informs
https://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE

Vol. 00, No. 0, Xxxxx 0000, pp. 000–000

Submitted to *Management Science*

ISSN 0025-1909, EISSN 1526-5501

Canaries in the Coal Mine? Seven Facts about the Recent Employment Effects of Artificial Intelligence

Erik Brynjolfsson
Stanford University and NBER, erikb@stanford.edu

Bharat Chandar
Stanford University, chandarb@stanford.edu

Ruyu Chen
Stanford University, ruyuchen@stanford.edu

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and are not intended to be a true representation of the article's final published form. Use of this template to distribute papers in print or online or to submit papers to another non-INFORM publication is prohibited.

Abstract. Using high-frequency administrative payroll data from ADP covering millions of U.S. workers through June 2026, we document seven facts about the labor market following the widespread adoption of generative AI. (1) We find no evidence of widespread, economy-wide job displacement. (2) However, employment of young workers (ages 22–25) in AI-exposed occupations now stands 19% below where it would be had it kept pace with that of their less-exposed peers; experienced workers show no comparable gap. (3) This divergence has widened steadily since we first documented it in August 2025. (4) It operates primarily through reduced hiring of young workers rather than increased separations. (5) Declines are concentrated in occupations where AI usage primarily substitutes for human tasks; where usage primarily complements workers, employment is flat or rising, especially for experienced workers. (6) The divergence is not explained by several prominent alternative factors: it persists when excluding technology firms and computer occupations, when controlling for exposure to interest-rate increases and for remote work, and across alternative measures of AI exposure. (7) Adjustment is occurring through employment rather than base compensation. These patterns attenuate when controlling for education, show some divergent trends predating generative AI, and are more pronounced in the ADP panel than in national survey benchmarks, with some evidence of consistent patterns in government administrative data. We interpret these facts as early, descriptive indicators—canaries in the coal mine—rather than causal estimates, and we provide a public set of AI Economic Indicators to facilitate ongoing tracking of changes in the economy.

Funding: This research was supported in part by the Stanford Digital Economy Lab, which receives financial support from a variety of individuals, groups and organizations through Stanford University. They are listed at `https://digitaleconomy.stanford.edu/home/support/`. ADP provided data for our research and supports the Stanford Digital Economy Lab as a member of its Corporate Affiliate Program, as well as through a sponsored research project unrelated to this paper. ADP has the right to review the paper to prevent the disclosure of confidential information.

Key words: Artificial Intelligence, Labor Economics, Future of Work

The proliferation of artificial intelligence (AI) tools, especially generative AI, has sparked a global debate about its potential impact on the labor market. This discourse spans predictions of enhanced productivity, fears of widespread job displacement, and skeptical views that AI will have small

economic effects. Historically, technologies have affected different occupations in heterogeneous ways, replacing work in some, complementing others, and transforming still others. However, the specific capabilities of generative AI, particularly its proficiency with codified knowledge work, suggest there may be "canaries in the coal mine" which will be affected first.[1]

AI capabilities have improved rapidly. For instance, between 2023 and 2024, AI performance on software engineering benchmarks surged from 4.4% to 71.7% (**?**). Recent work finds that current systems match or outperform industry professionals nearly half of the time on a pre-defined benchmark of economically valuable tasks (**?**). At the same time adoption rates among U.S. workers have approached 50% (**??**). Growing capabilities and adoption have sparked concerns that AI may displace human labor, particularly among entry-level workers in exposed fields like software engineering and customer service (**?**).[2] Amid this intense debate, a growing literature aims to estimate the employment effects of AI, either in aggregate or for specific subgroups like young workers. A key challenge in the literature is balancing three features that are core to measuring AI's labor market impact: timeliness, worker granularity, and administrative verifiability.

To help address this gap, we leverage high-frequency administrative payroll data from ADP, encompassing millions of workers through June 2026. The data allow us to document generative AI-era labor market shifts in close to real-time using administrative records, including for subgroups that have been the topic of particular interest. We identify a stark divergence in employment outcomes based on age and AI exposure. While overall employment remains robust, employment of young workers (ages 22–25) in AI-exposed jobs now stands 19% below where it would be had it kept pace with that of their less-exposed peers. Experienced workers show no comparable gap between more- and less-exposed occupations: much of what differs across these occupations affects them too, and their employment did not diverge. The declines are concentrated where AI is used to automate work; we do not find this same pattern in jobs where AI is used to complement workers. In levels, employment of 22–25-year-olds in the two most exposed quintiles fell about 11% between November 2022 and June 2026, while employment of the same age group in the three least-exposed quintiles grew about 10%. The slowdown in entry-level hiring in AI-exposed occupations has driven overall stagnation in employment growth for young workers. Earlier versions of this study headlined regression estimates adjusting for firm-level shocks (a 13% relative decline as of July

[1] This paper updates and extends "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence" (August 2025) with data through June 2026 and new analyses. Online Appendix Section EC.11 details the changes.

[2] Improved productivity of workers in an occupation could lead to either reduced or increased employment, depending on, among other things, how elastic demand is for the output of those workers.

2025 data; 16% as of September 2025 data). We now emphasize the simpler descriptive divergence, which requires no modeling choices: by this same measure, the kept-pace shortfall was 15% at the September 2025 data vintage and has since widened to 19% as of June 2026.

In sum, we document seven facts: (1) no evidence of widespread, economy-wide job displacement from AI; (2) employment of young workers in AI-exposed occupations stands 19% below where it would be had it kept pace with that of their less-exposed peers, with no comparable gap for experienced workers; (3) the divergence has widened steadily since we first documented it in August 2025; (4) it operates primarily through reduced hiring rather than increased separations; (5) declines are concentrated in occupations where AI usage substitutes for human tasks, while complementary usage is associated with flat or rising employment; (6) the divergence is not explained by several prominent alternative factors, including technology-sector dynamics, interest rates, and remote work; and (7) adjustment is occurring through employment rather than base compensation.

Additional analysis provides suggestive evidence about mechanisms. We find employment declines for young workers in occupations that involve codified knowledge. Occupations that involve tacit knowledge see faster employment growth for experienced workers. Finally, women face greater AI exposure on average, suggesting heterogeneity by gender worth monitoring going forward.

These patterns broadly remain when considering alternative analyses that exclude computer occupations, exclude firms in the technology sector, account for occupational interest rate exposure, or consider effects of work-from-home and outsourcing, and they remain directionally consistent with estimates that control for industry- and firm-level shocks through firm-time fixed effects. Three patterns in our own analysis merit particular attention, and we address them directly rather than in passing. First, more- and less-exposed occupations show some divergent trends that predate ChatGPT, particularly around the COVID-19 pandemic (Section **??**). Second, estimates attenuate when controlling for occupational education levels, which may reflect either an alternative explanation or the very channel through which AI operates (Section **??**). Third, the divergence is directionally consistent but more pronounced in the ADP panel than in national survey benchmarks (Section **??**). Against these concerns stand four countervailing findings: the divergence has continued to widen through mid-2026, long after interest rates peaked; by November 2022 the relative position of exposed occupations had already returned to approximately its pre-pandemic level, so the subsequent decline moves the gap below that baseline rather than back toward it; the declines load specifically on automative AI usage with a clear age gradient, a pattern that interest-rate, education, and remote-work stories do not predict; and US government administrative data show

consistent raw descriptive patterns by age and industry AI exposure (**?**).[**TODO:** Tucker shows model-free descriptive results at the industry x state x age level consistent with our findings, which is encouraging. However, Johnston and Makridis also have a paper looking at administrative data and find employment increases, though they don't split by age, and I don't think their approach is very good. Humlum and Vestergaard likewise use administrative data in Denmark. I'm not sure if it's worth mentioning them here, but leaving a note just in case. *EB: OK, lets' go with the Claude suggested approach*] We caution that this work does not estimate a causal impact of AI: these are descriptive facts, and ongoing work is needed to determine how much they are caused by the spread of generative AI rather than merely correlated with it. This study documents these facts and provides data infrastructure for tracking them going forward.

A rapidly growing body of evidence studies similar questions using different data and methods, finding a range of results (Chandar 2025, **?**, Eckhardt and Goldschlag 2025, **?**, **?**, **?**, **?**, **?**, **?**, **?**, **?**, **?**, **?**, **?**, Lambert and Schindler 2026, **?**).[3]

Our study combines three features that advance this body of knowledge in the US. First, the timeliness of our data series allows us to track changes in close to real-time, mirroring the rapid development of AI technologies. Second, the size and granularity of the data set enables tracking outcomes for important subgroups like young workers in AI-exposed occupations, in contrast to alternative sources such as the Current Population Survey, Current Employment Statistics, Quarterly Census of Employment and Wages, or Business Trends and Outlook Survey that either do not collect this information or lack statistical precision. Third, the administrative nature of our data mitigates representativeness and endogenous reporting concerns present in job posting or social media records, though we caution that there may also be selection in which firms use ADP.[4] The primary challenge in this study and the broader literature is causally separating AI from other changes in the economy, offering many avenues for further research. Our study is complementary to this expanding set of work.

Why might AI disproportionately affect entry-level workers? One possibility is that AI more effectively substitutes for knowledge that has been digitized and codified—the core of formal education—while acting as a complement to the tacit, experience-based knowledge held by senior employees (**?**). Consequently, AI may be automating the checkable, process-intensive tasks that

[3] See also Dominski and Lee (2025), Eckhardt and Goldschlag (2025), **?**), **?**), **?**), **?**), **?**), **?**), **?**), **?**.

[4] Generative AI can endogenously affect the matching process between workers and firms in ways that affect the interpretation of signals like job postings, such as by inducing employers to create more job postings that they have less intention to hire for (**?**).

historically justified entry-level headcount, while increasing the leverage of experienced staff (**??**). This slowdown in hiring at the entry level provides early suggestions that generative AI may be affecting the career ladder in the American labor market.

Implications The aim of this project is to highlight descriptive facts that are relevant to current discussions about the implications of artificial intelligence on the labor market. More research will help clarify how much of the patterns we document are caused by AI and through what mechanisms AI may change work. Causally identifying the effects of AI on labor markets is subtle and challenging. New methodologies and data sets will be needed to separate different factors.[5]

The data we present provide a snapshot of the labor market in the present. We do not make predictions about future impacts or claim that the trends we document will continue indefinitely. Future impacts of AI will be mediated by factors such as the continued progress of the technology; the rate of take up by firms and workers; policy responses; endogenous adjustment by firms and workers; and other macroeconomic forces. In prior waves of technological innovation the generality of human capability led to continued employment and rising real wages, with each wave having unique characteristics. Tracking the indices in this paper will help clarify how AI is similar and different from previous changes. We have released a public set of AI Economic Indicators to track the facts in this paper on an ongoing basis.[6]

1. Data

1.1. Payroll Data

We analyze anonymized administrative payroll data from ADP, providing a near real-time view of the US labor market. Our primary analysis sample is a balanced panel of firms observed monthly from January 2021 through June 2026, comprising between 3.7 and 4.5 million employees per month. To examine an extended pre-ChatGPT period we also consider a secondary balanced sample from January 2018 to June 2026. We restrict to full-time workers under age 70 with positive earnings. Employment and headcount both refer to a worker-firm match in the data. In a given month less than 0.1% of workers have full-time jobs at multiple firms.

[5] One promising direction for improving causal identification is to compare outcomes in firms that adopt AI compared to firms that do not rather than rely on differential trends by exposure indices. Several recent papers have taken this approach. However, as discussed in **?**, these analyses face challenges disambiguating AI impacts from endogeneity of adoption, poor measurement, anticipation, and general equilibrium effects. Frameworks for reconciling and integrating these analyses are a fruitful direction for research.

[6] The AI Economic Indicators are available at `https://digitaleconomy.stanford.edu/project/indicators/`.

