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

Erik Brynjolfsson* Bharat Chandar† Ruyu Chen‡

August 2026

For Latest Version, **Click Here**

Abstract

Using a sample of high-frequency administrative payroll data from ADP covering millions of U.S. workers through June 2026, we document six 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) Adjustment is occurring through employment rather than base compensation. 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. These patterns attenuate when controlling for education, show some divergent trends predating generative AI, and are more pronounced in the ADP analysis sample 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.

*Stanford University and NBER, erikb@stanford.edu.
†Stanford University, chandarb@stanford.edu.
‡Stanford University, ruyuchen@stanford.edu.

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% (Maslej et al., 2025). 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 (Patwardhan et al., 2025). At the same time adoption rates among U.S. workers have approached 50% (Hartley et al., 2025; Bick et al., 2024). 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 (Morris, 2025).[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 a sample of 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

[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 L 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.

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 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 July 2025 data vintage and has since widened to 19% as of June 2026.

In sum, we document six 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; and (6) 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.[3] 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 4). Second, estimates attenuate

[3] They are also robust to allowing for firm entry and exit, though the clientele in the ADP analysis sample may not fully reflect new business formation.

when controlling for occupational education levels, which may reflect either an alternative explanation or the very channel through which AI operates (Section 4). Third, the divergence is directionally consistent but more pronounced in the ADP panel than in national survey benchmarks (Section 5). 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 AI usage for automation 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 (Tucker, 2026). 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, 2025b; Gimbel et al., 2025; Eckhardt and Goldschlag, 2025; Hosseini Maasoum and Lichtinger, 2025; Klein Teeselink, 2025; Johnston and Makridis, 2025; Humlum and Vestergaard, 2025; Frank et al., 2026; Yotzov et al., 2026; Tucker, 2026; Crane and Soto, 2026; Massenkoff and McCrory, 2026; Bonney et al., 2026; Lambert and Schindler, 2026; Kharazian et al., 2026).[4]

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

[4]See also Dominski and Lee (2025); Eckhardt and Goldschlag (2025); Chen et al. (2025); Liu et al. (2025); Azar et al. (2025); Iscenko and Curto Millet (2026); Brynjolfsson et al. (2026); Aldasoro et al. (2026); Chen and Stratton (2026); Stevens (2026).

media records, though we caution that there may also be selection in which firms use ADP.[5] 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 (Acemoglu and Autor, 2011). Consequently, AI may be automating the checkable, process-intensive tasks that historically justified entry-level headcount, while increasing the leverage of experienced staff (Ide, 2025; Garicano and Rayo, 2025). 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.[6]

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

[5]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 (Wiles and Horton, 2025).

[6]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 Chandar and Klein Teeselink (2026), 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.

the facts in this paper on an ongoing basis.[7]

1 Data

1.1 Payroll Data

We analyze a sample of anonymized administrative payroll data from ADP, providing a near real-time view of a meaningful part of the US labor market. ADP provides payroll services to firms employing over 26 million U.S. workers. Our analysis sample is a subset of these records: a balanced panel of firms observed monthly from January 2021 through June 2026, comprising between 3.5 and 5 million employees per month. Throughout, the "ADP analysis sample" refers to this panel rather than to ADP's full client base. To examine an extended pre-ChatGPT period we also consider a 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 observed in the ADP sample.

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 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 A.1. We find that the descriptive facts are largely insensitive to the imputation.

In Online Appendix A.2 we also report comparisons between the ADP analysis sample and the broader US economy.[8] The analysis 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.[9] Online Appendix Fig. A.10 shows that, compared to the ACS and CPS, the ADP analysis 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 analysis sample, but

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

[8] See also Cajner et al. (2018); ADP Research (2025).

[9] Online Appendix Fig. A.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.

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 A and in Section 5. [10]

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.[11] These measures change over time as usage evolves. 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 F.2.[12]

In Online Appendix Section E 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 A.2 through A.6 list the largest occupations by ADP employment in each quintile.

For robustness analyses, we classify occupations as teleworkable using classifications from Dingel

[10] 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.

[11] 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.

[12] 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.

and Neiman (2020). We also use data on remote work by occupation from Lambert and Schindler (2026).[13] 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 analysis sample. Figure 1 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, 2025b; Eckhardt and Goldschlag, 2025; Gimbel et al., 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.

Figure 2 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 (Lahart and Chen, 2025; Federal Reserve Bank of New York, 2025).[14]

Fig. 3 shows employment growth interacting age with AI exposure quintile. The top left plot

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

[14] We compare employment trends in the ADP analysis sample to the CPS and ACS in Section 5.

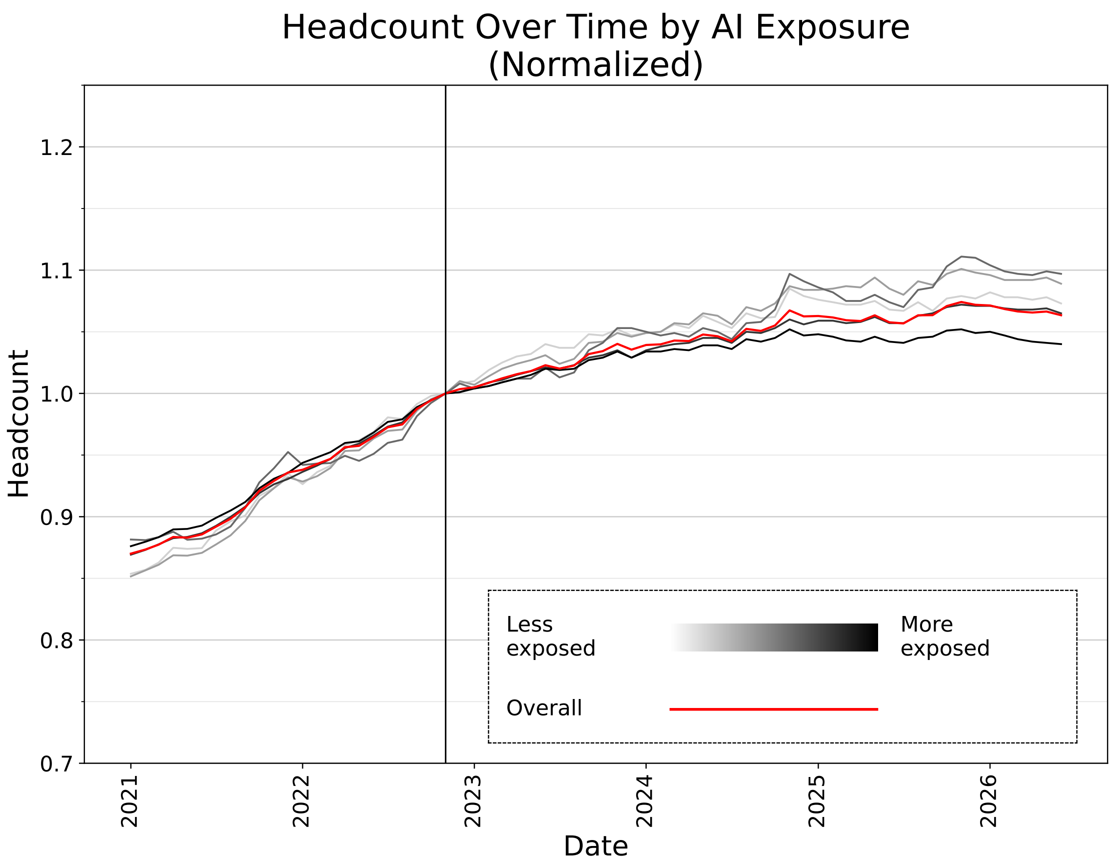

Figure 1: Employment changes by exposure quintile using measures from Eloundou et al. (2024), pooling across all age groups. Exposure quintiles are defined based on the GPT-4 β measures.

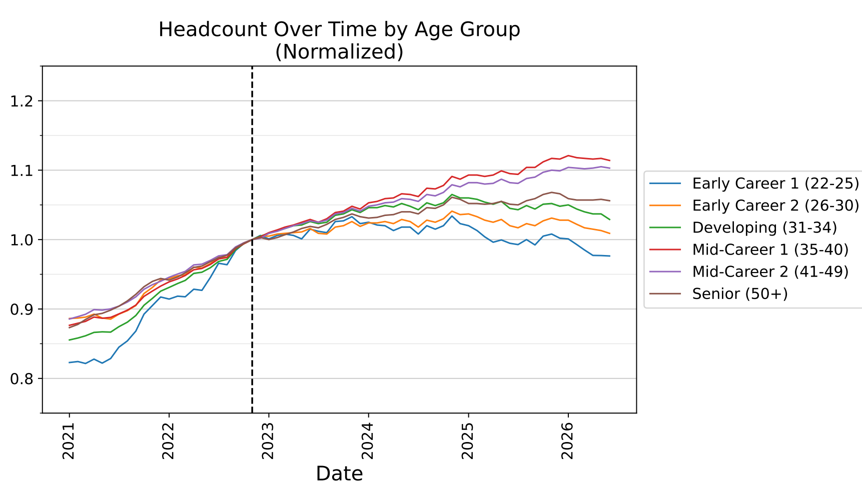

Figure 2: Employment changes by age in the ADP sample, normalized to 1 in November 2022.

shows a divergence in employment outcomes for more and less exposed occupations for workers aged 22–25, with more exposed occupations experiencing declining employment.[15] For older age groups, we find much less marked differences in employment growth across AI exposure quintiles. Online Appendix, Fig. B.3 through B.5 show case studies of employment changes for specific occupations, such as software development and customer service.

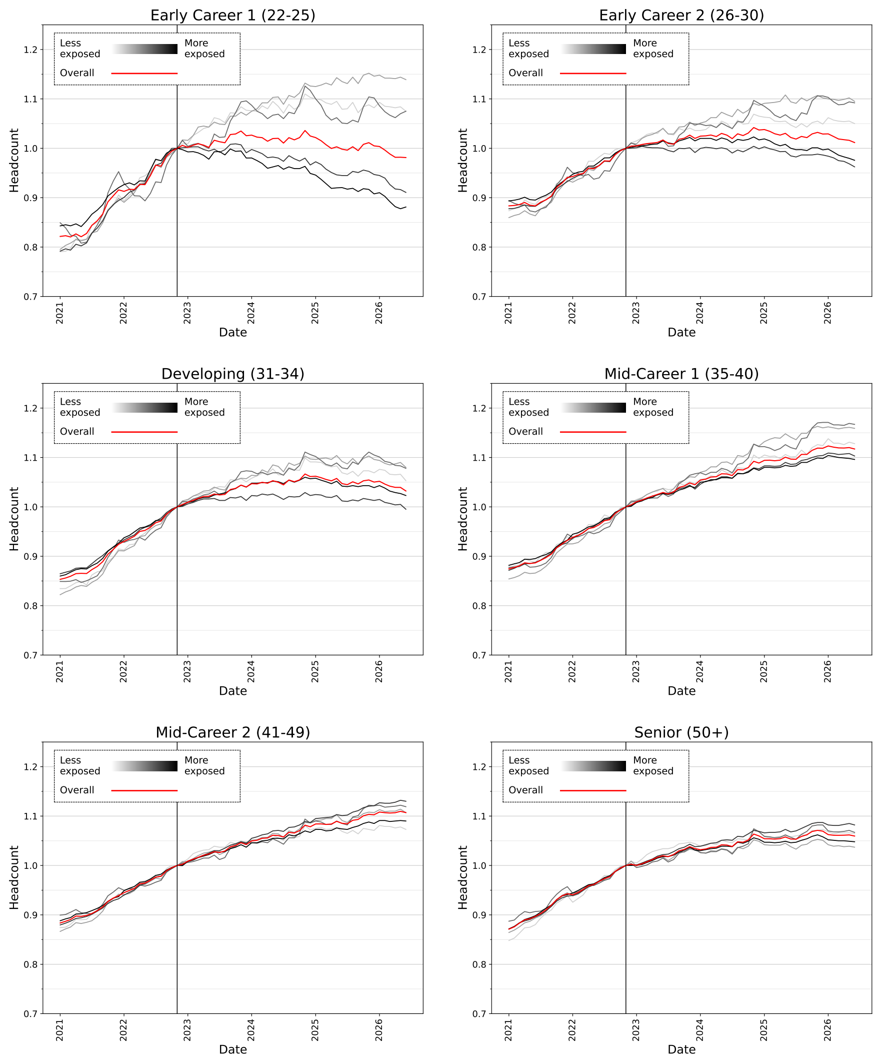

Figure 3: GPT-4 β. Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Exposure quintiles are defined based on the GPT-4 β measures. Darker lines are more exposed quintiles. The red line shows the overall trend pooling across quintiles.

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

[15] Online Appendix Figure B.1 shows a similar pattern for 18 to 21 year-olds, though the sample size is smaller and the employment changes are more volatile.

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. B.2 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. B.6). 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.[16] Table 1 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 2 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

[16] Clustering at the three-digit SOC level (about 90 groups) does not change the results qualitatively. The clustered standard errors are similar to the heteroskedasticity-robust standard errors or slightly narrower with additional controls.

work-from-home measure of Lambert and Schindler (2026), and teleworkability (Dingel and Neiman, 2020).

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. Section 3 summarizes the robustness controls in Panels B, D, and E; Section 4 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 5). 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 4 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 C.7. The prior versions of this study documented this divergence, and it has continued to widen since: by the kept-pace measure emphasized above, the shortfall was 15% at the July 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 3 and with evidence in other studies (Frank et al., 2026; Iscenko and Curto Millet, 2026; Tucker, 2026). Section 4 quantifies these pre-trends and discusses their implications for interpretation.

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, which permits assessing pre-trends (Figure 4); this is our primary sample. The panels differ only 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.

	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 1). The panels differ only 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.

