposure and changes in the unemployment rate across our 20 occupational groups, post- minus pre-AI availability, for each of the three exposure measures. All three measures show a positive relationship. Occupations with higher AI exposure saw larger relative increases in unemployment than

occupations with lower exposure. The pattern is clearest for Eloundou et al.'s (2023) LLM-rated task exposure metric (panel (b)), where the fitted line climbs from about -0.6 percentage points at the lowest exposure occupations to about 0.7 percentage points at the highest, giving a 1.3 percentage point predicted gap between least- and most-exposed occupations (see bold number to the right of the panel). Their human-rated task exposure metric (panel (a)) shows a similar positive slope with more scatter around the fitted line. Handa et al.'s (2025) usage-based measure (panel (c)) also shows a positive slope, though slightly flatter compared to the slopes of the other measures.

Figure 1: Change in Unemployment Rate by Occupational AI Exposure



Sources: Eloundou et al. 2023. "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models"; Handa et al. 2025. "Which Economic Tasks Are Performed With AI? Evidence From Millions of Claude Conversations"; Bureau of Labor Statistics Current Population Survey

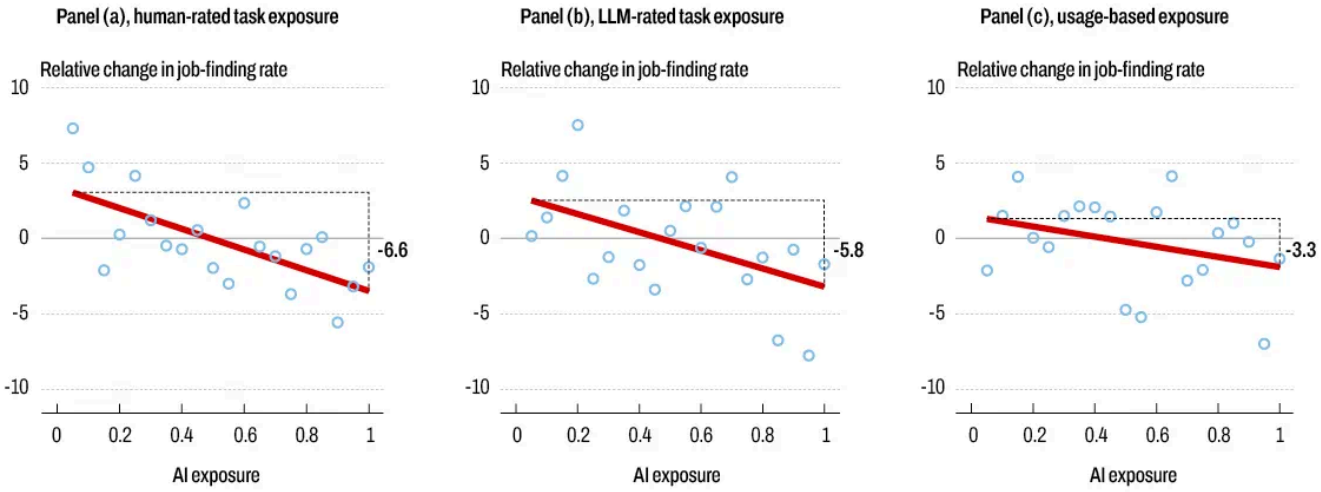
Notes: The figure plots the relative change in unemployment rates (percentage points) between the baseline period (2015–2019) and the comparison period (May 2025–April 2026) for each of 20 occupational groups (ventiles) ranked by AI exposure. Unemployed workers are assigned to occupational groups based on the occupation of their most recent job. The fitted regression line is shown in red along with the predicted gap between the least- and most-exposed occupational group displayed in bold to the right. Panels (a) and (b) use the human-rated and LLM-rated task exposures measures, respectively, from Eloundou et al. (2023); panel (c) uses the usage-based exposure measure from Handa et al. (2025). Estimates control for monthly seasonality.

To understand what is driving this divergence in unemployment rates, we turn to the underlying flows. Figure 2 plots the estimated change in the job-finding rate for each of the 20 exposure groups and for each of the three AI exposure measures. The red line is again a fitted regression line summarizing the overall slope.