To link workers to AI exposure, we utilize ADP's internal taxonomy, which maps standardized job titles to 2010 Standard Occupational Classification (SOC) codes based on job descriptions and industry data. Job titles are missing for roughly 30% of the raw sample; where possible we fill in missing occupation codes using a worker's own most recent (or, failing that, next) non-missing code. We show descriptive statistics around missingness in Online Appendix Section EC.1.1. We find that the descriptive facts are largely insensitive to the imputation.

In Online Appendix EC.1.2 we also report comparisons between the ADP data and the broader US economy.[7] The ADP sample overrepresents the manufacturing and wholesale sectors and underrepresents the retail and accommodation/food sectors compared to the Business Dynamics Statistics. It also overrepresents larger firms, with few firms having fewer than 10 employees.[8] Online Appendix Fig. EC.10 shows that, compared to the ACS and CPS, the ADP sample overrepresents occupations with higher AI exposure, for both men and women. The aggregate trends we document apply across the millions of workers represented in the ADP balanced sample, but overall labor market patterns in national statistics may vary due to the sample differences. Further details on sample construction, demographic imputation, and representativeness comparisons are provided in Online Appendix Section EC.1. [9]

1.2. Occupational AI Exposure and Covariates

We primarily employ two complementary measures of AI exposure. First, we use the GPT-4 based β exposure ratings from Eloundou et al. (2024), which estimate the susceptibility of O*NET tasks to Large Language Model (LLM) automation. Second, we utilize the Anthropic Economic Index (Handa et al. 2025), which tracks the share of actual queries to the model Claude related to specific O*NET tasks. A key feature of the Anthropic index is its classification of queries as either "automative" (substituting for labor) or "augmentative" (complementing labor); we leverage this distinction to test heterogeneous effects.[10] These measures change over time as usage evolves.

[7] See also Cajner et al. (2018), ADP Research (2025).

[8] Online Appendix Fig. EC.7 shows that the growth rate of employment for the balanced sample of firms exceeds employment growth rates in the unbalanced panel and in national statistics. By taking a balanced sample we condition on firm survival, which may predict faster employment growth.

[9] Relative to the previously circulated version of this paper, we have improved the data pipeline in two main ways. First, we improved the crosswalk that merges ADP occupation codes to external data sources such as the AI exposure measures. Second, we impute missing occupation codes as described above. The results are qualitatively similar after these changes, though the firm fixed-effects estimates are more sensitive than the raw descriptive results.

[10] The Anthropic Economic Index classifies AI usage as "automative" or "augmentative" (Handa et al. 2025), and earlier versions of this paper adopted that terminology. We now write "automation" and "complementary usage": in the growth literature, (labor-)augmenting technical change has a distinct formal meaning—technology that multiplies effective labor in the production function—whereas complementarity expresses directly the economic relationship we examine. We retain the Anthropic Economic Index's original labels when referring to its classifications in tables and figures.

We consider the automation and complementarity measures both from the March 2025 release of the Anthropic Economic Index, which we use for our main results, and from the September 2025, January 2026, and April 2026 releases, which we report in Online Appendix Section EC.5.2.[11]

In Online Appendix Section EC.4 we also consider alternative AI exposure measures from Brynjolfsson et al. (2018), Felten et al. (2021), Felten et al. (2023), and Tomlinson et al. (2025).

All exposure measures were aggregated to the occupational level and merged with payroll data using a BLS 2010-to-2018 SOC crosswalk. The original Eloundou et al. (2024) data are at the 8-digit 2018 SOC code level. To merge these to the ADP occupation codes, we convert the occupations to 6-digit SOC codes and crosswalk them to the 2010 vintage. Online Appendix Tables EC.2 through EC.6 list the largest occupations by ADP employment in each quintile.

For robustness analyses, we classify occupations as teleworkable using classifications from Dingel and Neiman (2020). We also use data on remote work by occupation from Lambert and Schindler (2026).[12] Real earnings are computed using the Personal Consumption Expenditure (PCE) index (base October 2017). We benchmark our findings against the monthly Current Population Survey (CPS) and with the American Community Survey (ACS).

2. Results

2.1. Fact 1: We do not observe widespread job displacement from AI

We first document descriptive facts about employment in the ADP data. Figure **??** shows the overall employment changes by AI exposure quintile, using occupational exposure from Eloundou et al. (2024). The data suggests a possible mild slowdown in employment growth for the most exposed jobs. Average employment in the ADP sample rose by about 6% between November 2022 and June 2026, while employment for the most exposed quintile grew by about 4%. These results are consistent with recent evidence that AI-exposed jobs have not seen widespread employment declines (Chandar 2025, Eckhardt and Goldschlag 2025, **?**). This aggregate stability is itself informative: claims of economy-wide AI-driven job losses are not visible in payroll data through June 2026.

2.2. Fact 2: Employment has declined for early-career workers in the most AI-exposed occupations

Employment of young workers in AI-exposed jobs now stands 19% below where it would be had it kept pace with that of their less-exposed peers; experienced workers show no comparable gap.

[11] We pool the data from the September 2025, January 2026, and April 2026 releases. While the March 2025 release only includes consumer application usage, the later releases also include API usage. For the later releases we combine consumer and API usage.

[12] We thank the authors for providing this data.

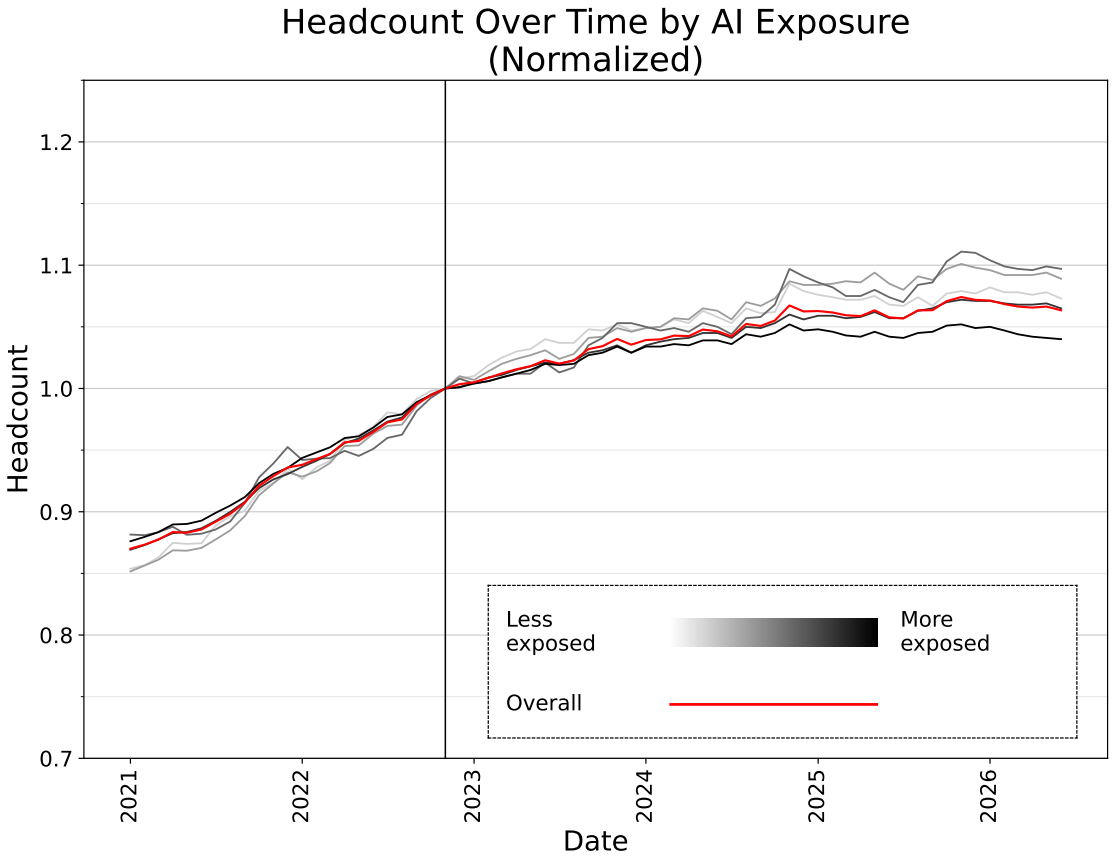


Figure 1 **Employment changes by exposure quintile using measures from Eloundou et al. (2024), pooling across all age groups. Exposure quintiles are defined based on the GPT-4 β measures.**

Figure **??** shows that employment growth for young workers has slowed relative to older workers since the start of 2023. By June 2026, employment for 22-25 year-olds in the balanced panel was about 2% below its November 2022 level, while every older age group had grown, by as much as 11% for workers aged 35-40. This is consistent with recent reports of a worsening job market for entry-level workers (**??**).[13]

Fig. **??** shows employment growth interacting age with AI exposure quintile. The top left plot shows a divergence in employment outcomes for more and less exposed occupations for workers aged 22-25, with more exposed occupations experiencing declining employment. For older age groups, we find much less marked differences in employment growth across AI exposure quintiles. Online Appendix, Fig. EC.12 through EC.14 show case studies of employment changes for specific occupations, such as software development and customer service.

[13] We compare employment trends in the ADP data to the CPS and ACS in Section **??**.

Figure 2 **Employment changes by age in the ADP sample, normalized to 1 in November 2022.**

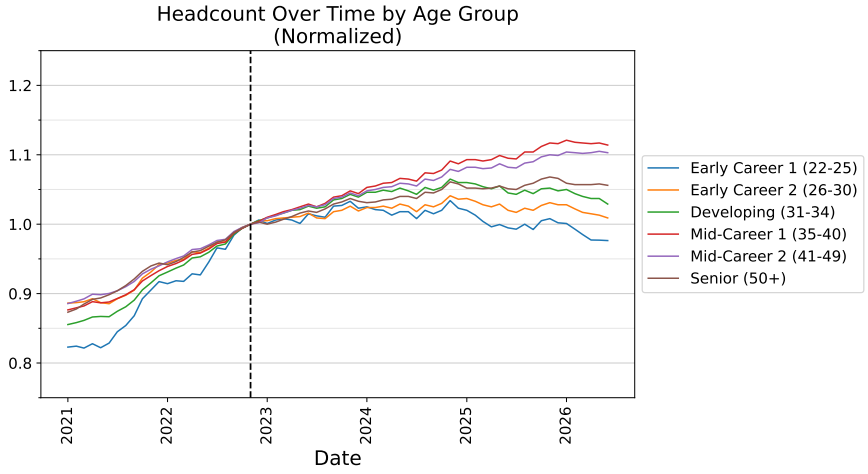
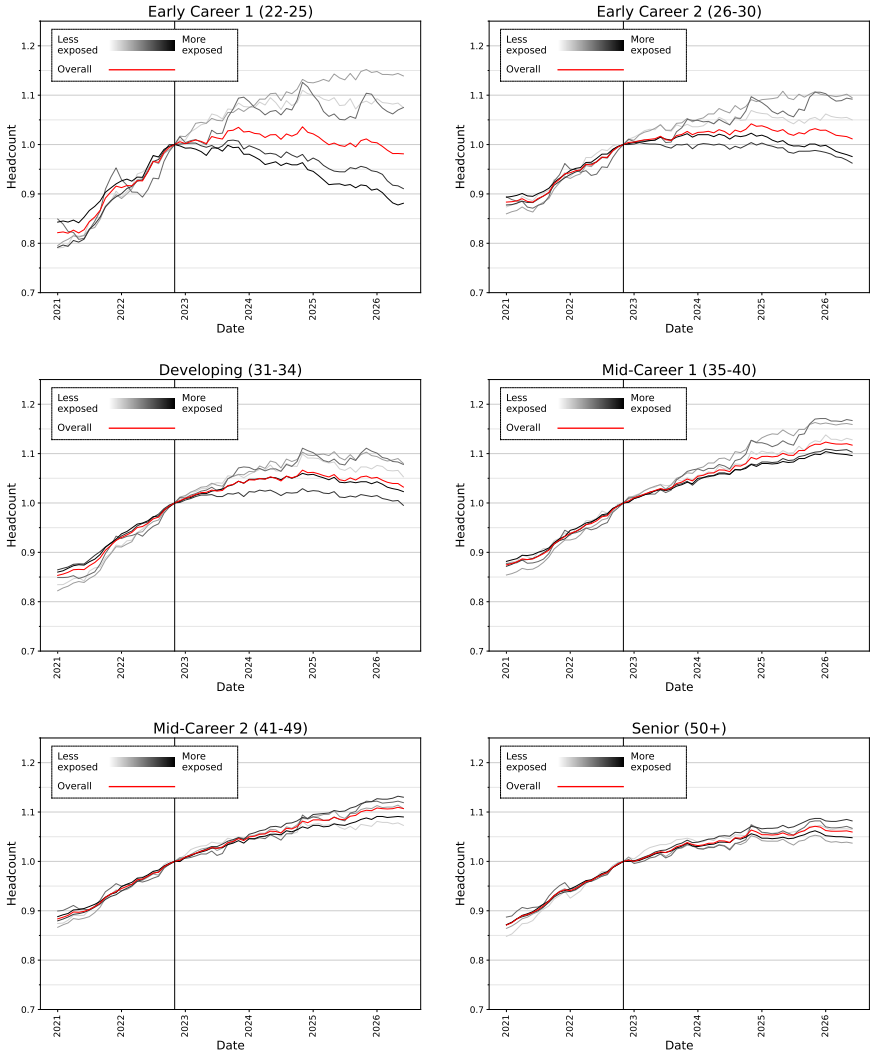