	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

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

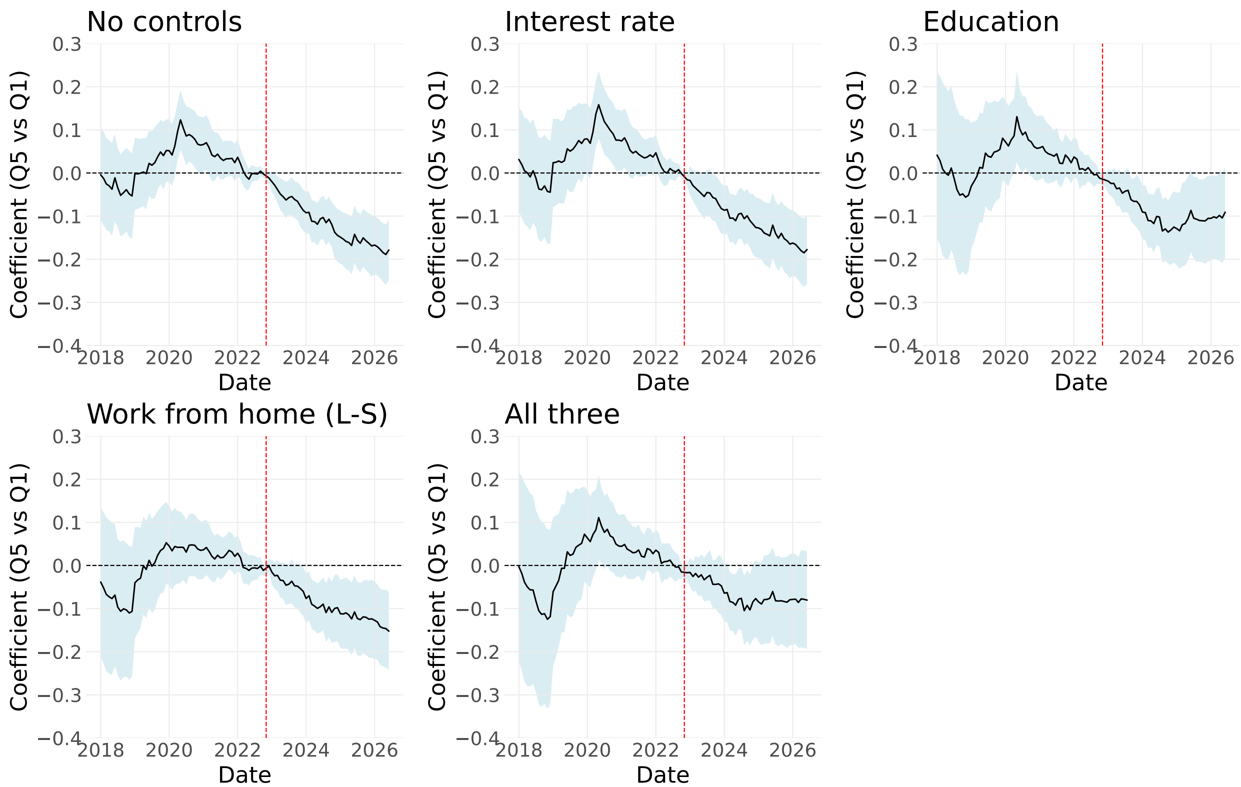

Figure 4: Employment-weighted occupation event study estimates for workers aged 22–25, most exposed quintile relative to quintile 1, tracing the occupation-level long-difference coefficient as the difference endpoint rolls forward over time. Estimated on the extended balanced panel of firms observed from January 2018 as in Table 1. The first panel has no controls; each of the next three adds a single set of occupational controls to that baseline, and the final panel includes all three together. 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 K gives the definitions.

Online Appendix Figure B.7 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 (Rodgers and Kassens, 2026). 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 Hosseini Maasoum and Lichtinger (2025) and Tucker (2026), who find decreased entry into entry-level roles, and Hyman et al. (2026), who find suggestive evidence of higher rates of layoffs for AI-exposed college-educated workers in California.[17] We caution that these quantities are determined in equilibrium, with changes in firms' hiring potentially a function of workers' exit rates and vice versa.[18]

[17] 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 Tucker (2026).

[18] A general slowdown in hiring could specifically impact AI-exposed young workers more because their entry into the labor market is most dependent on hiring (Rodgers and Kassens, 2026). A higher share of exposed workers than less exposed workers initially enter the labor market at age 22 or later because exposed jobs are more likely to require a college degree. However, changes in hiring rates by age and occupation are endogenous with these structural factors, with the evidence suggesting that firms in our sample have contracted hiring for exposed young workers beyond the rate at which their stock is structurally replenished. This is not the case for other types of workers. In addition, we find a similar divergence for jobs that do not require a college degree (Online Appendix Figure C.7) and for 18 to 21

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 (Brynjolfsson, 2022).

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 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. Online Appendix, Tables A.8 and A.9 show example occupations that are in the highest and lowest exposure category for each measure. Figure 5 shows results from the March 2025 release of the index, with results for later periods in Online Appendix Section F.2. Panel A of Fig. 5 shows employment changes by overall prevalence of related Claude queries for 22–25 year-olds.[19] 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 automation-oriented uses of AI substituting for labor while complementary uses are associated with flat or rising employment.[20] Online Appendix Section F.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

year-olds (Online Appendix Figure B.1), groups for which hiring rates are structurally more similar in more and less exposed jobs.

[19] See Online Appendix, Fig. F.1 through F.3 for other age groups.

[20] 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. F.4 and F.5 show that automation and augmentation results are similar when dropping occupations with low overall Claude usage.

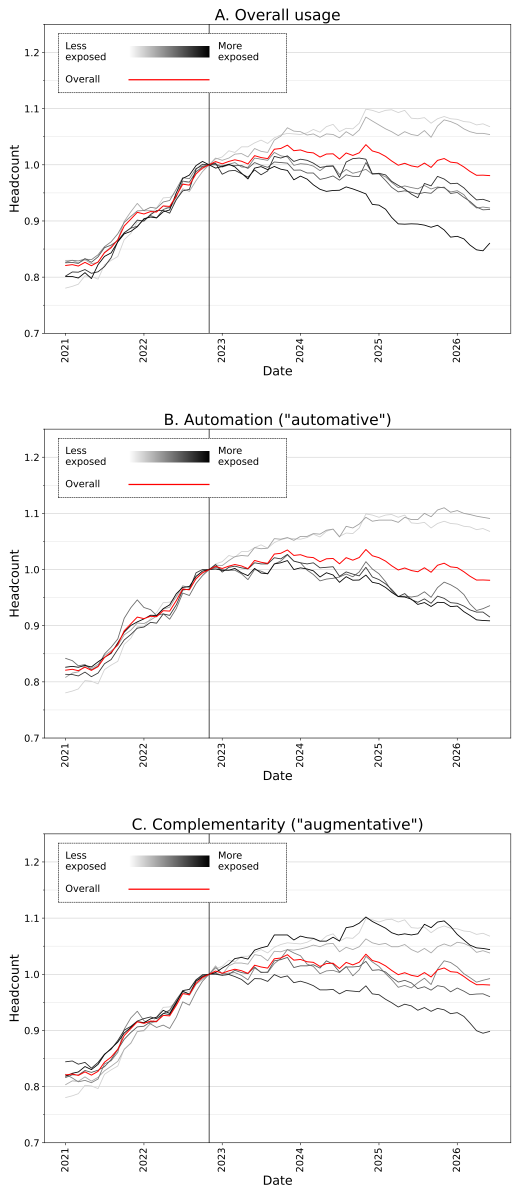

Figure 5: *Automation and Complementarity.* Employment changes by age and exposure quintile using measures from Handa et al. (2025). Panel A: *Overall usage.* Exposure quintiles are defined based on the share of queries to Claude that relate to tasks associated with an occupation, scaled by BLS employment 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. 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: *Complementarity.* Complementarity quintiles are defined based on the share of queries related to an occupation that are classified by Claude as "augmentative".

overall usage and with automation exposure. The complementarity gradient is mixed.

To distinguish the three usage measures, which are correlated across occupations, Table 3 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 (+0.024) and 50+ in both releases. This age gradient—automation associated with declines for the young, 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).

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 J details the construction.

Results are in Online Appendix Figure J.1. 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 (Ide, 2025; Garicano, 2025). We caution that these estimates are non-causal and purely

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	-0.098^{***}	0.016	-0.029
	(0.018)	(0.021)	(0.018)
26–30	-0.036^{***}	0.010	0.000
	(0.013)	(0.019)	(0.018)
31–34	-0.017	0.014	-0.017
	(0.011)	(0.015)	(0.021)
35–40	-0.014	0.004	-0.023
	(0.011)	(0.014)	(0.027)
41–49	-0.008	0.024^{**}	-0.020
	(0.010)	(0.010)	(0.025)
50+	-0.006	0.015^{*}	-0.037^{**}
	(0.009)	(0.008)	(0.016)
Panel B: Pooled September 2025, January 2026, and April 2026 AEI releases			
22–25	-0.084^{***}	-0.010	-0.023
	(0.014)	(0.019)	(0.016)
26–30	-0.041^{***}	-0.002	-0.005
	(0.011)	(0.015)	(0.010)
31–34	-0.018^{*}	-0.001	-0.012
	(0.009)	(0.013)	(0.024)
35–40	-0.018^{**}	-0.012	-0.010
	(0.008)	(0.010)	(0.037)
41–49	-0.008	0.024^{***}	-0.006
	(0.008)	(0.008)	(0.061)
50+	-0.004	0.014^{**}	-0.051
	(0.007)	(0.006)	(0.073)
Occupations		621–675	

suggestive, and more empirical work may help assess the importance of alternative mechanisms in the data.[21]

2.6 Fact 6: 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. G.1. The findings indicate less divergence in compensation compared to employment across more and less exposed occupations. Online Appendix, Fig. G.2 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. G.3). 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. G.4). These pay differences by exposure are comparatively small in both groups, so the adjustment documented in Fact 2 operates through employment rather than base compensation.

Prior work by Autor and Thompson (2025) 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

[21] 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.

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 Davis and Krolikowski (2025); Hazell and Taska (2025).[22]

3 Robustness

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 4. 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. C.4 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. C.5 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 slightly smaller when dropping technology firms (Online Appendix Table C.1). These results indicate that our findings are not specific to technology roles.[23]

[22]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. G.1). Appendix Fig. G.2 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.

[23]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.

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 (Iscenko and Curto Millet, 2026; Brynjolfsson et al., 2026). We use occupational interest rate exposure data from Zens et al. (2020) to test for the correlation between interest rate exposure and AI exposure.[24] Online Appendix, Fig. C.1 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. C.2 and C.3 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 1 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).[25] Using an alternative approach, Tucker (2026) 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; Emanuel et al., 2026). Online Appendix, Fig. C.8 and C.9 show results for occupations amenable to remote work (telework) and those that are not, according to Dingel and Neiman (2020).[26] 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.

We also consider an alternative measure of work-from-home amenability from Hansen et al. (2023) and Lambert and Schindler (2026), based on the share of job postings offering remote or

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

[25]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.

[26]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. C.8 we pool together the two lowest AI exposure quintiles into one group. In Online Appendix, Fig. C.9 we pool together the two highest AI exposure quintiles.

hybrid work for an occupation. Panels D and E in Table 1 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 4.3. [27]

Alternative AI Exposure Measures A number of studies seek to estimate occupational exposure to generative AI. In Online Appendix Section E 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. E.1 to E.5. The results are qualitatively similar across the different measures: for workers aged 22–25, employment falls in the most exposed quintile, by 10 to 13% 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.

Part-time and Temporary Workers Online Appendix, Fig. C.11 shows results when including part-time and temporary workers. Results are qualitatively similar.[28]

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. C.10 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 A.2, ADP tends to skew towards larger firms than the

[27]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.

[28]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 (Kochhar, 2023), making this channel unlikely in explaining our findings.

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 than pre-pandemic levels (U.S. Census Bureau, 2026). 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 (Gupta et al., 2026; Kim and Koning, 2026). Further work is needed to understand how the trends documented in this paper interact with the growth of new businesses and startups.

4 Alternative Explanations and Pre-Trends

4.1 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 (Kuhfeld and Lewis, 2025).[29] In Online Appendix, Fig. C.6 we show trends for occupations in which greater than 70% of workers have a college degree according to the 2017 American Community Survey (ACS).[30] In Online Appendix, Fig. C.7 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.[31] 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.

[29] Chandar (2025a) 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.

[30] 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. C.6. In addition, we pool together quintiles 2 and 3 into one group since quintile 2 has few occupations that require a college degree.

[31] See Ozimek and Goldschlag (2026) for a comparison to the CPS and discussions about outcome measurement.

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 -0.09 with college share added (Table 1, 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 2).

4.2 Pre-Trends

Figure 4 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 Frank et al. (2026) 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 contributors.

4.3 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 I, Fig. I.1 and Table I.1, 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.[32]

These analyses comparing AI and work-from-home 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 (Schubert, 2025; Deming, 2026).

4.4 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 C.6). The firm-time effects absorb aggregate firm shocks that impact each exposure

[32] One other difference between the studies is that the comparable estimate in Lambert and Schindler (2026) measures the junior hiring share outcome using a role-specific seniority variable coded by Revelio Labs. Our replication instead uses age to measure the junior hiring share. This difference in outcome variable is another possible source of varying results.

quintile equally; the firm-quintile effects adjust for baseline differences in hiring across quintiles within the firm. In Online Appendix Section C.6 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.[33] 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.

5 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 H 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 H compares our results to the 2024 American Community Survey (ACS). Online Appendix Figure H.8 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 the ADP analysis sample. Online Appendix Table H.1 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.[34] 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 H.2 compares the exposure gap across broad industry sectors.[35] 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

[33]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 L quantifies the sensitivity.

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

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

−0.074 while the ACS shows a positive gap of +0.245 [0.104, 0.386].

However, we caution that some of the aggregate discrepancy may reflect weighting changes in the ACS. The Census Bureau has documented substantial revisions to the population controls underlying the 2024 ACS weights, driven by an improved methodology for estimating net international migration (U.S. Census Bureau, 2025a). In the guidance issued alongside these estimates, the Census Bureau recommends that "whenever possible" users "compare population characteristics such as percentages, means, medians, and rates rather than estimates of population totals" (U.S. Census Bureau, 2025b). This poses challenges for comparing the results in ADP to the ACS because the primary outcome measured in this study is the change in employment totals across years. As a supplementary diagnostic, we also report the ratio of unweighted respondent counts. Doing so moves the ACS gap from −0.022 to −0.047, and from −0.019 to −0.042 on the sample restricted to full-time civilian wage and salary workers. Online Appendix Table H.3 reports the corresponding breakdown by sector. Weighting changes impact the education, health, and public administration sector in particular, with Q1 employment growth 16 percentage points higher when unweighted and the Q1 to Q5 gap shrinking by 13 percentage points.

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 (Tucker, 2026). 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.

6 Additional Heterogeneity

6.1 Gender

Women work in more AI-exposed occupations than men at every age, in the ADP analysis sample and in the CPS and ACS alike. Online Appendix Fig. A.10 and Table D.1 show each gender's employment 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 D.2. 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. D.1 and D.2 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 analysis sample, especially for women.