Task-based measures (panels (a) and (b)) tell a consistent and economically meaningful story. Both the human-rated and LLM-rated variants show a clear negative slope. Occupations with the highest AI exposure experienced job-finding rates that declined by roughly six percentage points more than occupations with the lowest exposure. This is a substantial shift. In a labor market where monthly job-finding rates for unemployed workers average around 34

percent, a 6 percentage point decline represents a meaningful deterioration in reemployment prospects for workers in exposed fields. The implied average unemployment duration increases from about 10.5 weeks to more than 13 weeks. [8]

Figure 2: Change in Job-Finding Rate by Occupational AI Exposure



Sources: Eloundou et al. 2023. "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models"; Handa et al. 2025. "Which Economic Tasks Are Performed With AI? Evidence From Millions of Claude Conversations"; Bureau of Labor Statistics Current Population Survey

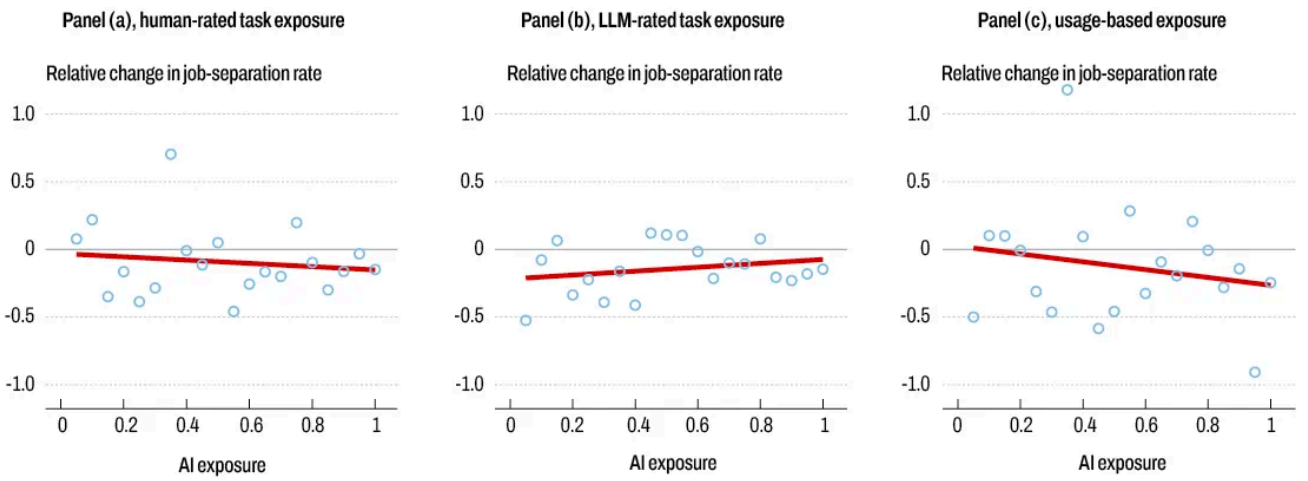
Notes: The figure plots the relative change in monthly job-finding rates (percentage points) between the baseline period (2015–2019) and the comparison period (January 2024–April 2026) for each of 20 occupational groups (ventiles) ranked by AI exposure. Unemployed workers are assigned to occupational groups based on the occupation of their most recent job. The fitted regression line is shown in red along with the predicted gap between the least- and most-exposed occupational group displayed in bold to the right. Panels (a) and (b) use the human-rated and LLM-rated task exposures measures, respectively, from Eloundou et al. (2023); panel (c) uses the usage-based exposure measure from Handa et al. (2025). Estimates control for monthly seasonality.

Usage-based exposure (panel (c)) shows a similar but somewhat weaker pattern with a flatter slope. This divergence likely reflects the nuance between what the two measures capture. While task-based measures capture the degree to which AI could, in principle, replace human effort on a given task, the usage-based measure potentially blends two very different relationships between AI and work: substitution, in which AI does the task instead of the worker, and augmentation, in which AI helps the worker do the task better or faster. Handa et al. (2025) estimate that roughly 43 percent of observed AI usage reflects substitution and 57 percent reflects augmentation. In occupations in which augmentation dominates, AI complements worker tasks, which can raise output per worker. A more productive worker generates more value for each dollar of wages paid, which makes those workers more attractive to hire, an effect that would work against finding a negative job-finding relationship. If the augmentation effect partially offsets the substitution effect within the usage-based measure, the flatter slope we observe is exactly what we would expect.

Another reason for the flatter relationship could be that actual usage reflects not only AI's capability to substitute for human work but also the costs of adoption, such as reorganizing workflows, retraining staff, and integrating AI tools into existing systems. Occupations in

which adoption is costly may show low current usage even if AI is fully capable of performing their tasks. If adoption costs are largely unrelated to AI's substitution potential, then the usage-based measure may simply be too noisy to detect a substitution signal. Under this interpretation, the flatter slope in panel (c) of Figure 2 reflects the uneven pace of adoption across occupations rather than an absence of underlying disruption.