Figure 3 **GPT-4 β. Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Exposure quintiles are defined based on the GPT-4 β measures. Darker lines are more exposed quintiles. The red line shows the overall trend pooling across quintiles.**



The employment slowdown for AI-exposed jobs is most evident for 22-25 year-old workers. Employment of 22–25-year-olds in the two most exposed quintiles fell about 11% between November 2022 and June 2026, while employment of the same age group in the three least-exposed quintiles grew about 10%—a divergence of 21 percentage points, or 19% relative to growth for the bottom three quintiles. Online Appendix, Fig. EC.11 shows that for each age group, employment growth from late 2022 to June 2026 was 5-16% for the lowest three AI exposure quintiles, with no clear ordering in employment growth by age, while employment for workers aged 35-49 in the highest two exposure quintiles grew by about 10%. About 60% of occupations in the first exposure quintile saw rising early-career employment over this period, compared to about 30% of occupations in the fifth quintile (Online Appendix, Fig. EC.15). These results show that declining employment in AI-exposed jobs is driving tepid overall employment growth for workers between the ages of 22 and 25. The two most exposed quintiles accounted for about 57% of employment in this age group in November 2022, so their roughly 11% decline through June 2026 subtracted about 6 percentage points from the age group's overall employment; growth in the less-exposed quintiles offset most of this, leaving total employment for 22–25-year-olds roughly flat (a 1.9% decline) over the period. Reallocation to less-exposed occupations does not fully offset this trend.

As 22-25 year-old workers account for less than 10% of the sample, differences in overall employment growth by AI exposure quintile are relatively muted. The notable interaction between age and AI exposure demonstrates that moderate aggregate changes can mask larger changes in specific subgroups like entry-level workers, demonstrating the value of large-scale microdata for tracking labor market impacts of AI.

Occupation-Level Regression Estimates As a further descriptive test of how age interacts with AI exposure, we estimate occupation-level long-difference regressions. For each age group, we regress the percent change in employment in each occupation between November 2022 and June 2026 on indicators for the occupation's AI exposure quintile, with heteroskedasticity-robust standard errors.[14] Table **??** reports the estimates for our main sample for this analysis—firms balanced since January 2018, the sample that permits assessing pre-trends—weighting occupations by employment; Table **??** repeats the exercise for firms balanced since 2021. Panel A includes no controls; Panels B–F add occupational controls for interest rate exposure (Zens et al. 2020), college share, the work-from-home measure of Lambert and Schindler (2026), and teleworkability (Dingel and Neiman 2020).

Table 1 Employment-weighted occupation-level long-difference estimates of the percent change in employment between November 2022 and June 2026, by age group and AI exposure quintile. Each cell is the coefficient on an exposure-quintile indicator (quintile 1 omitted), relative to the least-exposed quintile. The sample balances firms observed since January 2018. The panels differ in the occupational controls included: none (A); interest rate exposure (Zens et al. 2020) (B); college share from the 2017 ACS (C); the work-from-home measure of Lambert and Schindler (2026) (D); teleworkability (Dingel and Neiman 2020) (E); and interest rate, college share, and work from home jointly (F). All controls enter as quintile indicators except teleworkability, which is binary. Heteroskedasticity-robust standard errors in parentheses. Observation counts vary by age group because some occupations have no employment in certain age groups.

	Age group					
Exposure quintile	22–25	26–30	31–34	35–40	41–49	50+
Panel A: No controls						
Quintile 2	0.039 (0.040)	0.040 (0.035)	0.025 (0.032)	0.050* (0.030)	0.041 (0.033)	−0.029 (0.029)
Quintile 3	−0.013 (0.042)	0.025 (0.035)	0.017 (0.034)	0.053 (0.036)	0.065* (0.038)	0.013 (0.037)
Quintile 4	−0.145*** (0.041)	−0.064** (0.032)	−0.045 (0.028)	0.003 (0.026)	0.065** (0.030)	0.015 (0.027)
Quintile 5	−0.179*** (0.036)	−0.048 (0.033)	−0.014 (0.030)	0.001 (0.031)	0.031 (0.032)	−0.009 (0.029)
Panel B: + Interest rate exposure						
Quintile 2	0.038 (0.042)	0.036 (0.035)	0.019 (0.032)	0.049 (0.030)	0.041 (0.033)	−0.028 (0.030)
Quintile 3	−0.011 (0.045)	0.017 (0.036)	0.005 (0.036)	0.048 (0.037)	0.063 (0.039)	0.016 (0.037)
Quintile 4	−0.146*** (0.041)	−0.068** (0.032)	−0.050* (0.028)	0.002 (0.026)	0.064** (0.030)	0.013 (0.028)
Quintile 5	−0.178*** (0.041)	−0.062* (0.034)	−0.036 (0.030)	−0.011 (0.031)	0.025 (0.030)	−0.009 (0.029)
Panel C: + College share						
Quintile 2	0.065 (0.049)	0.057 (0.052)	0.047 (0.053)	0.062 (0.045)	0.034 (0.041)	−0.040 (0.033)
Quintile 3	0.049 (0.054)	0.034 (0.050)	0.024 (0.051)	0.064 (0.048)	0.047 (0.042)	−0.000 (0.038)
Quintile 4	−0.065 (0.056)	−0.066 (0.051)	−0.045 (0.050)	0.015 (0.045)	0.035 (0.038)	−0.006 (0.032)
Quintile 5	−0.091* (0.053)	−0.045 (0.051)	−0.015 (0.051)	0.013 (0.048)	0.000 (0.040)	−0.029 (0.034)
Panel D: + Work from home (Lambert and Schindler 2026)						
Quintile 2	0.037 (0.046)	0.025 (0.039)	0.011 (0.037)	0.055 (0.036)	0.032 (0.037)	−0.027 (0.032)
Quintile 3	−0.010 (0.050)	0.011 (0.041)	0.009 (0.040)	0.065 (0.044)	0.068 (0.044)	0.016 (0.043)
Quintile 4	−0.122** (0.050)	−0.080** (0.039)	−0.066* (0.037)	0.022 (0.039)	0.069* (0.039)	0.011 (0.035)
Quintile 5	−0.152*** (0.046)	−0.066* (0.038)	−0.043 (0.036)	0.015 (0.039)	0.021 (0.040)	−0.027 (0.035)
Panel E: + Teleworkability (Dingel and Neiman 2020)						
Quintile 2	0.041 (0.040)	0.039 (0.035)	0.024 (0.032)	0.051* (0.030)	0.041 (0.033)	−0.028 (0.029)
Quintile 3	−0.011 (0.043)	0.018 (0.036)	0.013 (0.035)	0.057 (0.036)	0.061 (0.038)	0.012 (0.038)
Quintile 4	−0.140*** (0.042)	−0.089*** (0.034)	−0.057* (0.030)	0.015 (0.029)	0.052 (0.032)	0.004 (0.029)
Quintile 5	−0.173*** (0.040)	−0.075** (0.034)	−0.028 (0.030)	0.015 (0.033)	0.015 (0.033)	−0.022 (0.031)
Panel F: + All controls (interest rate, college share, work from home)						
Quintile 2	0.062 (0.050)	0.049 (0.053)	0.037 (0.054)	0.064 (0.046)	0.033 (0.043)	−0.030 (0.033)
Quintile 3	0.056 (0.056)	0.029 (0.052)	0.024 (0.053)	0.066 (0.053)	0.054 (0.048)	0.011 (0.041)
Quintile 4	−0.051 (0.058)	−0.058 (0.053)	−0.041 (0.052)	0.028 (0.050)	0.054 (0.044)	0.002 (0.034)
Quintile 5	−0.080 (0.057)	−0.044 (0.052)	−0.031 (0.052)	0.014 (0.050)	0.008 (0.043)	−0.029 (0.034)
Occupations	608	630	642	653	661	671

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 2 Employment-weighted occupation-level long-difference estimates of the percent change in employment between November 2022 and June 2026, by age group and AI exposure quintile. Each cell is the coefficient on an exposure-quintile indicator (quintile 1 omitted), relative to the least-exposed quintile. The sample balances firms observed since 2021, shown for comparison with the primary 2018 sample (Table ??). The panels differ in the occupational controls included: none (A); interest rate exposure (Zens et al. 2020) (B); college share from the 2017 ACS (C); the work-from-home measure of Lambert and Schindler (2026) (D); teleworkability (Dingel and Neiman 2020) (E); and interest rate, college share, and work from home jointly (F). All controls enter as quintile indicators except teleworkability, which is binary. Heteroskedasticity-robust standard errors in parentheses. Observation counts vary by age group because some occupations have no employment in certain age groups.