6.2 Regional Variation

We also examine whether the patterns vary across the country. Online Appendix Section D.2 shows descriptive employment trends by Census region (Fig. D.4–D.7). The relative decline for the most exposed young workers appears in every region but is greatest in the West and Northeast, where employment for the most exposed quintile falls about 21% and 13% from its November 2022 level, and smallest in the South and Midwest (about −7% and −9%). Online Appendix Table D.3 reports the corresponding employment-weighted occupation-level regression estimates, which are consistent with the patterns in the figures.

6.3 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 (Hampole et al., 2025).

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 D.3 (Fig. D.8–D.11), 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 (Humlum and Vestergaard, 2025; Hosseini Maasoum and Lichtinger, 2025; Kharazian et al., 2026; Chandar and Klein Teeselink, 2026).

7 Discussion

We document a set of facts about the recent labor market effects of artificial intelligence using a sample of 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. Faster growth for experienced workers is associated with greater tacit knowledge; slower growth for young workers is correlated with greater codified knowledge. 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 (Autor et al., 2024). 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 (Bresnahan et al., 2002; Brynjolfsson et al., 2021).

The strongest form of complementarity operates through the creation of new products, services, and tasks in which human labor has comparative advantage (Acemoglu and Restrepo, 2019; Brynjolfsson, 2022; Autor et al., 2024). 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, Peter McCrory, 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.

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.

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 is available from the authors upon request.

References

Acemoglu, D. and D. Autor (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings*," in *Handbook of Labor Economics*, ed. by D. Card and O. Ashenfelter, Elsevier, vol. 4, 1043–1171. (document)

Acemoglu, D. and P. Restrepo (2019): "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33, 3–30. 7

ADP Research (2025): "ADP National Employment Report," . 8, 37

Aldasoro, I., L. Gambacorta, R. Pal, D. Revoltella, C. Weiss, and M. Wolski (2026): "AI adoption, productivity and employment: evidence from European firms," Working Paper 1325, Bank for International Settlements. 4

Appel, R., M. Massenkoff, P. McCrory, M. McCain, R. Heller, T. Neylon, and A. Tamkin (2026): "Anthropic Economic Index report: Economic primitives," . 2.5, A.3, A.9, F.2

Appel, R., P. McCrory, A. Tamkin, M. Stern, M. McCain, and T. Neylon (2025): "Anthropic Economic Index report: Uneven geographic and enterprise AI adoption," . 2.5, A.3, A.9, D.3, F.2

Autor, D., C. Chin, A. Salomons, and B. Seegmiller (2024): "New Frontiers: The Origins and Content of New Work, 1940–2018*," *The Quarterly Journal of Economics*, 139, 1399–1465. 7

Autor, D. and N. Thompson (2025): "Expertise," Working Paper, NBER. 2.6

Autor, D. H. (2015): "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, 29, 3–30. C.1

Azar, J., M. Gine, and J. Sanz-Espín (2025): "AI Is Already Eroding Wages: Quasi-Experimental Evidence From Occupational Exposure," . 4

Bick, A., A. Blandin, and D. J. Deming (2024): "The Rapid Adoption of Generative AI," Working Paper, National Bureau of Economic Research. (document)

Bonney, K., C. L. Breaux, E. Dinlersoz, L. S. Foster, J. C. Haltiwanger, and A. A. Pande (2026): "The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks," . (document)

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002): "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence*," *The Quarterly Journal of Economics*, 117, 339–376. 7

Brynjolfsson, E. (2022): "The Turing Trap: The Promise and Peril of Human-Like Artificial Intelligence," *Daedalus*, 151, 272–287. 2.5, 7

Brynjolfsson, E., B. Chandar, and R. Chen (2026): "Canaries, Interest Rates, and Timing: More on the Recent Drivers of Employment Changes for Young Workers," Stanford Digital Economy Lab. 4, 3

Brynjolfsson, E., T. Mitchell, and D. Rock (2018): "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings*, 108, 43–47. 1.2, 3, E.1

Brynjolfsson, E., D. Rock, and C. Syverson (2021): "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics*, 13, 333–372. 7

Cajner, T., L. Crane, R. Decker, A. Hamins-Puertolas, C. Kurz, and T. Radler (2018): "Using Payroll Processor Microdata to Measure Aggregate Labor Market Activity," Tech. rep., The Federal Reserve Board of Governors. 8, 37

Chandar, B. (2025a): "Shifts in the Composition of College Workers: Implications for US Inequality and Labor Demand," Tech. rep., Social Science Research Network. 29

——— (2025b): "Tracking Employment in AI-Exposed Jobs," Tech. rep., Social Science Research Network. (document), 2.1, H

Chandar, B. and B. Klein Teeselink (2026): "AI and Labor Around the World," Working paper. 6, 6.3

Chen, F. and J. Stratton (2026): "Artificial Intelligence in the Firm," Working Paper, Harvard University, preliminary working paper. 4

Chen, J. and J. Roth (2024): "Logs with Zeros? Some Problems and Solutions*," *The Quarterly Journal of Economics*, 139, 891–936. 41

Chen, W. X., S. Srinivasan, and S. Zakerinia (2025): "Displacement or Complementarity? The Labor Market Impact of Generative AI," . 4

Crane, L. D. and P. E. Soto (2026): "AI and Coder Employment: Compiling the Evidence," . (document)

Davis, S. J. and P. M. Krolikowski (2025): "Sticky Wages on the Layoff Margin," *American Economic Review*, 115, 491–524. 2.6

Deming, D. (2026): "Remote Work Containerized the Office, Paving the Way for AI," Substack newsletter. 4.3

Dingel, J. I. and B. Neiman (2020): "How many jobs can be done at home?" *Journal of Public Economics*, 189, 104235. 1.2, 2.2, 1, 2, 3, A.4, C.4, C.8, C.9

Dominski, J. and Y. S. Lee (2025): "Advancing AI Capabilities and Evolving Labor Outcomes," Tech. rep., arXiv, arXiv:2507.08244 [econ]. 4, H

Eckhardt, S. and N. Goldschlag (2025): "AI and Jobs: The Final Word (Until the Next One)," *Economic Innovation Group*. (document), 4, 2.1, H

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024): "GPTs are GPTs: Labor market impact potential of LLMs," *Science*, 384, 1306–1308, publisher: American Association for the Advancement of Science. 1.2, 2.1, 1, 3, 2.5, 2.6, 3, A.3, A.4, A.10, A.3, B.1, B.5, B.6, B.7, C.1, C.2, C.3, C.4, C.5, C.6, C.7, C.8, C.9, C.10, C.11, C.6, C.12, C.13, C.14, C.15, D.1, D.2, D.1, D.2, G.2, H.4, I, I.1

Emanuel, N., E. Harrington, and A. Pallais (2026): "The Power of Proximity to Coworkers," *The Quarterly Journal of Economics*, advance article. 3

Federal Reserve Bank of New York (2025): "The Labor Market for Recent College Graduates," . 2.2

Felten, E., M. Raj, and R. Seamans (2021): "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses," *Strategic Management Journal*, 42, 2195–2217, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/smj.3286. 1.2, 3, E.2

——— (2023): "How will Language Modelers like ChatGPT Affect Occupations and Industries?" Tech. rep., arXiv, arXiv:2303.01157 [econ]. 1.2, 3, E.3, E.4

Frank, M. R., A. J. Sabet, L. Simon, S. H. Bana, and R. Yu (2026): "AI-exposed jobs deteriorated before ChatGPT," ArXiv:2601.02554 [econ.GN]. (document), 2.3, 4.2

Garicano, L. (2025): "The AI Becker problem," . 2.5

Garicano, L. and L. Rayo (2025): "DP20634 Training in the Age of AI: A Theory of Apprenticeship Viability," CEPR Discussion Paper 20634, Paris & London. (document)

Gimbel, M., M. Kinder, J. Kendall, and M. Lee (2025): "Evaluating the Impact of AI on the Labor Market: Current State of Affairs | The Budget Lab at Yale," . (document), 2.1

Gupta, A., F. Qian, E. Simintzi, and Y. Sun (2026): "Generative AI and Entrepreneurship," Working paper, April 14, 2026. Available at SSRN: https://ssrn.com/abstract=5971474. 3

Hampole, M., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller (2025): "Artificial Intelligence and the Labor Market," Working Paper, National Bureau of Economic Research. 6.3

Handa, K., A. Tamkin, M. McCain, S. Huang, E. Durmus, S. Heck, J. Mueller, J. Hong, S. Ritchie, T. Belonax, K. K. Troy, D. Amodei, J. Kaplan, J. Clark, and D. Ganguli (2025): "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations," Tech. rep., arXiv, arXiv:2503.04761 [cs]. 1.2, 11, 2.5, 5, 3, A.3, 40, A.7, A.8, F.1, F.2, F.3, F.4, F.5

Hansen, S., P. J. Lambert, N. Bloom, S. J. Davis, R. Sadun, and B. Taska (2023): "Remote Work across Jobs, Companies, and Space," Working Paper 31007, National Bureau of Economic Research. 3

Hartley, J., F. Jolevski, V. Melo, and B. Moore (2025): "The Labor Market Effects of Generative Artificial Intelligence," SSRN Scholarly Paper, Social Science Research Network, Rochester, NY. (document)

Hazell, J. and B. Taska (2025): "Downward Rigidity in the Wage for New Hires," *American Economic Review*, 115, 4183–4217. 2.6

Hosseini Maasoum, S. M. and G. Lichtinger (2025): "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data," . (document), 2.4, 6.3

Humlum, A. and E. Vestergaard (2025): "Large Language Models, Small Labor Market Effects," Working Paper, National Bureau of Economic Research. (document), 6.3

Hyman, B., T. von Wachter, D. Herrera, R. Moghadam, J. Morris, S. Motghare, and K. Segal (2026): "Tracking AI-Related Job Loss Using Unemployment Insurance Claims Data in California," Tech. rep., California Policy Lab. 2.4

Ide, E. (2025): "Automation, AI, and the Intergenerational Transmission of Knowledge," *arXiv preprint arXiv:2507.16078*. (document), 2.5

Iscenko, Z. and F. Curto Millet (2026): "Looking for the Ladder: Is AI Impacting Entry-Level Jobs?" Working Paper, Economic Innovation Group. 4, 2.3, 3

Johnston, A. and C. Makridis (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. (document)

Kharazian, A., L. Simon, and R. Stevens (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. (document), 6.3

Kim, H. and R. Koning (2026): "AI-Native Firms," SSRN Scholarly Paper, Social Science Research Network, Rochester, NY. 3

Klein Teeselink, B. (2025): "Generative AI and Labor Market Outcomes: Evidence from the United Kingdom," . (document)

Kochhar, R. (2023): "Which U.S. Workers Are More Exposed to AI on Their Jobs?" *Pew Research Center*. 28

Kuhfeld, M. and K. Lewis (2025): "5 years after COVID-19 hit: Test data converge on math gains, stalled reading recovery," *Brookings*. 4.1

Lahart, J. and T.-P. Chen (2025): "Young Graduates Are Facing an Employment Crisis," *WSJ*, section: Economy. 2.2

Lambert, P. J. and Y. Schindler (2026): "The Broken Ladder: AI, Remote Work, and Early-Career Hiring," . (document), 1.2, 2.2, 2.3, 1, 2, 3, 27, 4.3, 32, A.4, I, I, I.1, I.1

Lim, S., D. Strauss, J. Burn-Murdoch, and C. Murray (2025): "Is AI killing graduate jobs?" *Financial Times*. H

Liu, Y., H. Wang, and S. Yu (2025): "Labor Demand in the Age of Generative AI: Early Evidence from the US Job Posting Data," *World Bank Working Paper*. 4

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

Massenkoff, M., E. Lyubich, P. McCrory, R. Appel, and R. Heller (2026): "Anthropic Economic Index report: Learning curves," . 2.5, A.3, A.9, F.2

Massenkoff, M. and P. McCrory (2026): "Labor market impacts of AI: A new measure and early evidence," Report, Anthropic. (document)

Morris, C. (2025): "Anthropic CEO warns AI could eliminate half of all entry-level white-collar jobs," *Fortune*. (document)

O'Brien, C. (2025): "A viral chart on recent graduate unemployment is misleading," *Agglomerations*. H

Ozimek, A. and N. Goldschlag (2026): "AI and Young-adult Jobs: The Real Mystery," . 31

Patwardhan, T., R. Dias, E. Proehl, G. Kim, M. Wang, O. Watkins, S. P. Fishman, M. Aljubeh, P. Thacker, L. Fauconnet, N. S. Kim, P. Chao, S. Miserendino, G. Chabot, D. Li, M. Sharman, A. Barr, A. Glaese, and J. Tworek (2025): "GDPval: Evaluating AI Model Performance on Real-World Economically Valuable Tasks," *arXiv preprint*. (document)

Rodgers, III, W. M. and A. L. Kassens (2026): "It's (Still) the Business Cycle: Young Adult Workers in a "Low-Hire, Low-Fire" Labor Market," Federal Reserve Bank of St. Louis, On the Economy blog; first post in a three-part series. 2.4, 18

Schubert, G. (2025): "Organizational Technology Ladders: Remote Work and Generative AI Adoption," SSRN Working Paper 5094265, Social Science Research Network, Rochester, NY. 4.3

Stevens, R. (2026): "Payrolls to Prompts: Firm-Level Evidence on the Substitution of Labor for AI," ArXiv:2602.00139 [econ]. 4

Tomlinson, K., S. Jaffe, W. Wang, S. Counts, and S. Suri (2025): "Working with AI: Measuring the Applicability of Generative AI to Occupations," ArXiv:2507.07935 [cs]. 1.2, 3, E.5

Tucker, L. C. (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. (document), 2.3, 2.4, 17, 3, 5

U.S. Census Bureau (2025a): "Guidance for Comparing the 2024 ACS 1-Year Estimates," American community survey research note, U.S. Census Bureau. 5, 43

——— (2025b): "Population Controls and Other Methodology for the 2024 ACS 1-Year Estimates," American Community Survey User Note 2025-03, U.S. Census Bureau. 5, H

——— (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. 3

Wiles, E. and J. J. Horton (2025): "Generative AI and Labor Market Matching Efficiency," . 5

Yotzov, I., J. M. Barrero, N. Bloom, P. Bunn, S. J. Davis, K. M. Foster, A. Jalca, B. H. Meyer, P. Mizen, M. A. Navarrete, P. Smietanka, G. Thwaites, and B. Z.