Figure 3: Change in Job-Separation Rate by Occupational AI Exposure



Sources: Eloundou et al. "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models"; Handa et al. 2025. "Which Economic Tasks Are Performed With AI? Evidence From Millions of Claude Conversations"; Bureau of Labor Statistics Current Population Survey

Notes: The figure plots the relative change in monthly job separation rates (percentage points) between the baseline period (2015–2019) and the comparison period (January 2024–April 2026) for each of 20 occupational groups (ventiles) ranked by AI exposure. The job separation rate is defined as the fraction of employed workers who transition to either unemployment or out of the labor force in the following month. Panels correspond to the human-rated, LLM-rated, and usage-based AI exposure measures, from left to right. See text for further details.

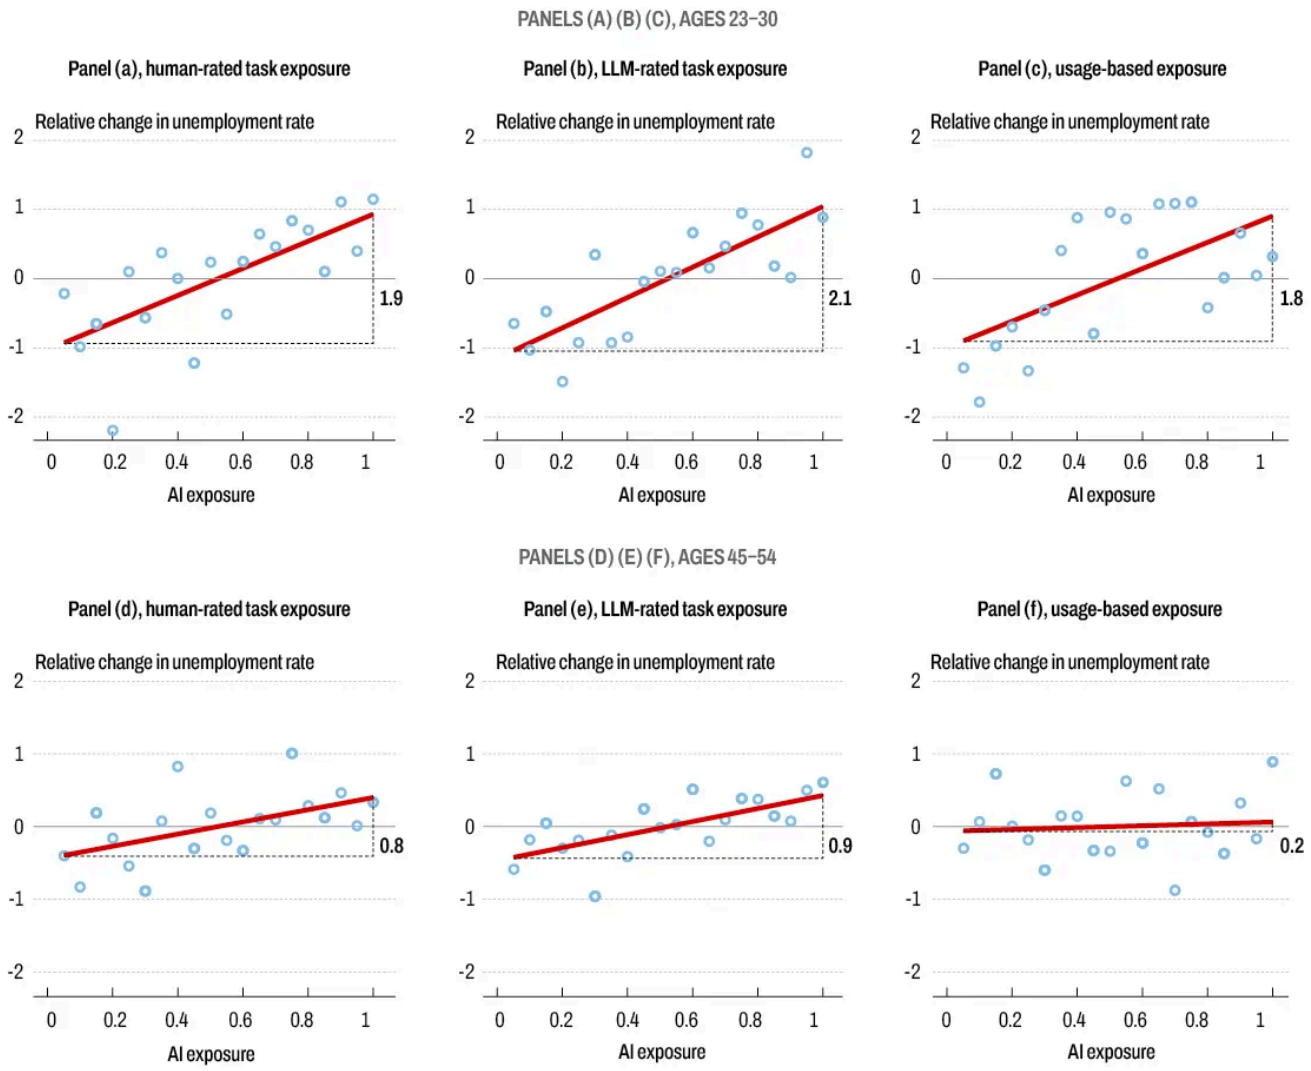
Figure 3 presents the corresponding results for job separation rates. The results here are mixed. The Eloundou et al. (2023) human-based measure shows essentially no relationship between AI exposure and changes in separation rates, while the other two measures suggest a mild relationship, but the signs are contradictory. One implies slightly higher separation in exposed occupations, the other slightly lower. None of these patterns are statistically robust.