	Age group					
Exposure quintile	22–25	26–30	31–34	35–40	41–49	50+
Panel A: No controls						
Quintile 2	0.064 (0.046)	0.045 (0.035)	0.027 (0.033)	0.031 (0.028)	0.034 (0.026)	−0.028 (0.023)
Quintile 3	0.000 (0.040)	0.043 (0.052)	0.025 (0.049)	0.038 (0.039)	0.046 (0.030)	0.002 (0.028)
Quintile 4	−0.164*** (0.038)	−0.087*** (0.028)	−0.058** (0.028)	−0.025 (0.023)	0.058** (0.023)	0.017 (0.020)
Quintile 5	−0.193*** (0.033)	−0.074*** (0.028)	−0.030 (0.028)	−0.032 (0.026)	0.017 (0.024)	−0.017 (0.023)
Panel B: + Interest rate exposure						
Quintile 2	0.066 (0.049)	0.046 (0.037)	0.024 (0.034)	0.032 (0.028)	0.035 (0.026)	−0.028 (0.023)
Quintile 3	0.009 (0.039)	0.048 (0.056)	0.023 (0.053)	0.039 (0.042)	0.045 (0.031)	0.003 (0.028)
Quintile 4	−0.161*** (0.037)	−0.087*** (0.029)	−0.061** (0.027)	−0.025 (0.023)	0.056** (0.023)	0.015 (0.021)
Quintile 5	−0.174*** (0.038)	−0.065** (0.033)	−0.035 (0.032)	−0.034 (0.029)	0.012 (0.022)	−0.020 (0.022)
Panel C: + College share						
Quintile 2	0.128 (0.082)	0.071 (0.067)	0.050 (0.062)	0.049 (0.046)	0.037 (0.037)	−0.036 (0.029)
Quintile 3	0.096 (0.072)	0.086 (0.070)	0.051 (0.067)	0.068 (0.051)	0.039 (0.037)	−0.011 (0.032)
Quintile 4	−0.034 (0.077)	−0.032 (0.062)	−0.028 (0.057)	0.015 (0.044)	0.040 (0.034)	−0.008 (0.029)
Quintile 5	−0.059 (0.075)	−0.014 (0.062)	0.001 (0.058)	0.009 (0.045)	−0.001 (0.035)	−0.041 (0.030)
Panel D: + Work from home (Lambert and Schindler 2026)						
Quintile 2	0.052 (0.050)	0.015 (0.040)	0.003 (0.038)	0.027 (0.033)	0.026 (0.031)	−0.030 (0.027)
Quintile 3	−0.020 (0.048)	0.014 (0.046)	0.003 (0.046)	0.041 (0.042)	0.044 (0.037)	−0.002 (0.034)
Quintile 4	−0.165*** (0.052)	−0.105*** (0.038)	−0.084** (0.037)	−0.013 (0.036)	0.055 (0.034)	0.001 (0.029)
Quintile 5	−0.185*** (0.047)	−0.095** (0.038)	−0.064* (0.037)	−0.026 (0.037)	0.004 (0.034)	−0.045 (0.030)
Panel E: + Teleworkability (Dingel and Neiman 2020)						
Quintile 2	0.066 (0.046)	0.045 (0.036)	0.026 (0.033)	0.032 (0.028)	0.034 (0.026)	−0.029 (0.023)
Quintile 3	0.005 (0.040)	0.042 (0.053)	0.024 (0.050)	0.042 (0.040)	0.042 (0.030)	−0.001 (0.028)
Quintile 4	−0.145*** (0.041)	−0.093*** (0.031)	−0.060** (0.030)	−0.012 (0.026)	0.043* (0.025)	0.002 (0.022)
Quintile 5	−0.171*** (0.037)	−0.080** (0.033)	−0.032 (0.030)	−0.017 (0.028)	−0.000 (0.025)	−0.035 (0.022)
Panel F: + All controls (interest rate, college share, work from home)						
Quintile 2	0.124 (0.081)	0.061 (0.068)	0.041 (0.063)	0.049 (0.048)	0.037 (0.040)	−0.029 (0.030)
Quintile 3	0.087 (0.063)	0.079 (0.061)	0.050 (0.061)	0.071 (0.051)	0.041 (0.040)	−0.006 (0.034)
Quintile 4	−0.040 (0.070)	−0.033 (0.058)	−0.031 (0.053)	0.021 (0.045)	0.052 (0.037)	−0.006 (0.030)
Quintile 5	−0.051 (0.075)	−0.012 (0.060)	−0.009 (0.055)	0.009 (0.046)	0.003 (0.037)	−0.047 (0.031)
Occupations	624	640	650	663	668	677

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

For workers aged 22-25, employment in the most exposed quintile declined by about 18 percentage points relative to the least exposed quintile (Panel A; -0.179). Estimates for older age groups are smaller, mirroring the descriptive patterns and consistent with the placebo role of experienced workers. Fact 6 (Section **??**) summarizes the robustness controls in Panels B, D, and E; Section **??** discusses the education control (Panels C and F), which produces the greatest attenuation.

The preceding analyses are descriptive results about the ADP sample, which may differ from the broader US labor market in important ways (Section **??**). The subsequent analyses test alternative explanations and put the results in context of broader labor market trends.

2.3. Fact 3: The entry-level divergence has grown over time

The gap in employment between the most and least AI-exposed occupations for young workers has grown steadily over time. Figure **??** plots occupation event study estimates for the most exposed quintile (relative to quintile 1) for workers aged 22-25, tracing the long-difference coefficient as the endpoint of the difference is rolled forward month by month, with no controls and under four control sets: occupational interest rate exposure, education (college share), the Lambert and Schindler (2026) work-from-home measure, and all three jointly. The relative decline grows to roughly 18 percentage points by mid-2026. Results for other age groups, which show no comparable persistent declines, are in Online Appendix Section EC.3.9. The initial version of this study (August 2025) documented this divergence, and it has continued to widen since: by the kept-pace measure emphasized above, the shortfall was 15% at the September 2025 data vintage and 19% as of June 2026.

The estimates show differential trends between more and less exposed occupations before ChatGPT, particularly during the COVID-19 Pandemic. This is consistent with the pattern in Figure **??** and with evidence in other studies (**???**). Section **??** quantifies these pre-trends and discusses their implications for interpretation.

[14] Clustering at the three-digit SOC level (about 90 groups) does not change the results qualitatively, with confidence intervals of similar width across specifications—if anything slightly narrower under the additional controls.

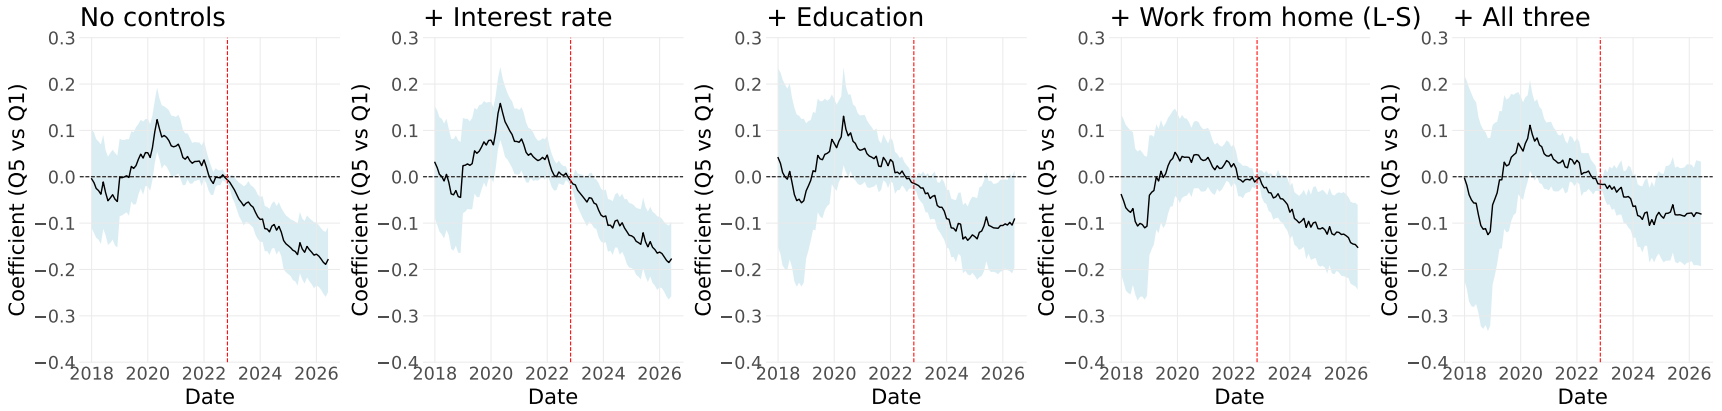


Figure 4 **Employment-weighted occupation event study estimates for workers aged 22-25, most exposed quintile relative to quintile 1, tracing the occupation-level long-difference coefficient as the difference endpoint rolls forward over time. Estimated on the extended balanced panel of firms observed from January 2018—the primary sample of Table ??—so the June 2026 endpoints coincide with the corresponding entries of Table ??. The first panel has no controls; each subsequent panel adds a different set of occupational controls. Shaded regions are 95% confidence intervals; standard errors clustered by occupation. The vertical line marks November 2022.**

2.4. Fact 4: The decline operates through reduced hiring rather than increased exits

The employment patterns above track the stock of employment. A decline in the stock can reflect reduced hiring or increased separations. The distinction matters: the first describes a narrowing entry point for new workers, the second the displacement of workers already employed. We measure hiring and separation rates for each (age group, exposure quintile) cell; Online Appendix Section EC.10 gives the definitions.

Online Appendix Figure EC.16 plots annual hiring and separation rates by age group, for the most and least exposed quintiles. During the COVID-19 Pandemic, hiring rates fell sharply. More exposed roles saw smaller declines in hiring and greater declines in separations, consistent with greater insulation for knowledge work. These differences indicate pretrends that warrant caution. By 2022 hiring and separation rates had largely recovered to pre-pandemic levels.

Since 2022, both hiring and separation rates have fallen across age groups and exposure levels, consistent with a "low hire, low fire" labor market that may particularly harm the prospects of young workers (**?**). We find no evidence that increased separations explain the divergence between more- and less-exposed young workers. Separation rates fell for both groups and, among young workers, fell at least as much in the most exposed occupations as in the least exposed—the opposite of what displacement would imply. The divergence instead reflects hiring, with the gap in hiring rates between more- and less-exposed young workers opening after 2022. These findings are broadly consistent with recent evidence from **?** and **?**, who find decreased entry into entry-level roles, and **?**, who find suggestive evidence of higher rates of layoffs for AI-exposed college-educated workers in California. We caution that these quantities are determined in equilibrium, with changes in firms' hiring potentially a function of workers' exit rates and vice versa.[15]

2.5. Fact 5: Declines are concentrated where AI usage substitutes for workers; complementary usage shows flat or rising employment

AI can either complement or substitute for labor. These may have very different implications for the labor market (**?**).

To assess how employment patterns differ based on the complementarity or substitutability of AI with labor, we use data on generative AI usage from the Anthropic Economic Index (Handa et al. 2025, Appel et al. 2025, 2026, Massenkoff et al. 2026). The Index provides an estimate of the share of queries that pertain to each occupation. In addition, for each task it reports estimates of

[15] Two alternative channels could generate these patterns: differential rates of aging out of an age band, and switching between exposure quintiles. Neither changed meaningfully over our sample and do not explain the changes we document, consistent with evidence from **?**.

the share of queries pertaining to that task that are "automative," "augmentative," or none of the above. We use these classifications as estimates of whether the usage of AI in an occupation is primarily a substitute or complement for labor. Throughout, the classification describes how AI is *used* in an occupation; whether employment rises or falls is a finding, not an implication of the label. Online Appendix, Tables EC.8 and EC.9 show example occupations that are in the highest and lowest exposure category for each measure. Figure **??** shows results from the March 2025 release of the index, with results for later periods in Online Appendix Section EC.5.2. Panel A of Fig. **??** shows employment changes by overall prevalence of related Claude queries for 22-25 year-olds.[16] The patterns match the findings using the measures from (Eloundou et al. 2024) closely. Panel B likewise shows that the occupations with the highest estimated automation shares have experienced declining employment for the youngest workers.