Wang (2026): "Firm Data on AI," Tech. Rep. w34836, National Bureau of Economic Research. (document)

Zens, G., M. Böck, and T. O. Zörner (2020): "The heterogeneous impact of monetary policy on the US labor market," *Journal of Economic Dynamics and Control*, 119, 103989. 2.2, 1, 2, 3, 24, 25, A.4, C.1, C.2, C.3

Online Appendix

A Data Appendix

This study uses a sample of data from ADP, a global leader in HR and payroll solutions. The company provides payroll services for firms that collectively employ over 26 million workers in the US. We use this to track employment changes for workers in occupations measured as more or less exposed to artificial intelligence.

We make several restrictions for our main analysis sample. We include only workers with positive earnings, exclude part-time employees, and subset to people under age 70.[36]

The set of firms using payroll services changes over time as companies join or leave ADP's platform. We maintain a consistent set of firms by keeping only companies that have employee earnings records for each month from January 2021 through June 2026. We also construct a balanced panel of firms observed from January 2018 through June 2026, which covers a longer pre-ChatGPT window and which we use where assessing pre-trends matters. Because it conditions on firm survival over a longer horizon, that panel is smaller.

ADP observes job titles for about 70% of workers in its system. Where a job title is missing we impute the occupation code from the worker's own employment history where possible, and exclude the records that remain unmatched. There are over 7,000 standardized job titles, including "Search engine optimization specialist," "Enterprise content management manager," and "Plant documentation control specialist." The company's internal research team maps each of these job titles to a 2010 Standard Occupational Classification (SOC) code, additionally using information such as the job description, industry, location, and other relevant data. We use these estimated SOC codes to merge our data to occupational AI exposure measures. Online Appendix Section A.1 examines the missing occupation codes and describes how we impute them where possible.

After these restrictions we have records on 3.5 to 5 million workers each month for our main analysis sample, though we consider robustness to alternative analyses such as allowing for firms to

[36]While we observe the year of birth for each worker, for privacy reasons we do not observe the exact date of birth. We impute month of birth from the distribution of birth months in the United States using data from the Centers for Disease Control and Prevention.

enter and leave the sample.

While the ADP data include millions of workers in each month, the distribution of firms using ADP services does not exactly match the distribution of firms across the broader US economy. Furthermore, our analysis sample is a subset of the universe of ADP data. Section A.2 compares the ADP analysis sample directly to official statistics. [37]

A.1 Missing Occupation Codes and Imputation

Because job titles are missing for roughly 30% of the sample, a natural concern is that missingness correlates with the outcomes we study. We cannot prove that occupation codes are missing at random, but we can correlate missingness with observables.

We consider two ways of imputing missing occupation codes. The first fills in a worker's missing O*NET code with their most recent non-missing value. The second additionally back-fills remaining missing values with the worker's first future non-missing code. Figures A.1 and A.2 show the share of workers with missing O*NET codes over time under each imputation rule, by age group and by BLS supersector. Trends in missingness do not correlate notably with age; differences are more notable across sectors.

The imputation methods reduce missingness and diminish the September–October 2025 spike in missing values. Descriptive results under both imputation methods closely mirror the results without imputation. We use the forward- and back-filled imputation as the primary specification in the paper.

Figures A.3 and A.4 reproduce our main descriptive result — employment changes by age and exposure quintile — under the two alternative treatments of missing occupation codes: no imputation at all, and forward-fill only. Both reproduce the pattern in the main text: within the youngest age group, employment declines in the most exposed quintiles and grows in the least exposed, while older workers show no comparable gradient. If anything the imputed series show slightly larger declines. For workers aged 22–25 in June 2026, employment in the most exposed quintile stands at 0.92 of its November 2022 level without imputation and 0.93 under forward-fill

[37] See also Cajner et al. (2018) and ADP Research (2025). Cajner et al. (2018) find a somewhat higher share of manufacturing and services establishments compared to the QCEW using data from March 2016. They also find that ADP somewhat overrepresents establishments in the Northeast. In addition, employers using ADP tend to grow faster on average than the typical establishment in the US economy.

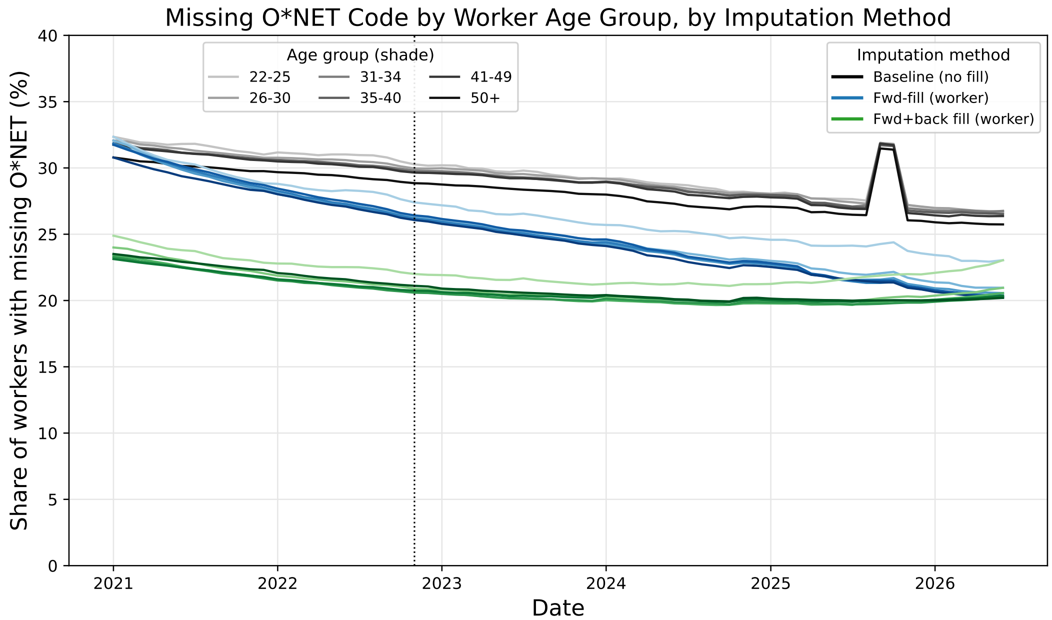

Figure A.1: Share of workers with missing O*NET occupation codes over time, by age group and imputation method. Line color denotes the imputation rule (no fill; forward-fill within worker; forward- and back-fill within worker) and shade denotes the age group, from lightest (22–25) to darkest (50+). The vertical line marks November 2022.

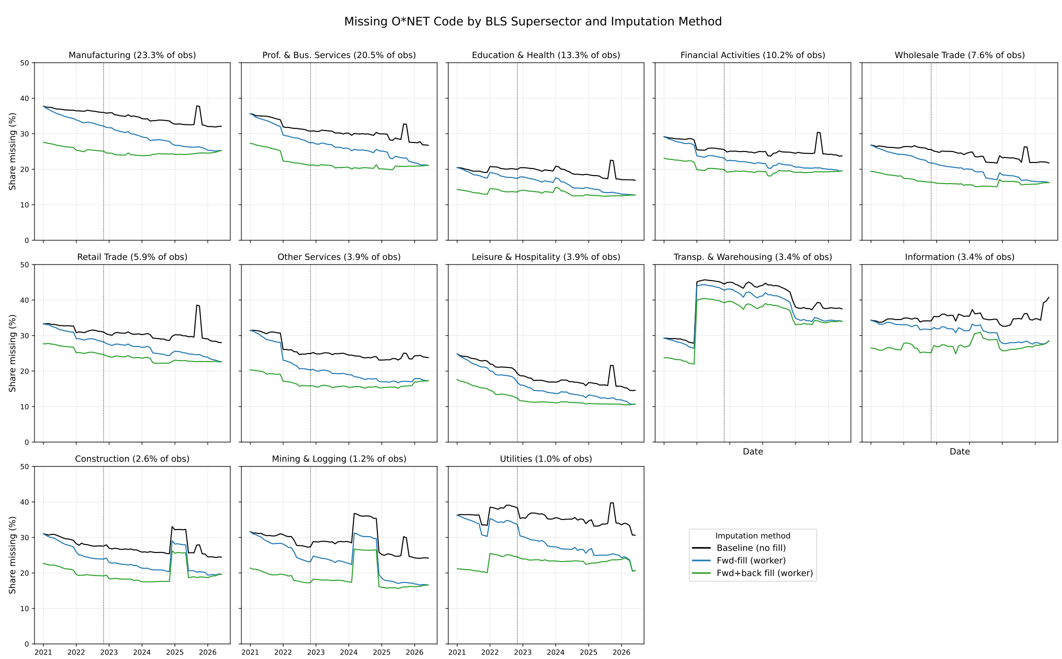

Figure A.2: Share of workers with missing O*NET occupation codes over time, by BLS supersector and imputation method. Panels are ordered by share of observations and each shows the three imputation rules (no fill; forward-fill within worker; forward- and back-fill within worker). Public administration and unclassified establishments are omitted. The vertical line marks November 2022.

alone, compared with 0.88 in our primary specification, while the least exposed quintiles grow in all three. The choice of imputation rule therefore does not drive our findings.

A.2 Representativeness

Figures A.5 and A.6 compare the composition of the ADP analysis sample to the national firm distribution from the Census Bureau's Business Dynamics Statistics (BDS). The Business Dynamics Statistics is an annual administrative dataset covering paid employment, constructed by linking annual U.S. Census Bureau business registers longitudinally.[38] The ADP analysis sample is broadly representative of the industry distribution, though it skews towards a higher share of manufacturing and wholesale firms and a relatively lower share of the retail and accommodation/food sectors. The ADP analysis sample also has a higher share of larger firms relative to the BDS; very few firms in the analysis sample have fewer than 10 employees.[39]

Table A.1 compares total employment growth in our analysis samples to official BLS payroll employment series. We measure growth from November 2022 through June 2026, the final month of our analysis sample and the latest month available in the BLS series at the time of writing. The balanced panel grows faster than national payroll employment over this period. This may be because conditioning on firm survival in the balanced panel pushes up growth rates. The unbalanced sample, which allows firms to enter and exit, grows more slowly than official statistics—indeed it declines outright over this window—and is noisier month-to-month. Figure A.7 plots the series over time.

Firm-time fixed effects sample Figures A.8 and A.9 show the firm-size and sector distributions of the fixed-effects sample with mimimal filters, separately by age group, alongside the national distribution from the BDS.

The ADP sample also contains a higher share of occupations with high AI exposure relative to the CPS and the ACS, for both men and women. Fig. A.10 compares each gender's employment across AI exposure quintiles in the three data sets.

[38] We benchmark against the BDS rather than the Quarterly Census of Employment and Wages (QCEW) because the BDS is constructed at the firm level, matching the unit of observation in our balanced panel, whereas the QCEW reports employment at the establishment level. Firm definitions need not align exactly across sources. ADP's client relationships are organized around payroll accounts, which may aggregate or split legal entities differently from the Census Bureau's firm definition.

[39] Comparisons with the BDS exclude NAICS code 92 (public administration) because the BDS only reports private-sector nonfarm employment. Public administration makes up 1.5% of employment in the ADP balanced