We interpret this lack of a consistent and substantive relationship as evidence that AI exposure has not (yet) meaningfully elevated the rate at which workers separate from jobs they currently hold. One interpretation is that firms are not yet replacing incumbent workers at scale; another is that any displacement is occurring through slower hiring rather than active layoffs or quits, something which would show up in job-finding rates rather than separation rates. Our results are consistent with this latter interpretation.

Young Workers Bear the Brunt

Figure 4 splits the sample into younger workers (ages 23–30) and older workers (ages 45–54) and replicates the unemployment rate analysis for each group. [9]

Figure 4: Change in Unemployment Rate by Occupational AI Exposure and Age Group



Sources: Eloundou et al. 2023. "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models"; Handa et al. 2025. "Which Economic Tasks Are Performed With AI? Evidence From Millions of Claude Conversations"; Bureau of Labor Statistics Current Population Survey

Notes: The figure plots the relative change in monthly unemployment rates (percentage points) between the baseline period (2015–2019) and the comparison period (May 2025–April 2026) for each of 20 occupational groups (ventiles) ranked by AI exposure. The top row restricts the sample to workers aged 23–30, and the bottom row does the same for workers aged 45–54. Panels correspond to human-rated, LLM-rated, and usage-based AI exposure measures, from left to right. See text for further details.

Younger workers experience a more pronounced positive relationship between AI exposure and unemployment. The fitted regression lines in the top row imply that moving from the lowest to the highest AI exposure group increases the predicted unemployment rate for younger workers by about two percentage points. For workers aged 45 and older (bottom row), the relationship between AI exposure and unemployment is positive but muted. The

fitted associations imply that moving from the lowest to the highest AI exposure group increases the predicted unemployment rate for older workers by, at most, one percentage point.

This age asymmetry has a natural interpretation. Younger workers are disproportionately competing for entry-level positions, which are specifically the roles in which AI tools may be most quickly substituting for routine cognitive tasks. Experienced workers in exposed occupations may be more insulated, at least in the short run, because their accumulated expertise, professional networks, and firm-specific knowledge collectively make them harder to replace. AI's disparate impact on younger workers is consistent with a gradual restructuring of occupations from the bottom up. [10]

Younger workers also experience a more negative relationship between AI exposure and job-finding prospects (not shown), but the age asymmetry is not as stark. Job separation rates (also not shown) display no distinctive pattern for either age group, consistent with the aggregate results.

Conclusion

Our findings for jobs vulnerable to potential AI-for-human substitution are consistent with what some observers have called a "no hire, no fire" labor market dynamic, in which firms respond to uncertainty about AI's transformative potential, not by laying off existing workers, but by quietly pulling back on new hiring (Ockerman, 2025). This pattern is visible in our data. Job separation rates in AI-exposed occupations are largely unchanged, while job-finding rates have fallen meaningfully. Firms appear to be waiting to see how AI reshapes their workflows before committing to new headcount, effectively using attrition and reduced hiring as a low-risk way to shrink their exposure to labor costs in occupations for which AI may soon substitute for human work. For many workers already inside these occupations, this may feel like business as usual. For those on the outside trying to get in—particularly younger workers entering the labor market for the first time—the door appears to be harder to open.

Our assessment is based on an early and necessarily incomplete look at what is likely to be a long and complex adjustment process. The short-run effects we document may look quite different from the equilibrium that emerges over the next couple of decades. Historically, general-purpose technologies have eventually raised productivity and labor demand broadly, even as they eliminated specific types of work in the short run (Hötte, Somers, and Theodorakopoulos, 2023). Whether the adoption of generative AI follows this pattern remains an open question.