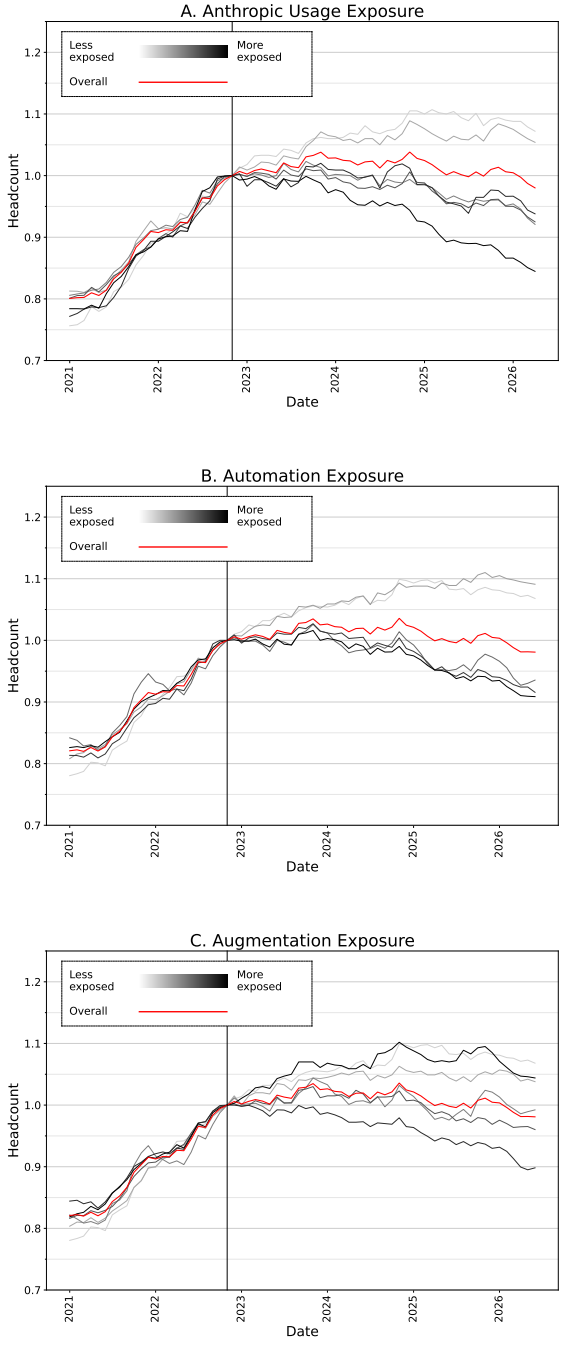
In contrast, Panel C indicates that the occupations with the highest estimated complementarity ("augmentative") shares have *not* experienced a similar pattern. Employment changes for young workers are not ordered by complementarity exposure, as the fifth quintile has among the fastest employment growth. The findings are consistent with automative uses of AI substituting for labor while complementary uses are associated with flat or rising employment.[17] Online Appendix Section EC.5.2 reports the same analysis using the more recent releases of the Anthropic Economic Index, pooling the September 2025, January 2026, and April 2026 data. Employment declines with overall usage and with automation exposure. The complementarity gradient is mixed.

To distinguish the three usage measures, which are correlated across occupations, Table **??** reports an occupation-level long-difference regression of the percent change in employment from November 2022 to June 2026 on standardized automation, complementarity ("augmentative"), and overall-usage exposures entered jointly, estimated separately by age group. For the youngest workers, only the automation coefficient is negative and statistically significant—about -0.10 per standard deviation for 22-25 year-olds—and it shrinks monotonically with age. The complementarity coefficient, by contrast, is positive and statistically significant for workers aged 41-49 ($+0.023$) and 50+ in both releases. This age gradient—automation associated with declines for the young,

[16] See Online Appendix, Fig. EC.54 through EC.56 for other age groups.

[17] Occupations in the first two augmentation quintiles have very low Claude usage (0.01% and 0.09% of conversations for the average occupation, respectively), with a high share of conversations classified as neither automative nor augmentative. In contrast, occupations in the third through fifth quintiles average 0.47%, 0.39%, and 0.33% of Claude conversations. For the automation measure, overall Claude usage increases on average with the automation share, with the lowest exposure group averaging 0.05% of conversations and the highest group averaging 0.73%. Online Appendix, Fig. EC.57 and EC.58 show that automation and augmentation results are similar when dropping occupations with low overall Claude usage.

Figure 5 ***Automation and Augmentation.*** **Employment changes by age and exposure quintile using measures from Handa et al. (2025). Panel A: *Overall usage.* Exposure quintiles are defined based on the share of queries to Claude that relate to tasks associated with an occupation, scaled by BLS employment in the occupation. Occupations whose associated tasks all have fewer than the minimum number of queries to appear in the usage data are treated as a separate category, coded as 0. Panel B: *Automation.* Automation levels are defined based on the share of queries related to an occupation that are classified by Claude as automative in nature. Occupations whose associated tasks all have fewer than the minimum number of queries to appear in the usage data are coded as 0. Note that greater than 20% of occupations above the minimum query threshold have an estimated automation share of 0. All occupations in the first and second quintile are consequently grouped together in level 1. The remaining quintiles are coded as 2, 3 and 4. Panel C: *Augmentation.* Augmentation quintiles are defined based on the share of queries related to an occupation that are classified by Claude as augmentative.**



complementarity with gains for the experienced—is the paper's most direct evidence on mechanism. The pattern holds under both the March 2025 release (Panel A) and the pooled later releases (Panel B).

Table 3 Occupation-level long-difference regression for usage, automation, and augmentation. Each panel reports a single employment-weighted regression per age group of the percent change in employment from November 2022 to June 2026 on the standardized automation ("automative"), complementarity ("augmentative"), and overall-usage exposures from the Anthropic Economic Index (Handa et al. 2025), all three entered jointly. Coefficients are the percent change in employment per standard deviation of each exposure. Panel A uses the March 2025 release; Panel B pools the September 2025, January 2026, and April 2026 releases. Standard errors clustered by occupation in parentheses. $^{*}p<0.1$**,** $^{**}p<0.05$**,** $^{***}p<0.01$**.**

Age group	Automation ("automative")	Complementarity ("augmentative")	Overall usage
Panel A: March 2025 AEI release			
22–25	-0.096^{***}	0.015	-0.026
	(0.017)	(0.020)	(0.017)
26–30	-0.035^{***}	0.010	-0.000
	(0.012)	(0.018)	(0.016)
31–34	-0.016^{*}	0.014	-0.015
	(0.010)	(0.014)	(0.018)
35–40	-0.013	0.003	-0.021
	(0.009)	(0.013)	(0.024)
41–49	-0.008	0.023^{***}	-0.018
	(0.008)	(0.009)	(0.020)
50+	-0.006	0.014^{*}	-0.033^{***}
	(0.007)	(0.008)	(0.012)
Panel B: Pooled September 2025, January 2026, and April 2026 AEI releases			
22–25	-0.082^{***}	-0.010	-0.021
	(0.015)	(0.018)	(0.014)
26–30	-0.040^{***}	-0.002	-0.005
	(0.011)	(0.014)	(0.009)
31–34	-0.017^{**}	-0.001	-0.011
	(0.008)	(0.012)	(0.019)
35–40	-0.018^{**}	-0.011	-0.009
	(0.007)	(0.010)	(0.029)
41–49	-0.008	0.023^{***}	-0.006
	(0.006)	(0.007)	(0.040)
50+	-0.004	0.013^{**}	-0.045
	(0.006)	(0.006)	(0.043)
Occupations		787–880	

Mechanism Evidence: Codified versus Tacit Knowledge Why might AI disproportionately affect entry-level workers? One possible mechanism is that AI substitutes more effectively for codified knowledge—formal, standardized, documented knowledge that can be taught through education, textbooks, or written procedures—while complementing tacit knowledge acquired through practice,

mentorship, and repeated exposure to real situations. As a simple descriptive exercise to evaluate this mechanism, we construct indices for codified and tacit knowledge requirements by occupation. We proxy an occupation's reliance on codified knowledge with its required level of formal education, supplemented by O*NET knowledge domains and work activities that reflect formally taught, standardized content (e.g., mathematics, law, analyzing data). We proxy reliance on tacit knowledge with required work experience and on-the-job training, supplemented by knowledge domains and activities that reflect experiential learning (e.g., mechanical knowledge, resolving conflicts, coaching others). Online Appendix Section EC.9 details the construction.

Results are in Online Appendix Figure EC.75. Occupations with higher codified knowledge have slower entry-level employment growth, while occupations with higher tacit knowledge have faster employment growth for mid-career and senior workers. These raw gradients are directionally consistent with the hypothesis that generative AI substitutes for codified knowledge and complements tacit knowledge, in line with recent theoretical literature studying potential effects of AI on the labor market (**??**). We caution that these estimates are non-causal and purely suggestive, and more empirical work may help assess the importance of alternative mechanisms in the data.[18]

2.6. Fact 6: The divergence is not explained by several prominent alternative factors

The divergence for young workers persists when excluding technology firms and computer occupations, when controlling for occupational exposure to interest-rate increases and for remote-work amenability, and across five alternative measures of AI exposure. Education is the one control that attenuates these estimates; we treat it separately in Section **??**. These are descriptive robustness results: they help clarify which alternatives are more or less likely to meaningfully explain the patterns we document.

Excluding Technology Occupations One possibility is that our results are explained by a general slowdown in technology hiring from 2022 to 2023 as firms recovered from the COVID-19 Pandemic. Online Appendix, Fig. EC.20 shows employment changes by age and exposure quintile after excluding computer occupations, corresponding to 2010 SOC codes that start with 15-1. Online Appendix, Fig. EC.21 shows results when excluding firms in information technology or computer systems design (NAICS codes 51 and 5415). Results are quite similar, consistent with the above case studies showing employment changes across various occupations. The occupation-level regression estimate for 22-25 year-olds is essentially unchanged when dropping computer occupations and only

[18] The codified-knowledge gradient is not significant when controlling for college share quintile, demonstrating its overlap with education. The tacit-knowledge gradient for experienced workers, by contrast, survives this control.

slightly smaller when dropping technology firms (Online Appendix Table EC.12). These results indicate that our findings are not specific to technology roles.[19]

Occupational Interest Rate Exposure Rapid advances in generative AI in late 2022 were closely preceded by increases in interest rates in March 2022. This raises interest rate changes as a potential alternative explanation for the patterns we document (**??**). We use occupational interest rate exposure data from Zens et al. (2020) to test for the correlation between interest rate exposure and AI exposure.[20] Online Appendix, Fig. EC.17 in fact finds that AI exposure and interest rate exposure are negatively correlated overall, with occupations such as construction having high interest rate exposure and low AI exposure. Online Appendix, Fig. EC.18 and EC.19 repeat our analysis separately for occupations below and above the median in interest rate exposure. Both cases are consistent with our main results, suggesting that occupational differences in interest rate exposure do not drive our results. Panel B of Table **??** shows that controlling for interest rate exposure does not meaningfully change the magnitude or precision of the estimated employment decline for 22-25 year-olds in the most exposed occupations (-0.178 versus -0.179 with no controls).[21] Using an alternative approach, **?** similarly finds a limited role for interest rate exposure in explaining the divergence in employment for young workers.

Remote Work Recent studies suggest that remote work may harm the labor market prospects of young workers and help to explain the slowdown in entry-level hiring (Lambert and Schindler 2026, **?**). Online Appendix, Fig. EC.24 and EC.25 show results for occupations amenable to remote work (telework) and those that are not, according to Dingel and Neiman (2020).[22] We find that, for young workers, more exposed occupations have slower employment growth, both in teleworkable occupations and in non-teleworkable occupations. The results for non-teleworkable occupations in particular suggest that our findings are not driven by outsourcing or work-from-home disruptions, at least solely.

[19] A related channel is the 2022 amendment to Internal Revenue Code §174 under the Tax Cuts and Jobs Act, which disallowed the immediate deduction of R&D expenditures, including software development costs, and may have restrained hiring in technology-intensive occupations. The robustness of our results to excluding technology firms and computer occupations also speaks against this channel.

[20] Thank you to Zanna Iscenko and Fabien Curto-Millet for the suggestion and to Zens et al. (2020) for generously sharing their data.

[21] The Zens et al. (2020) measure is estimated from pre-2020 monetary policy shocks and loads most heavily on goods-sector channels, so it may imperfectly capture the incidence of the 2022 rate increases on exposed knowledge-work occupations.

[22] Whether an occupation is amenable to remote work is positively correlated with AI exposure. Only two teleworkable occupations fall in the lowest quintile of estimated AI exposure according to the GPT-4 β measure. Likewise few non-teleworkable occupations fall in the highest quintile of AI exposure. For this reason in Online Appendix, Fig. EC.24 we pool together the two lowest AI exposure quintiles into one group. In Online Appendix, Fig. EC.25 we pool together the two highest AI exposure quintiles.