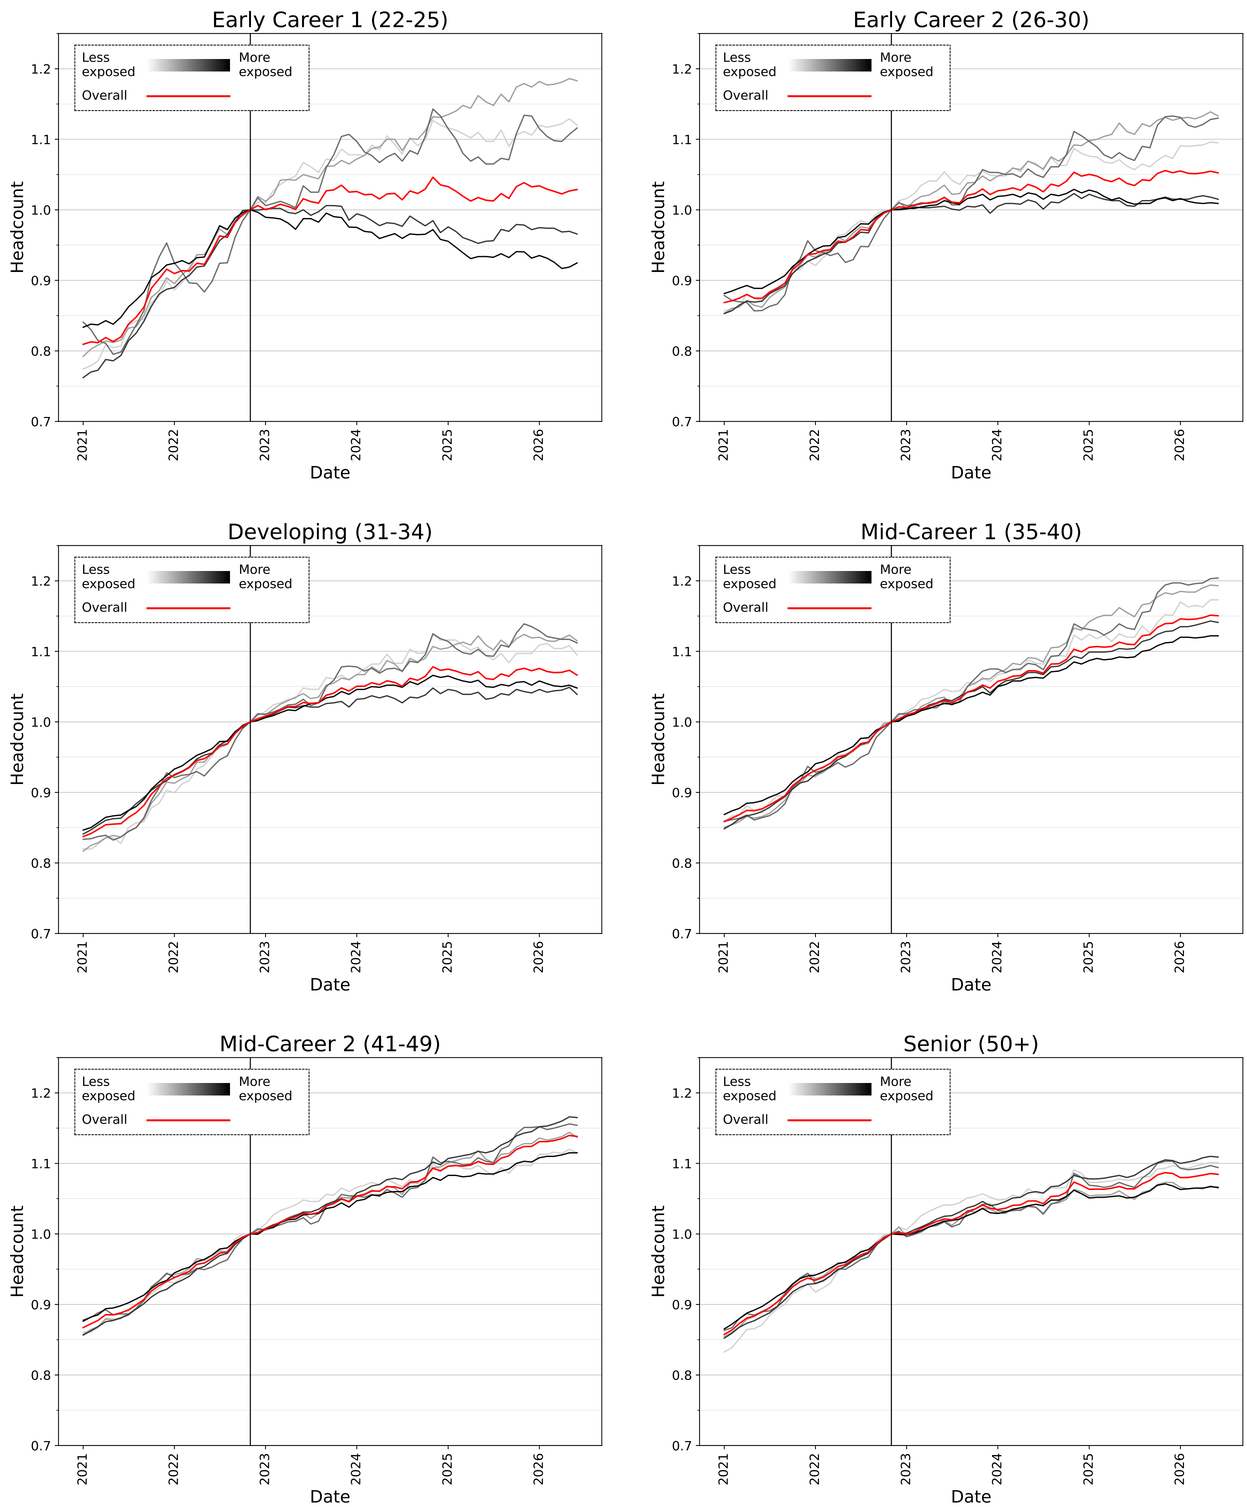

Figure A.3: *No imputation.* Employment changes by age and exposure quintile using measures from Eloundou et al. (2024), dropping workers with missing occupation codes rather than imputing them. Darker lines are more exposed quintiles; the red line pools across quintiles.

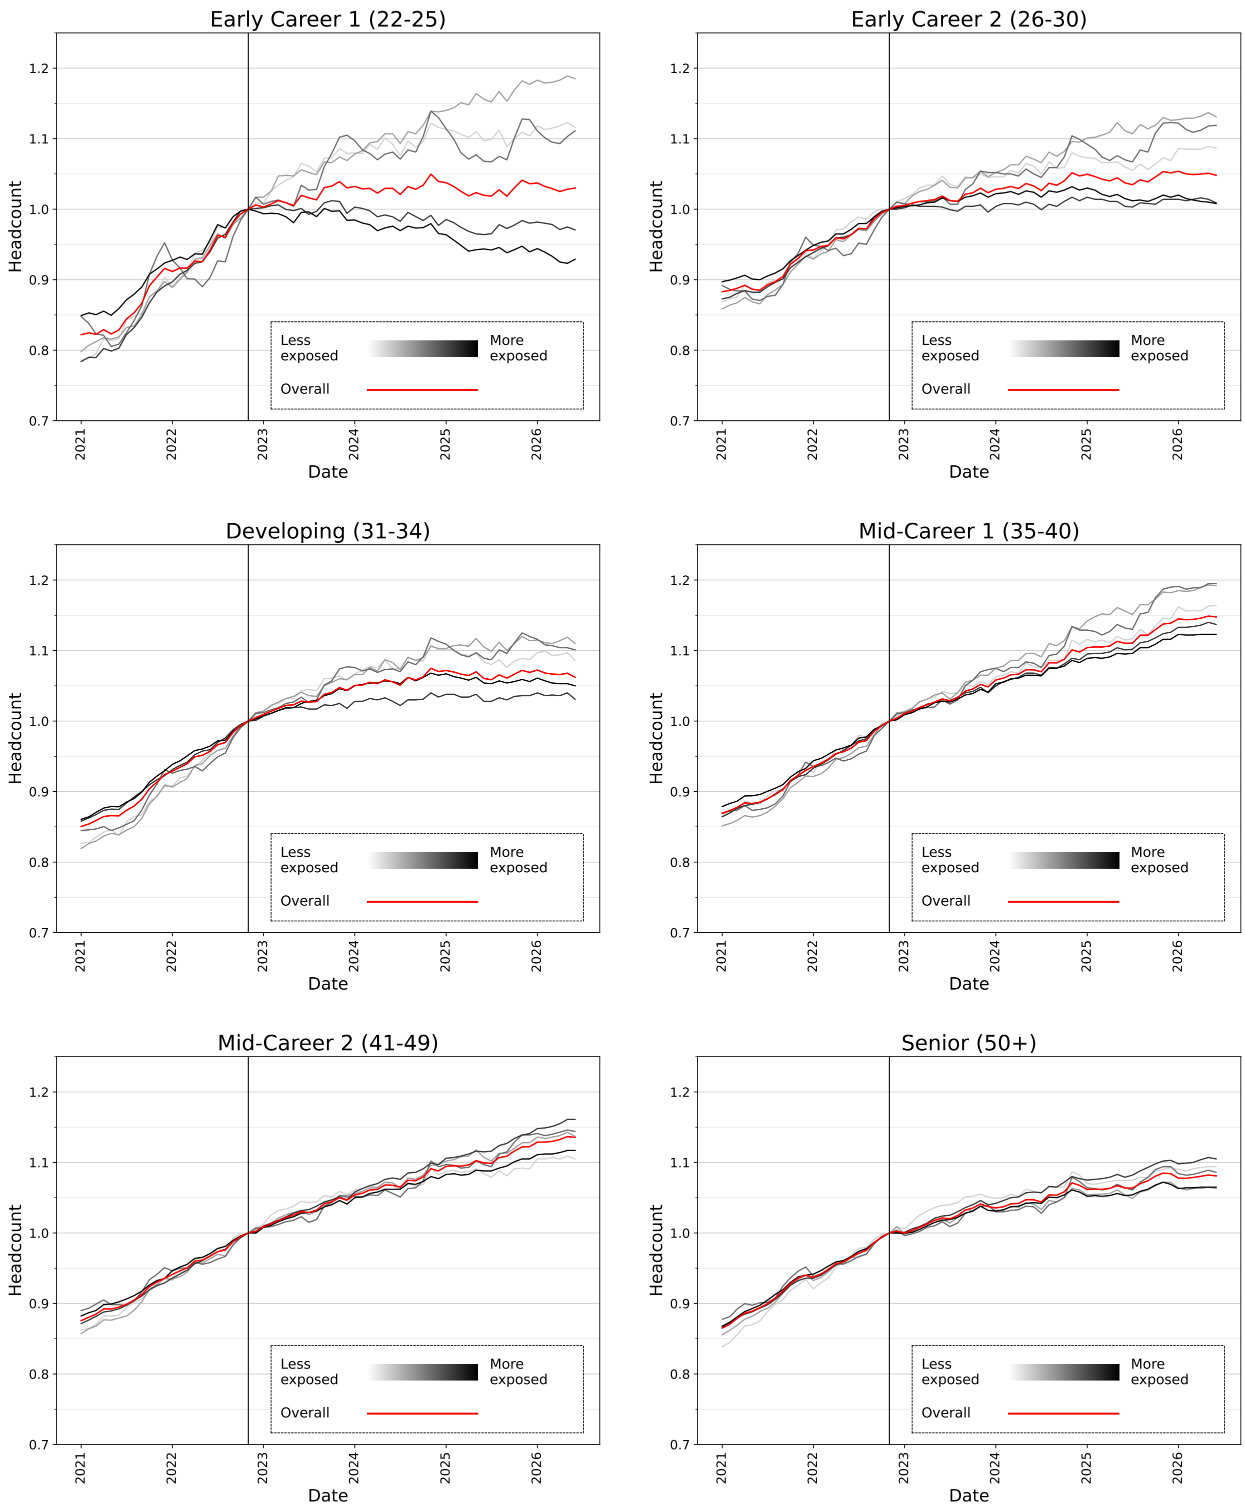

Figure A.4: *Forward-fill imputation only.* Employment changes by age and exposure quintile using measures from Eloundou et al. (2024), filling a worker's missing occupation code with their most recent non-missing value but not back-filling. Darker lines are more exposed quintiles; the red line pools across quintiles.

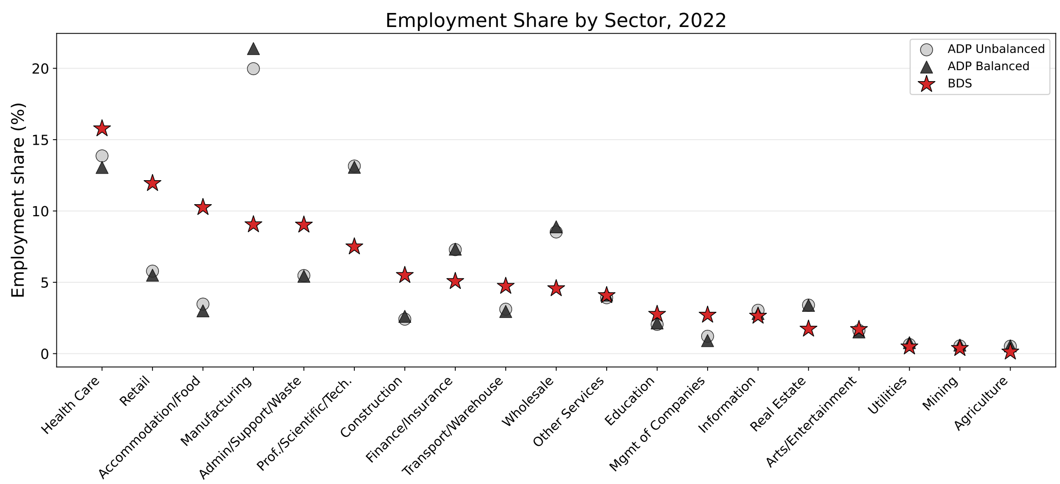

Figure A.5: Employment shares by sector in the ADP analysis sample compared to the Business Dynamics Statistics.

Figure A.6: Employment shares by firm size in the ADP analysis sample compared to the Business Dynamics Statistics.

Table A.1: Total employment growth, November 2022 to June 2026. ADP series are computed from the analysis samples in this paper; official series are BLS payroll employment from FRED (PAYEMS: total nonfarm; USPRIV: total private).

Series	Growth
ADP balanced panel (main analysis sample)	+6.3%
ADP unbalanced sample	−3.7%
BLS total nonfarm payrolls (PAYEMS)	+3.1%
BLS total private employment (USPRIV)	+2.8%

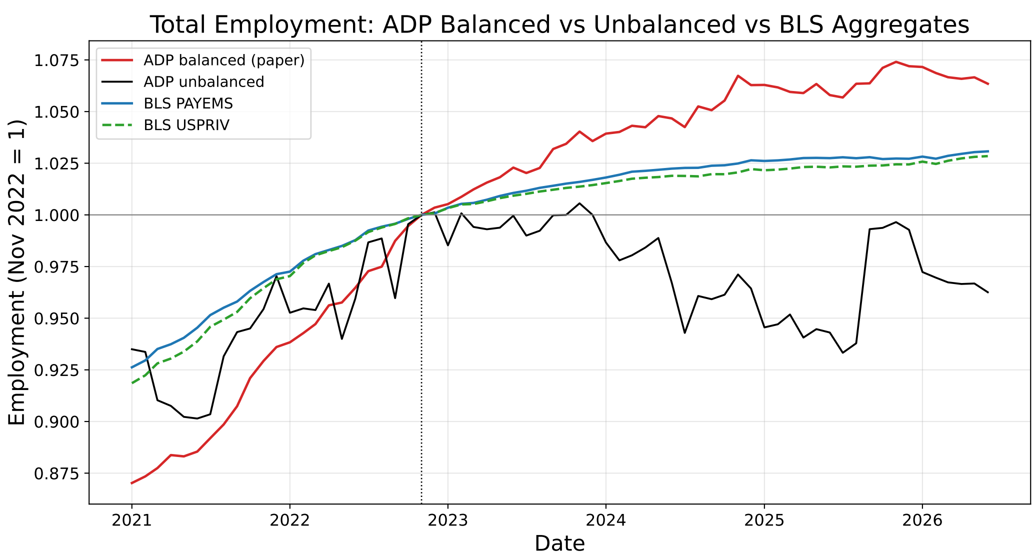

Figure A.7: Total employment over time in the ADP balanced and unbalanced analysis samples compared to BLS payroll employment (FRED series PAYEMS and USPRIV), normalized to 1 in November 2022.

Figure A.8: Firm-count share by size class in the minimal filters firm fixed-effects estimation sample in 2022, shown separately for each age group, with the national firm distribution from the Business Dynamics Statistics for comparison.

analysis sample.