References

- Bick, Alexander, Adam Blandin, and David Deming. 2026. “The Rapid Adoption of Generative AI.” *Management Science* 72(10): 8309–9253, iv–vi. doi.org/10.1287/mnsc.2025.02523.
- 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. Stanford Digital Economy Lab. digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. “GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models.” Working paper. arxiv.org/abs/2303.10130.
- Eckhardt, Sarah, and Nathan Goldschlag. 2025. “AI and Jobs: The Final Word (Until the Next One).” Economic Innovation Group, August 10. eig.org/ai-and-jobs-the-final-word/ .
- Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Etienne Breton, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, David Van Riper, and Kari C.W. Williams. 2025. “Integrated Public Use Microdata Series, Current Population Survey: Version 13.0.” Minneapolis, MN: IPUMS. doi.org/10.18128/D030.V13.0.
- Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin, Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, Kevin Troy, Dario Amodei, Jared Kaplan, Jack Clark, and Deep Ganguli. 2025. “Which Economic Tasks Are Performed with AI? Evidence from Millions of Claude Conversations.” Anthropic. Working paper. arxiv.org/abs/2503.04761.
- Hosseini, Seyed, and Guy Lichtinger. 2026. “Generative AI as Seniority-Biased Technological Change: Evidence from US Resume and Job Posting Data.” Harvard University. Working paper. papers.ssrn.com/sol3/papers.cfm?abstract_id=5425555.
- Hötte, Kerstin, Melline Somers, and Angelos Theodorakopoulos. 2023. “Technology and Jobs: A Systematic Literature Review.” *Technological Forecasting & Social Change* 194: 122750. doi.org/10.1016/j.techfore.2023.122750.
- Ockerman, Emma. 2025. “The US Labor Market Ground to a Halt in 2025. The Risk in 2026 Is That It ‘Cracks.’” Yahoo! Finance, December 16. finance.yahoo.com/news/the-us-labor-market-ground-to-a-halt-in-2025-the-risk-in-2026-is-that-it-cracks-140026614.html .
- Simon, Lisa. 2025. “Is AI Responsible for the Rise in Entry-Level Unemployment?” Revelio Labs, August 4. reveliolabs.com/news/macro/is-ai-responsible-for-the-rise-in-entry-level-

unemployment/.

Endnotes

1. Generative AI refers to artificial intelligence that creates new content, such as text or images, based on patterns in its training data. Return to 1
2. The CPS is the official source for US unemployment statistics and surveys approximately 60,000 households each month. Return to 2
3. We follow the methodology of Eckhardt and Goldschlag (2025) in creating a crosswalk to match AI exposure scores with the CPS occupation codes in the earlier years of our sample. Return to 3
4. This commonly adopted approach misses new labor market entrants that lack a previous occupation. In addition, workers may search for jobs in a different occupational group than the one they last held, and to the extent that they are, we expect AI's true effect on labor market outcomes to be stronger than what we document here. Return to 4
5. We exclude the November 2025 survey because the relevant information for linking individuals across monthly surveys is not available in October 2025 because of disruptions in data collection during the US government shutdown. Return to 5
6. Because labor market flows may exhibit seasonal fluctuations, we include month-of-year fixed effects in all regressions. This allows us to control for monthly seasonality throughout the analysis and isolate variation in labor market outcomes that is unrelated to predictable seasonal patterns. Return to 6
7. While our difference-in-differences approach accounts for any changes in aggregate macroeconomic conditions between the baseline and comparison periods, it is reassuring that the average national unemployment rate for people aged 23–54 was about 4.5 percent in both periods. The average change in the figures is always normalized to zero. Return to 7
8. Letting f denote the monthly job-finding rate, the implied average unemployment duration in weeks is $-4.33/\log(1-f)$, where $4.33=52/12$ (the average number of weeks in a month). Return to 8
9. Sample excludes the first unemployment spell at labor market entry for young workers since they do not have a previous occupation. Return to 9
10. Our findings resonate with Brynjolfsson, Chandar, and Chen (2025), Simon (2025), and Hosseini and Lichtinger (2026), among others, who highlight AI's disproportionate effect on employment levels among young workers. Return to 10

Suggested Citation

Cline, Alexander, and Barış Kaymak. 2026. “Job Finding and Unemployment in the Age of AI.” Federal Reserve Bank of Cleveland, Economic Commentary 2026-26.
https://doi.org/10.26509/frbc-ec-202626

This work by Federal Reserve Bank of Cleveland is licensed under Creative Commons Attribution-NonCommercial 4.0 International