We also consider an alternative measure of work-from-home amenability from **?** and Lambert and Schindler (2026), based on the share of job postings offering remote or hybrid work for an occupation. Panels D and E in Table **??** show that the general descriptive patterns persist in magnitude when adjusting for work-from-home, whether using the Lambert and Schindler (2026) remote-work measure (Panel D) or the Dingel and Neiman (2020) teleworkability classification (Panel E). These results contrast with findings from Lambert and Schindler (2026). We discuss this further in Section **??**. [23]

Alternative AI Exposure Measures A number of studies seek to estimate occupational exposure to generative AI. In Online Appendix Section EC.4 we repeat the analysis using five alternatives: suitability for machine learning (Brynjolfsson et al. 2018), the AI Occupational Exposure index and its language-modeling and image-generation variants (Felten et al. 2021, 2023), and the AI applicability score derived from Bing Copilot conversations (Tomlinson et al. 2025). These results are in Online Appendix, Fig. EC.49 to EC.50. The results are qualitatively similar across the different measures: for workers aged 22-25, employment falls in the most exposed quintiles, by 9 to 14% relative to November 2022, while the least exposed quintiles grow. As with the measure from Eloundou et al. (2024), differences by exposure quintile for other age groups are more muted.

Other Robustness Checks Online Appendix, Fig. EC.27 shows results when including part-time and temporary workers. Results are qualitatively similar.[24]

Firm Entry and Exit The preceding results consider a balanced sample of firms. This removes entry and exit of firms as a margin of adjustment. While this removes trends that may be specific to growth in the ADP platform, it may be that it masks patterns in how workers sort across incumbent and exiting firms by seniority and AI exposure. Online Appendix, Fig. EC.26 shows that results are similar when we do not take a balanced sample of firms, suggesting that holding the set of firms fixed is not central to the documented trends. However, as noted in Online Appendix EC.1.2, ADP tends to skew towards larger firms than the overall US economy. One possibility is that some of the changes in exposed entry level jobs come from reallocation across new, smaller businesses that are not well-represented in ADP. New business applications have been substantially higher

[23] We thank Lambert and Schindler (2026) for generously sharing their work-from-home data. Teleworkable jobs with relatively low AI exposure include work that requires human interaction, such as marriage and family therapists, mental health/substance abuse social workers, sales and financial managers, and lawyers. Non-teleworkable jobs with relatively high AI exposure include tellers, office and file clerks, bookkeeping and auditing clerks, and receptionists and information clerks.

[24] Another possibility is that employment trends reflect COVID-19 stimulus checks distorting labor supply. However, these payments were income-conditioned, and AI-exposed occupations average higher incomes (**?**), making this channel unlikely in explaining our findings.

than pre-pandemic levels (**?**). A very small, though economically consequential, share of these new businesses are venture-backed startups. Recent evidence suggests AI-native startups in particular may be leaner, more technical, and more senior (**??**). Further work is needed to understand how the trends documented in this paper interact with the growth of new businesses and startups.

2.7. Fact 7: Adjustment is occurring through employment rather than base compensation

In addition to employment we observe workers' annual base compensation, which we use to test for labor market adjustment along the compensation margin. An important limitation is that base salary excludes bonuses, overtime pay, commissions, equity, and tips—components that are largest in precisely the most exposed, high-income occupations, and through which compensation adjustment could occur unobserved. Salary data are deflated to 2017 dollars using the PCE index.

Results for various occupations are in Online Appendix, Fig. EC.62. The findings indicate less divergence in compensation compared to employment across more and less exposed occupations. Online Appendix, Fig. EC.63 shows results by age and Eloundou et al. (2024)-based exposure quintile. We find little difference in compensation trends by age or exposure quintile.

This average comparison could nonetheless be distorted by composition: if hiring of young exposed workers contracts, the remaining pool skews toward higher-tenure, higher-paid incumbents. We consider two subsets of workers to explore this possibility. First, among job-stayers—workers continuously employed at the same firm—average real base pay has grown more slowly in more exposed jobs, though this is not specific to young workers. These differences are small relative to the roughly 20-percentage-point divergence in employment (Online Appendix, Fig. EC.64). Second, among new hires, real starting pay shows no obvious relationship with AI exposure for young workers. It rose modestly faster for older hires in exposed jobs, the opposite pattern as for stayers (Online Appendix, Fig. EC.65). These pay differences by exposure are comparatively small in both groups, so the adjustment documented in Fact 4 operates through employment rather than base compensation.

Prior work by **?** notes that jobs are bundles of tasks and technology that replaces tasks that do not require much expertise may reduce occupational employment but increase occupational wages. Conversely, technology that replaces expert tasks may decrease wages. The sign of the wage effect depends not only on the share of tasks displaced but also whether these tasks are expert or inexpert. The limited changes we find for wages suggest that these effects may be offsetting, at least in the

short run. Alternatively, the results could be explained by wage stickiness in the short run, consistent with recent evidence from **??**.[25]

3. Alternative Explanations and Pre-Trends

3.1. Changes in Education

Another possibility is that our results stem from worsening education outcomes during the COVID-19 Pandemic, as more educated workers have greater average AI exposure (**?**).[26] In Online Appendix, Fig. EC.22 we show trends for occupations in which greater than 70% of workers have a college degree according to the 2017 American Community Survey (ACS).[27] In Online Appendix, Fig. EC.23 we show trends for occupations in which fewer than 30% of workers have a college degree.

Occupations with a high share of college graduates have declining employment overall, with muted differences between more-exposed and less-exposed occupations compared to our main results.[28] In contrast, occupations with a low share of college graduates have rising overall employment, with the least AI-exposed occupations growing and the most exposed occupations declining in employment. Further, for lower college share occupations, the dispersion in employment outcomes is visible in higher age groups as well, with workers up to age 30 showing separation in employment trends by AI exposure. These findings suggest that deteriorating education outcomes cannot fully explain our main results. They also suggest that for non-college workers, experience may serve as less of a buffer to labor market disruption than for college workers.

Occupational college share is correlated with AI exposure—no occupation in the least-exposed quintile has a college share above 35%—and education is plausibly a channel through which AI affects work rather than only a confounder, since generative AI substitutes best for precisely the codified knowledge taught through formal schooling. Estimates with and without the education control therefore bracket a range. In our primary sample, which balances firms since 2018 and permits assessing pre-trends, the most-exposed-quintile estimate is -0.18 with no controls and

[25] Perhaps one slight exception to the muted compensation patterns is software developers, for whom compensation for 22-25 year-olds grew somewhat faster than for older workers in the period after ChatGPT, though the gap shrinks toward the end of the sample (Online Appendix, Fig. EC.62). Appendix Fig. EC.63 indicates this is an idiosyncratic pattern for software developers: aggregating to exposure quintiles, the most exposed young workers do not see faster compensation growth on average. For software developers specifically, rising relative pay may partly reflect composition effects if firms primarily contracted hiring of lower-paid engineers, or negative supply shocks that simultaneously raised entry-level wages and reduced hiring.

[26] **?** finds that declines in average skill levels for college graduates explain a sizable share of the slowdown in the growth of the college wage gap in recent decades.

[27] Not a single occupation in the first quintile of GPT-4 β based exposure measure has a college share above 35%, so that quintile is excluded from the results in Online Appendix, Fig. EC.22. In addition, we pool together quintiles 2 and 3 into one group since quintile 2 has few occupations that require a college degree.

[28] See **?** for a comparison to the CPS and discussions about outcome measurement.

-0.09 with college share added (Table **??**, Panels A and C), the latter remaining statistically significant at the 10 percent level; the former is preferred if college share mostly proxies exposure to AI-automatable knowledge work, the latter if it proxies independent shocks to educated labor markets, such as pandemic-era disruptions to skill formation. The range is similar in the sample balanced since 2021 (-0.19 with no controls and -0.06 with college share, though the latter is no longer statistically significant; Table **??**).

3.2. Pre-Trends

Figure **??** shows differential trends before ChatGPT: the most exposed occupations boomed relative to the least exposed during the pandemic, peaking around mid-2020, and declined relative to them thereafter. Two quantifications discipline interpretation. First, by November 2022 the relative position of the most exposed occupations had returned approximately to its 2018–2019 level; the subsequent decline of roughly 18 percentage points therefore takes the gap well below its pre-pandemic baseline rather than back toward it. Decomposing the full swing, the relative gap rose about 13 percentage points from its 2018–2019 baseline to its mid-2020 peak and has since fallen about 30 percentage points; of that decline, roughly 13 points unwind the pandemic boom back to baseline and the remaining 17 points carry the gap below it. Second, a reversion of the pandemic-era boom would predict the relative gap to stabilize once it returned to baseline, which it did by late 2022; instead the decline continued for more than three years—including through the most recent year of data and over a period in which interest rates stabilized and then declined—carrying the gap well below its pre-pandemic level. Nonetheless, occupations with high AI exposure differ from others in many ways—they are more likely to be high-income knowledge work—and **?** document deterioration in AI-exposed jobs predating ChatGPT; these patterns cannot be read as causal effects of AI. The continued trends do, however, suggest weighing structural changes that specifically affect young workers over transitory or business-cycle explanations. Candidate explanations include AI, with COVID-19 education disruptions and work-from-home as other plausible possibilities.

3.3. Reconciling Remote Work Estimates

Our findings contrast with estimates from Lambert and Schindler (2026), who find that accounting for work-from-home significantly attenuates the correlation between AI exposure and employment changes. These differences may stem from alternative variable definitions and analyses or differing data sets. In Online Appendix Section EC.8, Fig. EC.74 and Table EC.16, we follow the empirical specification in their work as closely as possible, finding that AI exposure continues to predict a

decline in the junior share of new hires after accounting for work-from-home. This suggests that the results vary because of differences across the data sets, with Lambert and Schindler (2026) using data from online professional social media profiles and job postings. These analyses are descriptive prediction exercises and not causal. Future work may seek to further separate the causal impacts of AI and work-from-home on work given their correlated exposure. These channels may interact, with prior work suggesting that firms that adopted remote work were subsequently more likely to adopt generative AI (**??**).

3.4. Within-Firm Estimates

The results so far could be explained by industry- or firm-level shocks correlated with sorting patterns by age and measured AI exposure. For example, one possibility is that young workers with high measured AI exposure are disproportionately likely to sort to firms heavily susceptible to interest rate increases.

We test for this class of confounders with Poisson event study regressions of within-firm employment by exposure quintile, controlling for firm-time and firm-quintile fixed effects (Online Appendix Section EC.3.8). The firm-time effects absorb aggregate firm shocks that impact each exposure quintile equally; the firm-quintile effects adjust for baseline differences in hiring across quintiles within the firm. In Online Appendix Section EC.3.8 we estimate the model across different samples that vary the stringency for firm size. For workers aged 22-25, changes are directionally consistent with employment declines in the two most exposed quintiles, though precision varies across samples. Estimates for older age groups are small and statistically indistinguishable from zero.[29] Alternative confounders that would not be controlled for with firm-time effects are occupation characteristics unrelated to AI that could lead to changes in hiring even conditional on the firm.