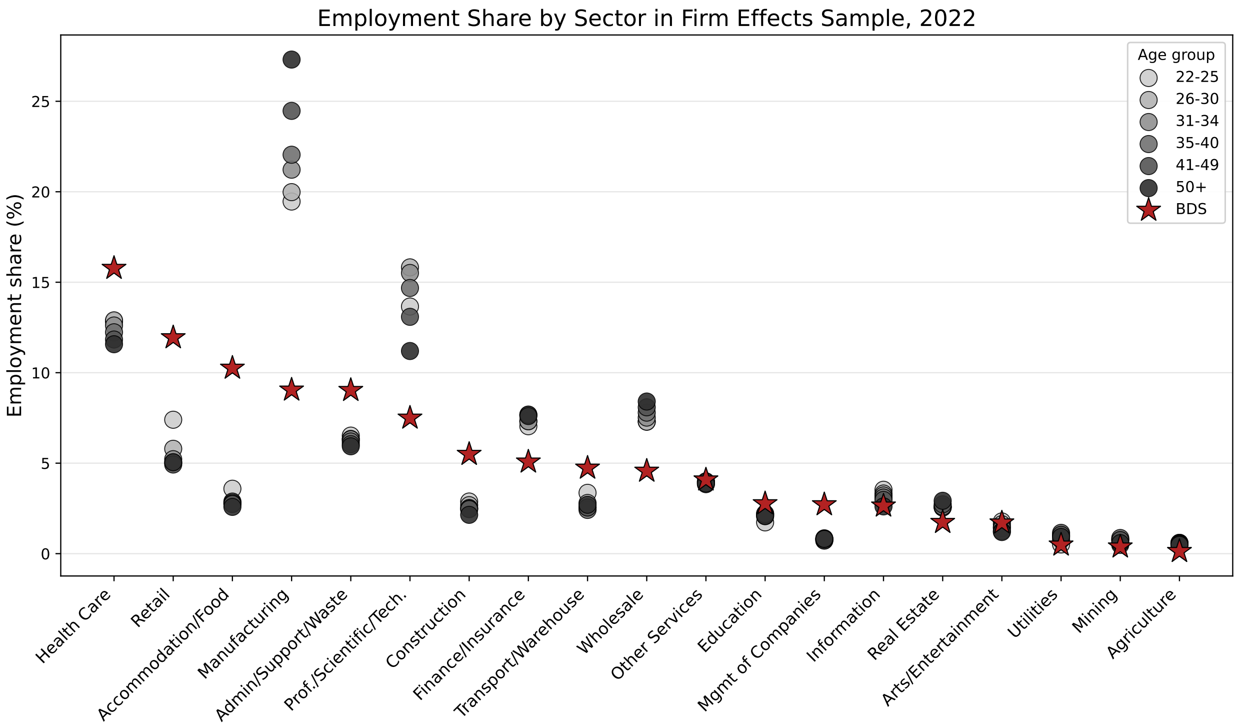

Figure A.9: Employment share by sector in the minimal filters firm fixed-effects estimation sample in 2022, shown separately for each age group, with the national employment distribution from the Business Dynamics Statistics for comparison. Sectors are ordered by descending BDS employment share.

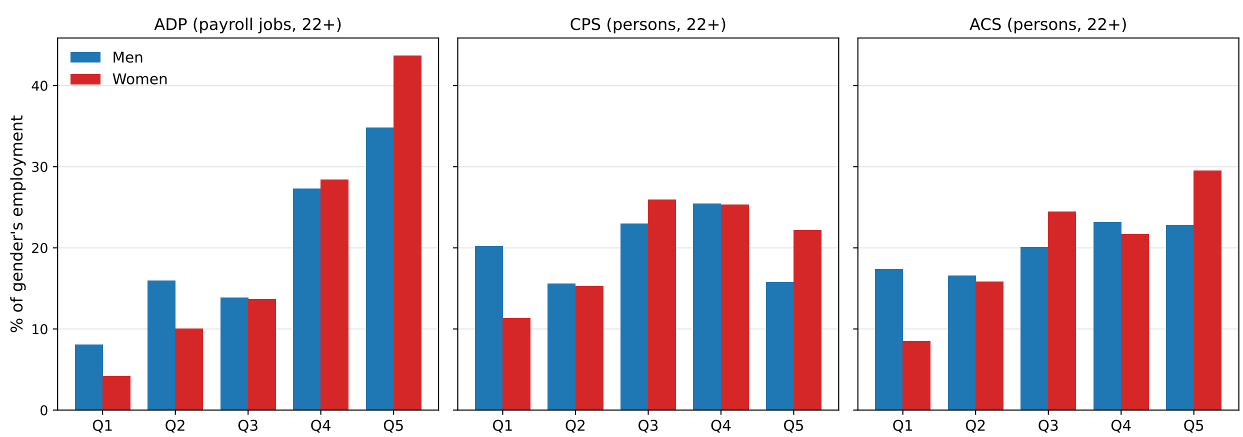

Figure A.10: Share of each gender's employment by AI exposure quintile in 2022, in the ADP analysis sample and in the CPS and ACS. Exposure quintiles are defined based on the GPT-4 β measures from Eloundou et al. (2024).

A.3 Occupational AI Exposure

We use two primary approaches for measuring occupational exposure to AI. The first uses exposure measures from Eloundou et al. (2024), who estimate AI exposure by O*NET task using ChatGPT validated with human labeling. They construct occupational exposure measures by aggregating the task data to the 2018 SOC code level. We focus on the GPT-4 based β exposure measures from their paper. Tables A.2 through A.6 show the top occupations by Eloundou et al. (2024) AI exposure.

The second approach we take uses generative AI usage data from the Anthropic Economic Index (Handa et al., 2025; Appel et al., 2025, 2026; Massenkoff et al., 2026). This index reports the estimated share of queries pertaining to each O*NET task based on several million conversations with Claude, Anthropic's generative AI model. It aggregates the data to the occupational level based on these task shares. One feature of the Anthropic Economic Index is that for each task it reports estimates of the share of queries pertaining to that task that are "automative," "augmentative," or none of the above. We use this as an estimate of whether AI usage for an occupation is primarily complementary or substitutable with labor.[40]

We use a 2010 SOC code to 2018 SOC code crosswalk from the BLS to merge the exposure measures to the payroll data. Table A.8 shows example occupations in the highest and lowest categories of each Anthropic Economic Index usage measure for the March 2025 release, and Table A.9 shows the same for the pooled September 2025 through April 2026 releases.

[40] Handa et al. (2025) use Claude to classify conversations into six categories. Directive and Feedback Loop conversations are considered Automative; Task Iteration, Learning, and Validation are considered Augmentative. They also instruct the model to choose None of the above "liberally." Table A.7 reproduces Table 1 from Handa et al. (2025) and shows more details about the automation and augmentation measures.

Table A.2: Top 50 occupations by ADP employment, Quintile 1 (lowest exposure) (October 2022).

Rank	Occupation
1	Laborers and Freight, Stock, and Material Movers, Hand
2	Maids and Housekeeping Cleaners
3	Industrial Truck and Tractor Operators
4	Packaging and Filling Machine Operators and Tenders
5	Grounds Maintenance Workers (incl. Landscaping and Groundskeeping Workers; Pesticide Handlers, Sprayers, and Applicators, Vegetation; Tree Trimmers and Pruners)
6	Molding, Coremaking, and Casting Machine Setters, Operators, and Tenders, Metal and Plastic
7	Janitors and Cleaners, Except Maids and Housekeeping Cleaners
8	Home Health Aides
9	Cutting, Punching, and Press Machine Setters, Operators, and Tenders, Metal and Plastic
10	Automotive Master Mechanics
11	Construction Laborers
12	Cleaners of Vehicles and Equipment
13	Automotive Body and Related Repairers
14	Dishwashers
15	Refuse and Recyclable Material Collectors (incl. Tank Car, Truck, and Ship Loaders)
16	Electromechanical Equipment Assemblers
17	Laundry and Dry-Cleaning Workers
18	Shoe and Leather Workers and Repairers
19	Paving, Surfacing, and Tamping Equipment Operators
20	Farmworkers, Farm, Ranch, and Aquacultural Animals
21	Plumbers
22	Food Preparation Workers
23	Extruding and Drawing Machine Setters, Operators, and Tenders, Metal and Plastic
24	Glaziers
25	Fitness Trainers and Aerobics Instructors

Rank	Occupation
26	Paper Goods Machine Setters, Operators, and Tenders
27	Sewing Machine Operators
28	Grinding and Polishing Workers, Hand
29	Surgical Technologists
30	Excavating and Loading Machine and Dragline Operators
31	Personal Care and Service Workers, All Other
32	Grinding, Lapping, Polishing, and Buffing Machine Tool Setters, Operators, and Tenders, Metal and Plastic
33	Sawing Machine Setters, Operators, and Tenders, Wood
34	Material Moving Workers, All Other
35	Sheet Metal Workers
36	Coating, Painting, and Spraying Machine Setters, Operators, and Tenders
37	Rotary Drill Operators, Oil and Gas
38	Structural Metal Fabricators and Fitters
39	Electrical Power-Line Installers and Repairers
40	Painters, Construction and Maintenance
41	Hairdressers, Hairstylists, and Cosmetologists
42	Slaughterers and Meat Packers
43	Helpers–Extraction Workers
44	Textile Winding, Twisting, and Drawing Out Machine Setters, Operators, and Tenders
45	Cutting and Slicing Machine Setters, Operators, and Tenders
46	Riggers
47	Dental Laboratory Technicians
48	Woodworking Machine Setters, Operators, and Tenders, Except Sawing
49	Adhesive Bonding Machine Operators and Tenders
50	Pipe Fitters and Steamfitters

Table A.3: Top 50 occupations by ADP employment, Quintile 2 (October 2022).

Rank	Occupation
1	Maintenance and Repair Workers, General
2	Nursing Assistants
3	Helpers–Production Workers
4	Stock Clerks- Stockroom, Warehouse, or Storage Yard
5	Welders, Cutters, and Welder Fitters
6	Cooks, Restaurant
7	Security Guards
8	Waiters and Waitresses
9	Machinists
10	Baggage Porters and Bellhops
11	Stock Clerks and Order Fillers
12	Cashiers
13	Electricians
14	Medical and Clinical Laboratory Technicians
15	Order Fillers, Wholesale and Retail Sales
16	Preschool Teachers, Except Special Education
17	Physical Therapists
18	Emergency Medical Technicians and Paramedics
19	Construction Carpenters
20	Counter Attendants, Cafeteria, Food Concession, and Coffee Shop
21	Firefighters
22	Farmworkers and Laborers, Crop, Nursery, and Greenhouse
23	Heating, Air Conditioning, and Refrigeration Mechanics and Installers
24	Bartenders (incl. Baristas; Fast Food and Counter Workers)
25	Demonstrators and Product Promoters
26	Physical Therapist Assistants
27	Phlebotomists
28	Plant and System Operators, All Other
29	Dental Hygienists
30	Veterinary Technologists and Technicians
31	Water and Wastewater Treatment Plant and System Operators
32	Electrical and Electronic Equipment Assemblers
33	Flight Attendants
34	Recreation Workers
35	Tool and Die Makers
36	Petroleum Pump System Operators, Refinery Operators, and Gaugers
37	Psychiatric Technicians
38	Parking Lot Attendants
39	Crane and Tower Operators
40	Automotive and Watercraft Service Attendants
41	Mail Clerks and Mail Machine Operators, Except Postal Service
42	Personal Care Aides
43	Print Binding and Finishing Workers
44	Childcare Workers (incl. Nannies)
45	Aircraft Mechanics and Service Technicians
46	Plating and Coating Machine Setters, Operators, and Tenders, Metal and Plastic
47	Food Batchmakers
48	Mixing and Blending Machine Setters, Operators, and Tenders
49	Food Servers, Nonrestaurant
50	Gaming Dealers

Table A.4: Top 50 occupations by ADP employment, Quintile 3 (October 2022).

Rank	Occupation
1	Taxi Drivers and Chauffeurs
2	Registered Nurses (incl. Acute Care Nurses; Advanced Practice Psychiatric Nurses; Clinical Nurse Specialists; Critical Care Nurses)
3	Retail Salespersons
4	First-Line Supervisors of Housekeeping and Janitorial Workers
5	Computer Numerically Controlled Machine Tool Programmers, Metal and Plastic
6	Heavy and Tractor-Trailer Truck Drivers
7	Administrative Services Managers
8	Social and Community Service Managers
9	Medical Assistants
10	Manufacturing Production Technicians
11	Licensed Practical and Licensed Vocational Nurses
12	Chefs and Head Cooks
13	Industrial and Refractory Machinery Mechanics (incl. Industrial Machinery Mechanics; Maintenance Workers, Machinery; Millwrights; Refractory Materials Repairers, Except Brickmasons)
14	Driver/Sales Workers
15	Teacher Assistants
16	Police Patrol Officers
17	Engineering Technicians, Except Drafters, All Other
18	Mental Health and Substance Abuse Social Workers
19	Chemists
20	Secondary School Teachers, Except Special and Career/Technical Education
21	Production Workers, All Other
22	Child, Family, and School Social Workers
23	Bus Drivers, Transit and Intercity
24	Dental Assistants
25	Merchandise Displayers and Window Trimmers
26	Computer-Controlled Machine Tool Operators, Metal and Plastic
27	Physician Assistants (incl. Anesthesiologist Assistants)
28	Elementary School Teachers, Except Special Education
29	Educational, Guidance, School, and Vocational Counselors
30	Construction and Building Inspectors (incl. Energy Auditors)
31	First-Line Supervisors of Transportation and Material-Moving Machine and Vehicle Operators
32	Radiologic Technologists
33	Printing Press Operators
34	Social and Human Service Assistants
35	Athletes, Coaches, Umpires, and Related Workers (incl. Athletes and Sports Competitors; Coaches and Scouts; Umpires, Referees, and Other Sports Officials)
36	Occupational Therapists (incl. Low Vision Therapists, Orientation and Mobility Specialists, and Vision Rehabilitation Therapists)
37	Residential Advisors
38	Mental Health Counselors
39	Nonfarm Animal Caretakers
40	Telecommunications Equipment Installers and Repairers, Except Line Installers
41	Bailiffs, Correctional Officers, and Jailers (incl. Bailiffs; Correctional Officers and Jailers)
42	Middle School Teachers, Except Special and Career/Technical Education
43	Clergy
44	Diagnostic Medical Sonographers
45	Mechanical Engineering Technicians
46	Light Truck or Delivery Services Drivers
47	Ophthalmic Medical Technicians
48	Airline Pilots, Copilots, and Flight Engineers
49	Healthcare Social Workers
50	Bakers

Table A.5: Top 50 occupations by ADP employment, Quintile 4 (October 2022).

Rank	Occupation	Rank	Occupation
1	Sales Managers	26	Construction Managers
2	General and Operations Managers	27	Training and Development Managers
3	First-Line Supervisors of Production and Operating Workers	28	Family and General Practitioners
4	Marketing Managers	29	Food Service Managers
5	Medical and Health Services Managers	30	Logistics Managers
6	First-Line Supervisors of Office and Administrative Support Workers	31	Education Administrators, Postsecondary
7	Industrial Production Managers	32	First-Line Supervisors of Construction Trades and Extraction Workers (incl. Solar Energy Installation Managers)
8	First-Line Supervisors of Retail Sales Workers	33	Civil Engineers (incl. Transportation Engineers; Water/Wastewater Engineers)
9	Financial Managers	34	Sales Representatives, Wholesale and Manufacturing, Technical and Scientific Products
10	Human Resources Managers	35	Hotel, Motel, and Resort Desk Clerks
11	Management Analysts	36	Lodging Managers
12	Financial Managers, Branch or Department	37	Real Estate Sales Agents
13	First-Line Supervisors of Mechanics, Installers, and Repairers	38	Cost Estimators
14	Lawyers	39	Nurse Practitioners
15	Shipping, Receiving, and Traffic Clerks	40	Producers
16	Architectural and Engineering Managers (incl. Biofuels/Biodiesel Technology and Product Development Managers)	41	Physicians and Surgeons, All Other
17	Storage and Distribution Managers	42	Meeting, Convention, and Event Planners
18	Quality Control Systems Managers	43	Electronics Engineering Technicians
19	Financial Analysts	44	Transportation Managers
20	Property, Real Estate, and Community Association Managers	45	Art Directors
21	Inspectors, Testers, Sorters, Samplers, and Weighers	46	Commercial and Industrial Designers
22	Personal Financial Advisors	47	Special Education Teachers, All Other
23	Treasurers and Controllers	48	Pharmacy Technicians
24	Graphic Designers	49	Compensation, Benefits, and Job Analysis Specialists
25	Advertising and Promotions Managers	50	Education Administrators, Elementary and Secondary School

Table A.6: Top 50 occupations by ADP employment, Quintile 5 (highest exposure) (October 2022).

Rank	Occupation
1	Chief Executives (incl. Chief Sustainability Officers)
2	Customer Service Representatives
3	Managers, All Other
4	Accountants
5	Human Resources Specialists
6	Sales Representatives, Wholesale and Manufacturing, Except Technical and Scientific Products
7	Software Developers, Systems Software
8	Computer and Information Systems Managers
9	Receptionists and Information Clerks
10	Market Research Analysts and Marketing Specialists (incl. Search Marketing Strategists)
11	Computer User Support Specialists
12	Industrial Engineers
13	Miscellaneous Life, Physical, and Social Science Technicians (incl. Agricultural Technicians; Biological Technicians; Chemical Technicians; Environmental Science and Protection Technicians, Including Health; and others)
14	Loan Officers
15	Office Clerks, General
16	Computer Systems Analysts
17	Engineers, All Other
18	Production, Planning, and Expediting Clerks
19	Network and Computer Systems Administrators
20	Training and Development Specialists
21	Business Operations Specialists, All Other
22	Purchasing Managers
23	Bookkeeping, Accounting, and Auditing Clerks
24	Billing, Cost, and Rate Clerks
25	Patient Representatives

Rank	Occupation
26	Tellers
27	Computer Systems Engineers/Architects
28	Telemarketers
29	Executive Secretaries and Executive Administrative Assistants
30	Medical Records and Health Information Technicians
31	Paralegals and Legal Assistants
32	Computer Hardware Engineers
33	Supply Chain Managers
34	Computer Programmers
35	Computer Network Architects
36	Insurance Adjusters, Examiners, and Investigators
37	Bill and Account Collectors
38	Mechanical Engineers (incl. Automotive Engineers; Fuel Cell Engineers)
39	Compliance Officers (incl. Coroners; Customs Brokers; Environmental Compliance Inspectors; Equal Opportunity Representatives and Officers; and others)
40	Dispatchers (incl. Dispatchers, Except Police, Fire, and Ambulance; Public Safety Telecommunicators)
41	Legal Secretaries
42	Insurance Claims Clerks
43	Information Security Analysts
44	Information Technology Project Managers
45	Public Relations Specialists
46	Investment Fund Managers
47	Manufacturing Engineers
48	Logisticians (incl. Logistics Analysts; Logistics Engineers; Project Management Specialists)
49	Travel Agents
50	Medical Secretaries

Automative Behaviors *AI directly executes tasks with minimal human involvement*	**Augmentative Behaviors** AI enhances human capabilities through collaboration
Directive: Complete task delegation with minimal interaction *Illustrative Example: "Format this technical documentation in Markdown"*	**Task Iteration:** Collaborative refinement process *Illustrative Example: "Let's draft a marketing strategy for our new product. ... Good start, but can we add some concrete metrics?"*
Feedback Loop: Task completion guided by environmental feedback *Illustrative Example: "Here's my Python script for data analysis – it's giving an IndexError. Can you help fix it? ... Now I'm getting a different error..."*	**Learning:** Knowledge acquisition and understanding *Illustrative Example: "Can you explain how neural networks work?"*
	Validation: Work verification and improvement *Illustrative Example: "I've written this SQL query to find duplicate customer records. Can you check if my logic is correct and suggest any improvements?"*

Table A.7: Table 1 from Handa et al. (2025). Handa et al. (2025) classify conversations from Claude, the LLM, into five distinct patterns across two broad categories based on how people integrate AI into their workflow.

Table A.8: Example occupations by Anthropic Economic Index exposure category, March 2025 Vintage (Handa et al., 2025)

Metric	Least exposed (examples)	Most exposed (examples)
Per-Capita Usage	• Administrative Services Managers • Industrial Truck and Tractor Operators • Packaging and Filling Machine Operators and Tenders • Bus Drivers, Transit and Intercity	• Business Operations Specialists, All Other • Computer Systems Analysts • Network and Computer Systems Administrators • Advertising and Promotions Managers
Automation	• Maintenance and Repair Workers, General • Marketing Managers • First-Line Supervisors of Production and Operating Workers • Nursing Assistants	• Business Operations Specialists, All Other • Accountants and Auditors • First-Line Supervisors of Office and Administrative Support Workers • Receptionists and Information Clerks
Augmentation	• First-Line Supervisors of Production and Operating Workers • Laborers and Freight, Stock, and Material Movers, Hand • Welders, Cutters, Solderers, and Brazers • Tellers	• Chief Executives • Maintenance and Repair Workers, General • Registered Nurses • Sales Reps, Wholesale and Manufacturing

Table A.9: Example occupations by Anthropic Economic Index exposure category, Sept. 2025 - April 2026 Vintage (Appel et al., 2025, 2026; Massenkoff et al., 2026)

Metric	Least exposed (examples)	Most exposed (examples)
Per-Capita Usage	• Laborers and Freight, Stock, and Material Movers, Hand • Industrial Truck and Tractor Operators • Welders, Cutters, Solderers, and Brazers • Packaging and Filling Machine Operators and Tenders	• Business Operations Specialists, All Other • Computer Systems Analysts • Network and Computer Systems Administrators • Advertising and Promotions Managers
Automation	• Maintenance and Repair Workers, General • First-Line Supervisors of Production and Operating Workers • Home Health Aides • Mental Health and Substance Abuse Social Workers	• Business Operations Specialists, All Other • Customer Service Representatives • Sales Managers • Accountants and Auditors
Augmentation	• First-Line Supervisors of Production and Operating Workers • Administrative Services Managers • Food Service Managers • Insurance Underwriters	• Chief Executives • Marketing Managers • First-Line Supervisors of Office and Administrative Support Workers • Computer and Information Systems Managers

A.4 Other Data

To compare employment changes for teleworkable versus non-teleworkable occupations, we use data from Dingel and Neiman (2020) and Lambert and Schindler (2026). We use occupational interest rate estimates from Zens et al. (2020). We use the Personal Consumption Expenditure index from the BLS to compute real earnings, indexed to October 2017. We use monthly Current Population Survey (CPS) and American Community Survey (ACS) data as a comparison for our main findings.

B Supplementary Employment Results

This section provides additional descriptive employment results supporting the main text.

Figure B.1 shows employment by AI exposure quintile for workers aged 18–21, a younger band than the 22–25 group used throughout the main text, in the main sample of firms balanced since January 2021.

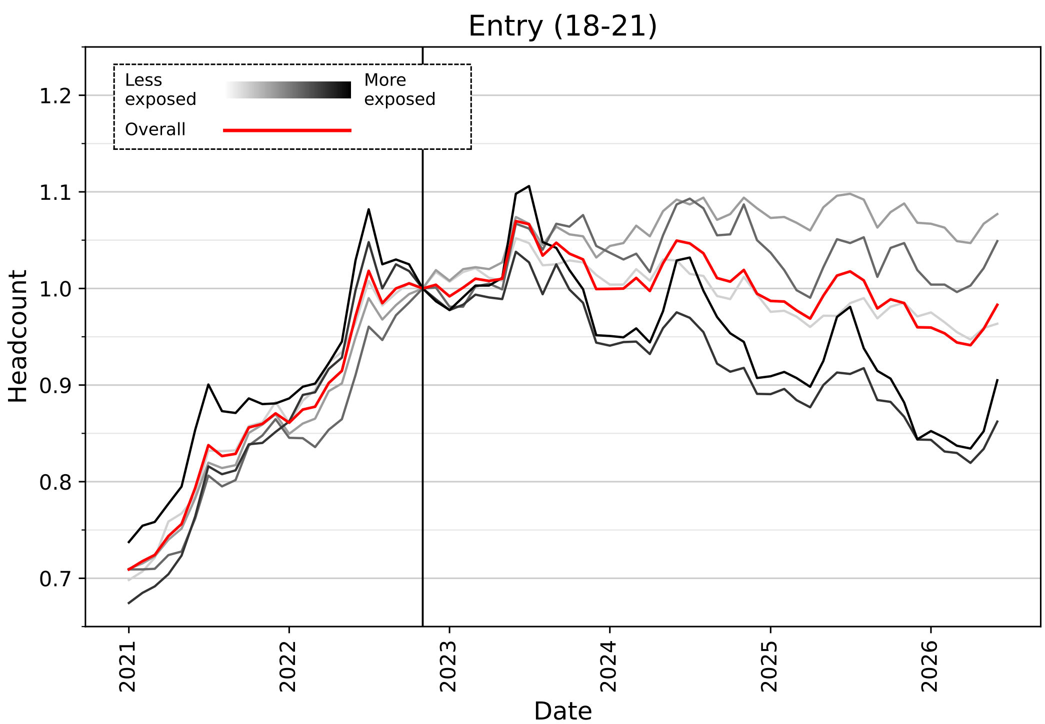

Figure B.1: Employment changes by AI exposure quintile for workers aged 18–21, normalized to 1 in November 2022, using measures from Eloundou et al. (2024). Darker lines are more exposed quintiles; the red line pools across quintiles. Firms balanced since January 2021.

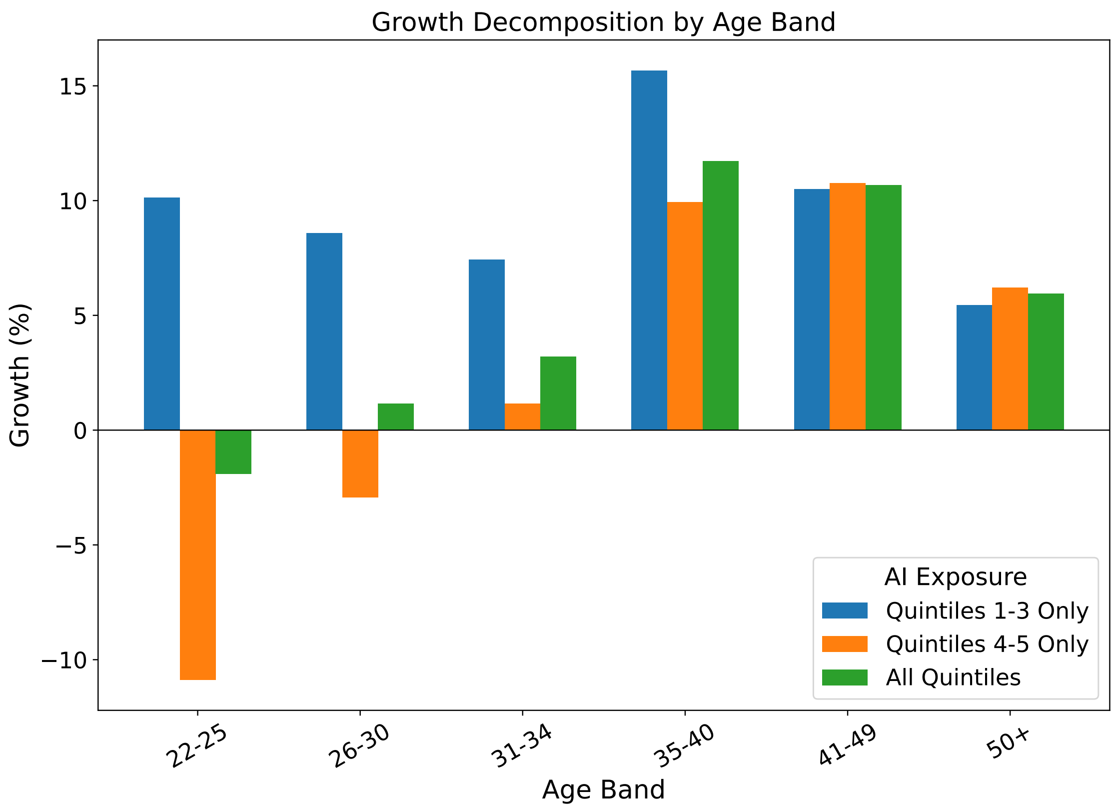

Figure B.2: Growth in employment between November 2022 and June 2026 by age and GPT-4 β-based AI exposure group.