4. External Validity: Comparison to National Benchmarks

A useful benchmark for our findings is to compare them to estimates from the monthly Current Population Survey (CPS). In Online Appendix, Section EC.7 we perform similar analyses in the CPS, finding high levels of noise consistent with small sample sizes in fine age-occupation cells. Online Appendix Section EC.7 compares our results to the 2024 American Community Survey (ACS). Online Appendix Figure EC.73 shows that, among the youngest workers, the two most exposed quintiles show the weakest employment growth by 2024, but the trends are less pronounced than in

[29] The firm fixed-effects results were more sensitive to changes we made in the data pipeline since prior versions of this study. Online Appendix Section EC.11 quantifies the sensitivity.

the ADP data. Online Appendix Table EC.14 shows that the gap between the most and least exposed quintiles is -0.022 [-0.055, $+0.011$] in the ACS, compared to -0.132 in the ADP sample when averaging its monthly series to a comparable annual (2022-to-2024) change.[30] This comparison uses the 2022–2024 window over which the ACS is available; the ADP gap has continued to widen since, reaching about -0.205 by mid-2026. Restricting the ACS sample to a population similar to ADP—full-time, civilian, wage and salary workers—changes the gap between Q1 and Q5 only slightly, to -0.019 [-0.059, $+0.020$], so differences in worker coverage account for little of the discrepancy.

Online Appendix Table EC.15 compares the exposure gap across broad industry sectors.[31] Within professional, information, and financial services, the two sources agree closely: the gap between Q1 and Q5 is -0.231 in ADP and -0.213 [-0.353, -0.072] in the ACS. The discrepancy between them is concentrated in education, health, and public administration, where ADP shows a small gap of -0.074 while the ACS shows a positive gap of $+0.245$ [0.104, 0.386].

Some of the aggregate discrepancy reflects uncertainty in the ACS itself. The Census Bureau has documented substantial revisions to the 2024 ACS population weights (U.S. Census Bureau 2025); recomputing the overall comparison using each cell's average weight from 2022 instead moves the ACS gap from -0.022 to -0.042. The 2024 ACS weight revisions are largest in education, health, and public administration.

Taken together, the magnitude of the divergence appears specific to the ADP sample, but its direction is consistent. The two agree for white collar occupations. The sign of the age and exposure pattern is corroborated by descriptive evidence in representative US administrative data by age and industry (**?**). These magnitudes merit further investigation and urge caution against extrapolating them to the economy as a whole; a particular concern is that changes in the ADP sample might correspond to reallocation to firms outside the sample in equilibrium. The ideal data set for studying how these trends generalize would be microdata on the universe of employment, measured with high frequency and with age and occupation codes.

5. Additional Heterogeneity

5.1. Gender

Women work in more AI-exposed occupations than men at every age, in the ADP data and in the CPS and ACS alike. Online Appendix Fig. EC.10 and Table EC.10 show each gender's employment

[30] Confidence intervals are computed using the 80 Census successive-difference replicate weights.

[31] Industries were aggregated into five broad sectors so that each ACS cell has at least 500 survey respondents.

across exposure quintiles. This suggests that AI could have differential effects on the labor market by gender. Even within exposure quintiles, we find notable occupational differences between men and women, as shown in Online Appendix Table EC.11. Within the least exposed occupations, men concentrate in manual labor roles while women concentrate in cleaning and healthcare roles. Within the most exposed quintile, men have greater concentration in technical work such as software development and information systems management. Online Appendix, Fig. EC.28 and EC.29 show results separately for men and women. The results are qualitatively similar, though we find steeper average declines for young women than young men consistent with greater concentration in declining high-exposure roles. We caution that these average trends need not exactly match results in national statistics given the greater occupational exposure on average in the ADP data, especially for women.

5.2. Regional Variation

We also examine whether the patterns vary across the country. Online Appendix Section EC.3.10 shows descriptive employment trends by Census region (Fig. EC.41–EC.44). The relative decline for the most exposed young workers appears in every region but is largest in the West, where employment for the most exposed quintile falls about 21% from its November 2022 level, and smallest in the South (about -8%). Online Appendix Table EC.13 reports the corresponding employment-weighted occupation-level regression estimates, which are consistent with the patterns in the figures.

5.3. AI Adoption

The analyses above compare occupations by measured AI exposure. A complementary approach is to ask whether employment changes concentrate in markets that are more likely to adopt AI. A direct firm-level adoption analysis is not feasible with the current ADP data because firm identifiers are anonymized, which precludes merging external adoption measures based on job postings or earnings calls (**?**).

As a step toward looking at where adoption is most intense, we use state-level relative usage data from the Anthropic Economic Index (February 2026 release). Following their documentation, we divide each state's relative usage by the size of its labor force to construct a state-level adoption index, and group states into tiers (leading, upper-middle, lower-middle, and emerging). We then recompute our descriptive employment statistics within each tier. As shown in Online Appendix Section EC.3.11 (Fig. EC.45–EC.48), the decline in employment for exposed young workers is

sharpest in the leading-adoption states (about -19% for the most exposed quintile) and is smallest in the emerging-adoption states. Ongoing and future work may help assess the effect of AI adoption on employment changes (**????**).

6. Discussion

We document a set of facts about the recent labor market effects of artificial intelligence using payroll data from ADP. Overall, we do not find widespread job displacement in occupations most exposed to generative AI. However, employment has declined substantially for early-career workers in the occupations most exposed to AI, including but not limited to software development and customer support. While economy-wide employment continues to grow, employment growth for young workers overall has been stagnant. These declines are concentrated in applications of AI that *automate* rather than *complement* work, while faster growth for experienced workers is associated with tacit knowledge; the parallel codified-knowledge gradient for young workers is difficult to separate from occupational education. Labor-market adjustment has been more visible in employment than in compensation. The patterns are broadly similar after controlling for occupational interest rate exposure and remote-work amenability; after conditioning on firm-time shocks; and when excluding technology-oriented industries or occupations. They attenuate when adjusting for education.

While our main estimates may be influenced by factors other than generative AI, our results are consistent with the hypothesis that generative AI has begun to affect entry-level employment.

The adoption of new technologies typically leads to heterogeneous effects across workers, resulting in adjustment periods as workers reallocate from displaced forms of work to new forms with growing labor demand (**?**). Past transitions such as the IT revolution ultimately led to robust growth in employment and real wages following physical and human capital adjustments, with some workers benefiting more than others (**??**).

The strongest form of complementarity operates through the creation of new products, services, and tasks in which human labor has comparative advantage (**???**). Our usage-based measures capture how AI is applied to existing tasks and will only partially reflect this margin; tracking the emergence of new kinds of work is an important direction for the AI Economic Indicators and for future research.

Tracking employment trends on an ongoing basis will help determine if the adjustment to AI follows a similar pattern with ultimately rising employment and compensation. Our public

AI Economic Indicators will continue monitoring these outcomes to assess whether the trends documented in the paper persist in the future.

Acknowledgments

We thank David Autor, Sarah Bana, Eric Bergman, Nick Bloom, Cody Cook, Diane Coyle, Fabien Curto Millet, Chris Forman, Joshua Gans, Martha Gimbel, Nathan Goldschlag, Basil Halperin, Anders Humlum, Zanna Iscenko, Molly Kinder, Peter John Lambert, Christina Langer, Fei-Fei Li, Frank Li, Omeed Maghzian, James Manyika, Adam Ozimek, Jiaxin Pei, Nela Richardson, Daniel Rock, Brad Ross, Phil Trammell, Andrew Wang, and participants at the Stanford Digital Economy Lab workshop and numerous seminars for helpful feedback. We are grateful to ADP for access to the data and the Stanford Digital Economy Lab for financial support. All errors are our own.

Data Availability. The data used in this study are proprietary payroll records provided by ADP under a data-use agreement and cannot be made publicly available. Derived, occupation-level results underlying the paper's figures are published and updated monthly through the public AI Economic Indicators (`https://digitaleconomy.stanford.edu/project/indicators/canaries-dashboard/`), where they can be downloaded. Code for all analyses, and the occupation-level crosswalks linking AI-exposure measures to ADP occupations, are available from the authors upon request.

References

Acemoglu D, Autor D (2011) Skills, Tasks and Technologies: Implications for Employment and Earnings*. Card D, Ashenfelter O, eds., *Handbook of Labor Economics*, volume 4, 1043–1171 (Elsevier), URL `http://dx.doi.org/10.1016/S0169-7218(11)02410-5`.

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

ADP Research (2025) ADP National Employment Report. URL `http://adpemploymentreport.com`.

Aldasoro I, Gambacorta L, Pal R, Revoltella D, Weiss C, Wolski M (2026) AI adoption, productivity and employment: evidence from European firms. Working Paper 1325, Bank for International Settlements, URL `https://www.bis.org/publ/work1325.pdf`.

Appel R, Massenkoff M, McCrory P, McCain M, Heller R, Neylon T, Tamkin A (2026) Anthropic economic index report: Economic primitives. URL `https://www.anthropic.com/research/anthropic-economic-index-january-2026-report`.

Appel R, McCrory P, Tamkin A, Stern M, McCain M, Neylon T (2025) Anthropic economic index report: Uneven geographic and enterprise AI adoption. URL `https://www.anthropic.com/research/anthropic-economic-index-september-2025-report`.

Autor D, Chin C, Salomons A, Seegmiller B (2024) New Frontiers: The Origins and Content of New Work, 1940–2018*. *The Quarterly Journal of Economics* 139(3):1399–1465, ISSN 0033-5533, URL http://dx.doi.org/10.1093/qje/qjae008.

Autor D, Thompson N (2025) Expertise. Working Paper, NBER, URL http://dx.doi.org/10.3386/w33941.

Azar J, Gine M, Sanz-Espín J (2025) AI Is Already Eroding Wages: Quasi-Experimental Evidence From Occupational Exposure. URL http://dx.doi.org/10.2139/ssrn.5842084.

Bick A, Blandin A, Deming DJ (2024) The Rapid Adoption of Generative AI. Working Paper, National Bureau of Economic Research, URL http://dx.doi.org/10.3386/w32966.

Bonney K, Breaux CL, Dinlersoz E, Foster LS, Haltiwanger JC, Pande AA (2026) The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks. URL http://dx.doi.org/10.3386/w35141.

Bresnahan TF, Brynjolfsson E, Hitt LM (2002) Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence*. *The Quarterly Journal of Economics* 117(1):339–376, ISSN 0033-5533, URL http://dx.doi.org/10.1162/003355302753399526.

Brynjolfsson E (2022) The turing trap: The promise and peril of human-like artificial intelligence. *Daedalus* 151(2):272–287, URL http://dx.doi.org/https://doi.org/10.1162/daed_a_01915.

Brynjolfsson E, Chandar B, Chen R (2026) Canaries, Interest Rates, and Timing: More on the Recent Drivers of Employment Changes for Young Workers. Stanford Digital Economy Lab, URL https://digitaleconomy.stanford.edu/news/canaries-interest-rates-and-timinga-more-on-recent-drivers-of-employment-changes-for-y

Brynjolfsson E, Mitchell T, Rock D (2018) What Can Machines Learn, and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings* 108:43–47, ISSN 2574-0768, URL http://dx.doi.org/10.1257/pandp.20181019.

Brynjolfsson E, Rock D, Syverson C (2021) The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *American Economic Journal: Macroeconomics* 13(1):333–372, ISSN 1945-7707, URL http://dx.doi.org/10.1257/mac.20180386.

Cajner T, Crane L, Decker R, Hamins-Puertolas A, Kurz C, Radler T (2018) Using Payroll Processor Microdata to Measure Aggregate Labor Market Activity. Technical report, The Federal Reserve Board of Governors, URL https://www.federalreserve.gov/econres/feds/using-payroll-processor-microdata-to-measure-aggregate-labor-market-activity.htm.

Chandar B (2025a) Shifts in the Composition of College Workers: Implications for US Inequality and Labor Demand. Technical report, Social Science Research Network, URL https://ssrn.com/abstract=5384528.

Chandar B (2025b) Tracking Employment in AI-Exposed Jobs. Technical report, Social Science Research Network, URL https://ssrn.com/abstract=5384519.

Chandar B, Klein Teeselink B (2026) AI and labor around the world, working paper.

Chen F, Stratton J (2026) Artificial Intelligence in the Firm. Working Paper, Harvard University, URL https://fion.ac/jellyfish.pdf, preliminary working paper.

Chen WX, Srinivasan S, Zakerinia S (2025) Displacement or Complementarity? The Labor Market Impact of Generative AI.

Crane LD, Soto PE (2026) AI and Coder Employment: Compiling the Evidence. URL http://dx.doi.org/10.17016/FEDS.2026.018.

Davis SJ, Krolikowski PM (2025) Sticky Wages on the Layoff Margin. *American Economic Review* 115(2):491–524, ISSN 0002-8282, URL http://dx.doi.org/10.1257/aer.20240309.

Deming D (2026) Remote work containerized the office, paving the way for AI. URL https://forklightning.substack.com/p/remote-work-containerized-the-office, substack newsletter.

Dingel JI, Neiman B (2020) How many jobs can be done at home? *Journal of Public Economics* 189:104235, ISSN 0047-2727, URL http://dx.doi.org/10.1016/j.jpubeco.2020.104235.

Dominski J, Lee YS (2025) Advancing AI Capabilities and Evolving Labor Outcomes. Technical report, arXiv, URL http://dx.doi.org/10.48550/arXiv.2507.08244, arXiv:2507.08244 [econ].

Eckhardt S, Goldschlag N (2025) AI and Jobs: The Final Word (Until the Next One). *Economic Innovation Group* URL https://eig.org/ai-and-jobs-the-final-word/.

Eloundou T, Manning S, Mishkin P, Rock D (2024) GPTs are GPTs: Labor market impact potential of LLMs. *Science* 384(6702):1306–1308, URL http://dx.doi.org/10.1126/science.adj0998, publisher: American Association for the Advancement of Science.

Emanuel N, Harrington E, Pallais A (2026) The Power of Proximity to Coworkers. *The Quarterly Journal of Economics* URL http://dx.doi.org/10.1093/qje/qjag027, advance article.

Federal Reserve Bank of New York (2025) The Labor Market for Recent College Graduates. URL https://www.newyorkfed.org/research/college-labor-market#--:explore:unemployment.

Felten E, Raj M, Seamans R (2021) Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal* 42(12):2195–2217, ISSN 1097-0266, URL http://dx.doi.org/10.1002/smj.3286, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/smj.3286.

Felten E, Raj M, Seamans R (2023) How will Language Modelers like ChatGPT Affect Occupations and Industries? Technical report, arXiv, URL http://dx.doi.org/10.48550/arXiv.2303.01157, arXiv:2303.01157 [econ].

Frank MR, Sabet AJ, Simon L, Bana SH, Yu R (2026) AI-exposed jobs deteriorated before ChatGPT. URL http://dx.doi.org/10.48550/arXiv.2601.02554, arXiv:2601.02554 [econ.GN].

Garicano L (2025) The AI Becker problem. URL https://www.siliconcontinent.com/p/the-ai-becker-problem.

Garicano L, Rayo L (2025) DP20634 Training in the Age of AI: A Theory of Apprenticeship Viability. CEPR Discussion Paper 20634, Paris & London, URL https://cepr.org/publications/dp20634.

Gimbel M, Kinder M, Kendall J, Lee M (2025) Evaluating the Impact of AI on the Labor Market: Current State of Affairs | The Budget Lab at Yale. URL https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs.

Gupta A, Qian F, Simintzi E, Sun Y (2026) Generative AI and entrepreneurship, working paper, April 14, 2026. Available at SSRN: https://ssrn.com/abstract=5971474.

Hampole M, Papanikolaou D, Schmidt LD, Seegmiller B (2025) Artificial Intelligence and the Labor Market. Working Paper, National Bureau of Economic Research, URL http://dx.doi.org/10.3386/w33509.

Handa K, Tamkin A, McCain M, Huang S, Durmus E, Heck S, Mueller J, Hong J, Ritchie S, Belonax T, Troy KK, Amodei D, Kaplan J, Clark J, Ganguli D (2025) Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations. Technical report, arXiv, URL http://dx.doi.org/10.48550/arXiv.2503.04761, arXiv:2503.04761 [cs].

Hansen S, Lambert PJ, Bloom N, Davis SJ, Sadun R, Taska B (2023) Remote Work across Jobs, Companies, and Space. Working Paper 31007, National Bureau of Economic Research, URL http://dx.doi.org/10.3386/w31007.

Hartley J, Jolevski F, Melo V, Moore B (2025) The Labor Market Effects of Generative Artificial Intelligence. SSRN Scholarly Paper, Social Science Research Network, Rochester, NY, URL http://dx.doi.org/10.2139/ssrn.5136877.

Hazell J, Taska B (2025) Downward rigidity in the wage for new hires. *American Economic Review* 115(12):4183–4217.

Hosseini Maasoum SM, Lichtinger G (2025) Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data. URL http://dx.doi.org/10.2139/ssrn.5425555.

Humlum A, Vestergaard E (2025) Large Language Models, Small Labor Market Effects. Working Paper, National Bureau of Economic Research, URL http://dx.doi.org/10.3386/w33777.

Hyman B, von Wachter T, Herrera D, Moghadam R, Morris J, Motghare S, Segal K (2026) Tracking AI-related job loss using unemployment insurance claims data in California. Technical report, California Policy Lab, URL https://capolicylab.org/wp-content/uploads/2026/06/Tracking-AI-Related-Job-Loss-Using-Unemployment-Insurance-Claims-Data-in-California.pdf.

Ide E (2025) Automation, ai, and the intergenerational transmission of knowledge. *arXiv preprint arXiv:2507.16078* .

Iscenko Z, Curto Millet F (2026) Looking for the Ladder: Is AI Impacting Entry-Level Jobs? Working Paper, Economic Innovation Group, URL https://eig.org/wp-content/uploads/2026/01/TAWP-Iscenko-Millet.pdf.

Johnston A, Makridis C (2025) The Labor Market Effects of Generative AI: A Difference-in-Differences Analysis of AI Exposure. SSRN Scholarly Paper, Social Science Research Network, Rochester, NY, URL http://dx.doi.org/10.2139/ssrn.5375017.

Kharazian A, Simon L, Stevens R (2026) A New Look at AI's Impact on Jobs: Firm-Level AI Spending and Workforce Adjustment. Ramp Economics Lab Working Paper ramp-ai-jobs-2026, Ramp Economics Lab, URL https://ramp.com/data/ai-jobs-impact.

Kim H, Koning R (2026) AI-Native Firms. SSRN Scholarly Paper, Social Science Research Network, Rochester, NY, URL http://dx.doi.org/10.2139/ssrn.6905079.

Klein Teeselink B (2025) Generative AI and Labor Market Outcomes: Evidence from the United Kingdom. URL https://papers.ssrn.com/abstract=5516798.

Kochhar R (2023) Which U.S. Workers Are More Exposed to AI on Their Jobs? *Pew Research Center* URL https://www.pewresearch.org/social-trends/2023/07/26/which-u-s-workers-are-more-exposed-to-ai-on-their-jobs/.

Kuhfeld M, Lewis K (2025) 5 years after COVID-19 hit: Test data converge on math gains, stalled reading recovery. *Brookings* URL https://www.brookings.edu/articles/5-years-after-covid-19-hit-test-data-converge-on-math-gains-stalled-reading-recovery/

Lahart J, Chen TP (2025) Young Graduates Are Facing an Employment Crisis. *WSJ* URL https://www.wsj.com/economy/jobs/jobs-unemployment-rise-young-people-ce4704d8, section: Economy.

Lambert PJ, Schindler Y (2026) The Broken Ladder: AI, Remote Work, and Early-Career Hiring. URL http://dx.doi.org/10.2139/ssrn.6787638.

Liu Y, Wang H, Yu S (2025) Labor demand in the age of generative ai: Early evidence from the us job posting data. *World Bank Working Paper* .

Maslej N, Fattorini L, Perrault R, Gil Y, Parli V, Kariuki N, Capstick E, Reuel A, Brynjolfsson E, Etchemendy J, et al. (2025) Artificial intelligence index report 2025. *arXiv preprint arXiv:2504.07139* .

Massenkoff M, Lyubich E, McCrory P, Appel R, Heller R (2026) Anthropic economic index report: Learning curves. URL https://www.anthropic.com/research/economic-index-march-2026-report.

Massenkoff M, McCrory P (2026) Labor market impacts of AI: A new measure and early evidence. Report, Anthropic, URL https://cdn.sanity.io/files/4zrzovbb/website/2b5bbaf2c1eb81dbf6e6fb813c1a24e35a64d376.pdf.

Morris C (2025) Anthropic ceo warns ai could eliminate half of all entry-level white-collar jobs. *Fortune* URL https://fortune.com/2025/05/28/anthropic-ceo-warning-ai-job-loss/.

Ozimek A, Goldschlag N (2026) AI and Young-adult Jobs: The Real Mystery. URL https://eig.org/ai-and-young-adult-jobs-the-real-mystery/.

Patwardhan T, Dias R, Proehl E, Kim G, Wang M, Watkins O, Fishman SP, Aljubeh M, Thacker P, Fauconnet L, Kim NS, Chao P, Miserendino S, Chabot G, Li D, Sharman M, Barr A, Glaese A, Tworek J (2025) Gdpval: Evaluating ai model performance on real-world economically valuable tasks. *arXiv preprint* URL http://dx.doi.org/10.48550/arXiv.2510.04374.

Rodgers WM III, Kassens AL (2026) It's (Still) the Business Cycle: Young Adult Workers in a "Low-Hire, Low-Fire" Labor Market. URL https://www.stlouisfed.org/on-the-economy/2026/jun/its-still-business-cycle-young-adult-workers-low-hire-low-fire-labor-market, federal Reserve Bank of St. Louis, On the Economy blog; first post in a three-part series.

Schubert G (2025) Organizational technology ladders: Remote work and generative AI adoption. SSRN Working Paper 5094265, Social Science Research Network, Rochester, NY, URL http://dx.doi.org/10.2139/ssrn.5094265.

Stevens R (2026) Payrolls to Prompts: Firm-Level Evidence on the Substitution of Labor for AI. URL http://dx.doi.org/10.48550/arXiv.2602.00139, arXiv:2602.00139 [econ].

Tomlinson K, Jaffe S, Wang W, Counts S, Suri S (2025) Working with AI: Measuring the Applicability of Generative AI to Occupations. URL http://dx.doi.org/10.48550/arXiv.2507.07935, arXiv:2507.07935 [cs].

Tucker LC (2026) You're (not) Hired: Artificial Intelligence and Early Career Hiring in the Quarterly Workforce Indicators. Working Paper CES-26-27, U.S. Census Bureau, Center for Economic Studies, URL https://www.census.gov/library/working-papers/2026/adrm/CES-WP-26-27.html.

US Census Bureau (2025) Guidance for comparing the 2024 ACS 1-year estimates. American community survey research note, U.S. Census Bureau, URL https://www.census.gov/content/dam/Census/programs-surveys/acs/tech-doc/user-notes/2024-acs-1-year-research-note.pdf.

US Census Bureau (2026) Business formation statistics. https://www.census.gov/econ/bfs/index.html, monthly data release, June 2026 (released July 9, 2026). Accessed July 20, 2026.

Wiles E, Horton JJ (2025) Generative AI and Labor Market Matching Efficiency. URL http://dx.doi.org/10.2139/ssrn.5187344.

Yotzov I, Barrero JM, Bloom N, Bunn P, Davis SJ, Foster KM, Jalca A, Meyer BH, Mizen P, Navarrete MA, Smietanka P, Thwaites G, Wang BZ (2026) Firm Data on AI. Technical Report w34836, National Bureau of Economic Research, URL http://dx.doi.org/10.3386/w34836.

Zens G, Böck M, Zörner TO (2020) The heterogeneous impact of monetary policy on the US labor market. *Journal of Economic Dynamics and Control* 119:103989, ISSN 0165-1889, URL http://dx.doi.org/10.1016/j.jedc.2020.103989.