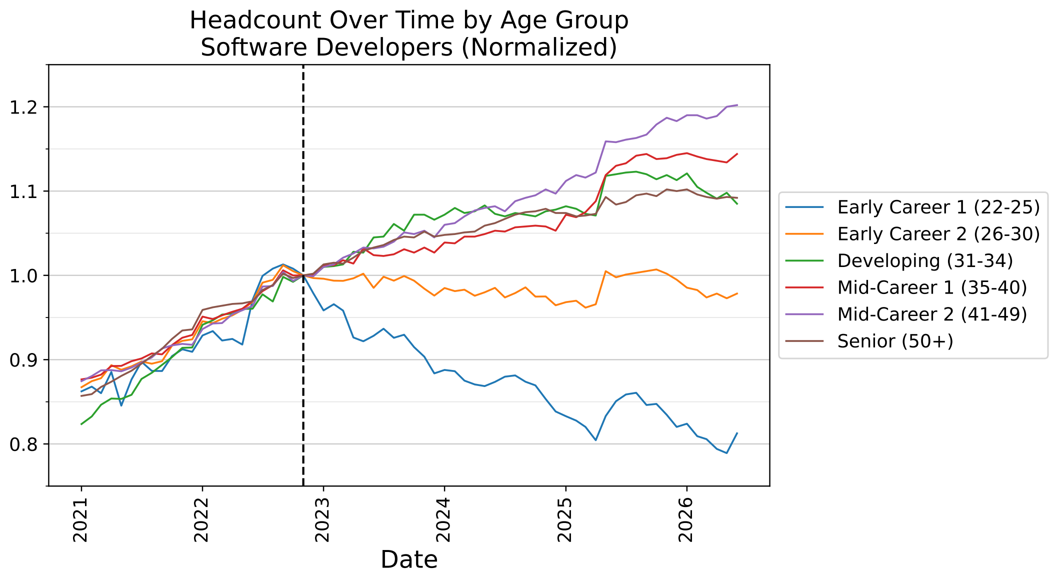

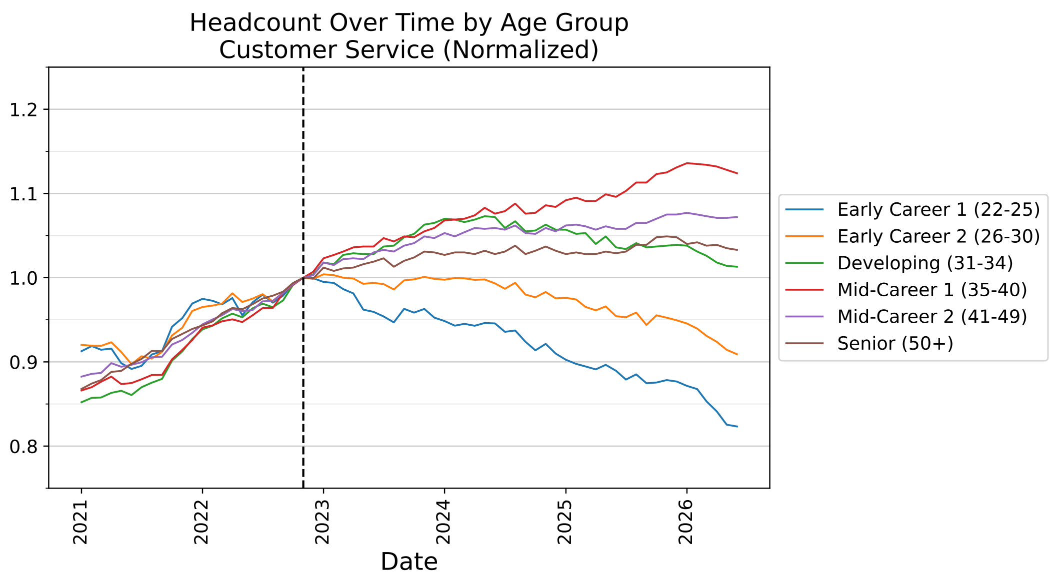

Figure B.3: Employment changes for software and customer service agents by age, normalized to 1 in November 2022.

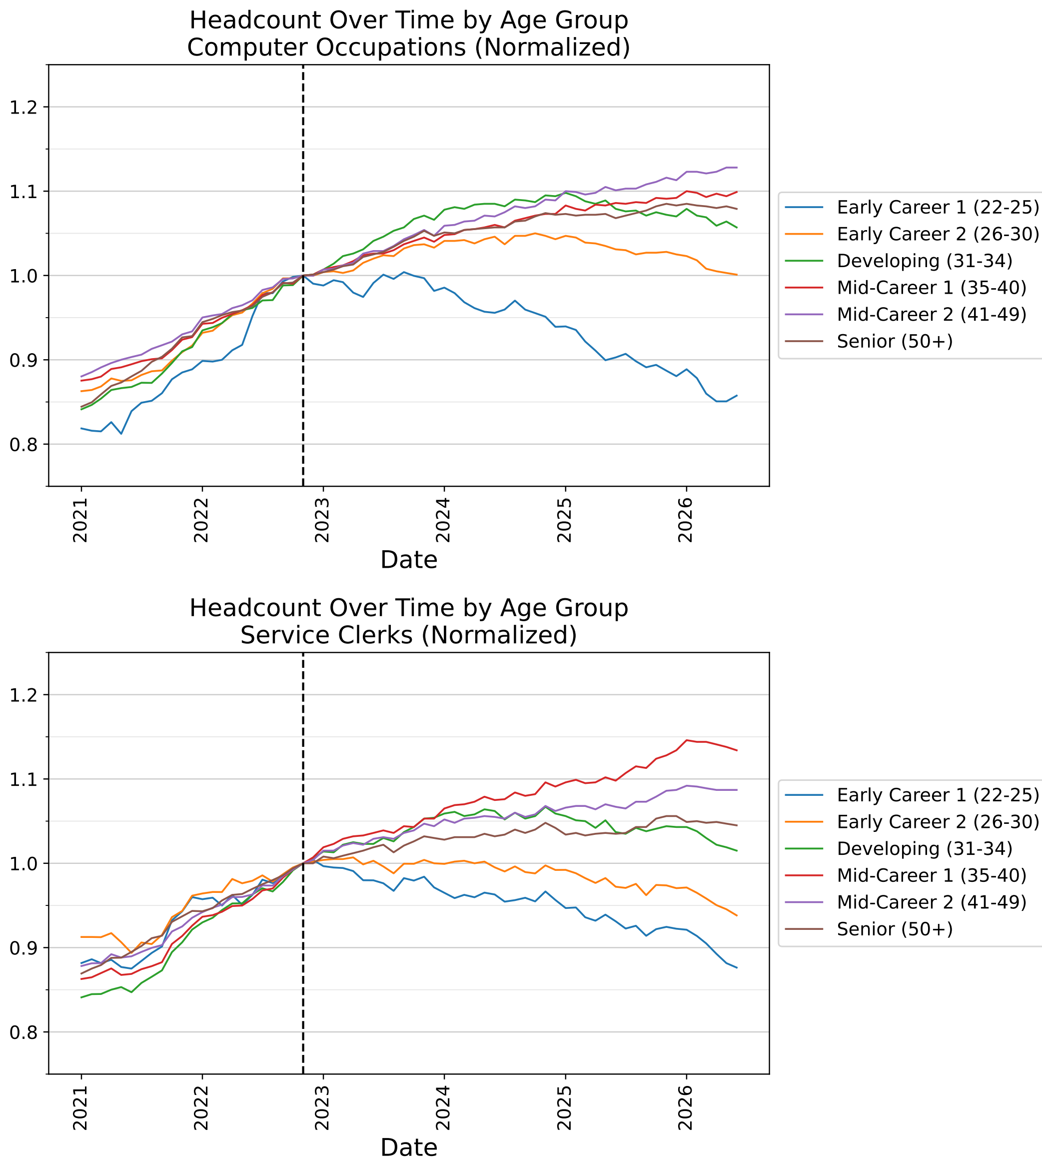

Figure B.4: Employment changes for computer occupations (2010 SOC codes starting with 15-1) and service clerks (starting with 43-4), normalized to 1 in November 2022.

Figure B.5: Employment changes for marketing and sales managers (exposure quintile 4), first-line production supervisors (exposure quintile 3), stock clerks and order fillers (exposure quintile 2), and health aides (exposure quintile 1), normalized to 1 in November 2022. Exposure quintiles are defined based on the Eloundou et al. (2024) GPT-4 β measure.

Figure B.6: Percent of occupations with an increase in employment for 22–25 year old workers between November 2022 and June 2026. Exposure quintiles are defined based on the Eloundou et al. (2024) GPT-4 β measure.

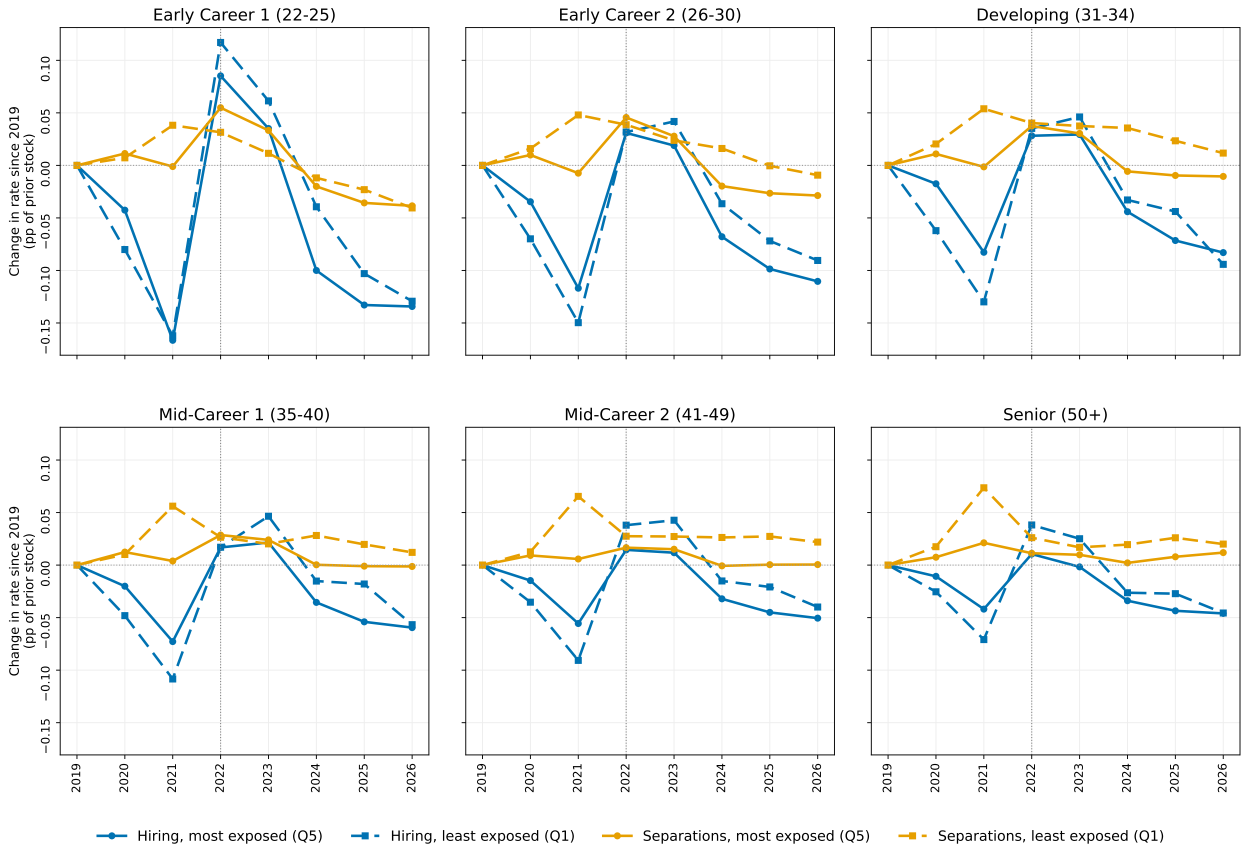

Figure B.7: Annual hiring and separation rates by age group and AI exposure, plotted as the change since 2019 in percentage points of the prior-year stock. The hiring rate is the number of new worker–firm matches in a cell over the prior year divided by the cell's headcount a year earlier; the separation rate is the analogous quantity for workers leaving the firm. Rates are measured January over January on the extended balanced panel of firms observed from January 2018. Blue lines are hiring and orange lines separations; solid lines with circles are the most exposed quintile (Q5) and dashed lines with squares the least exposed quintile (Q1). Exposure quintiles are defined based on the Eloundou et al. (2024) GPT-4 β measure.

C Robustness to Alternative Explanations

This section collects the robustness analyses referenced in the main text.

C.1 Occupational Interest Rate Exposure

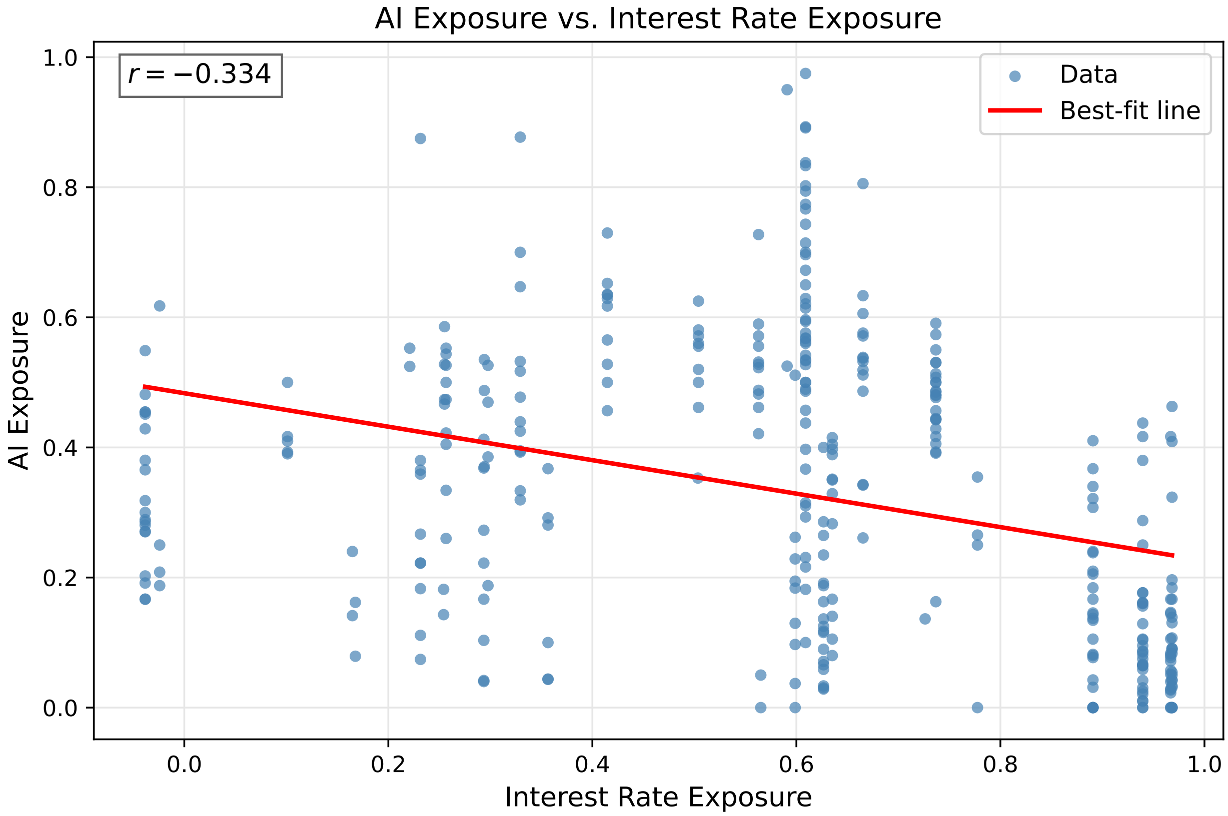

Figure C.1: Occupational AI exposure according to Eloundou et al. (2024) compared to interest rate exposure according to Zens et al. (2020). We measure interest rate exposure using the 2-year cumulative impulse response function identified via Cholesky decomposition. We use the crosswalk from Autor (2015) available on David Dorn's website to convert from 1990 occupation codes to 2010 Census codes. We then use a crosswalk to convert the 2010 Census codes to 2010 SOC codes.

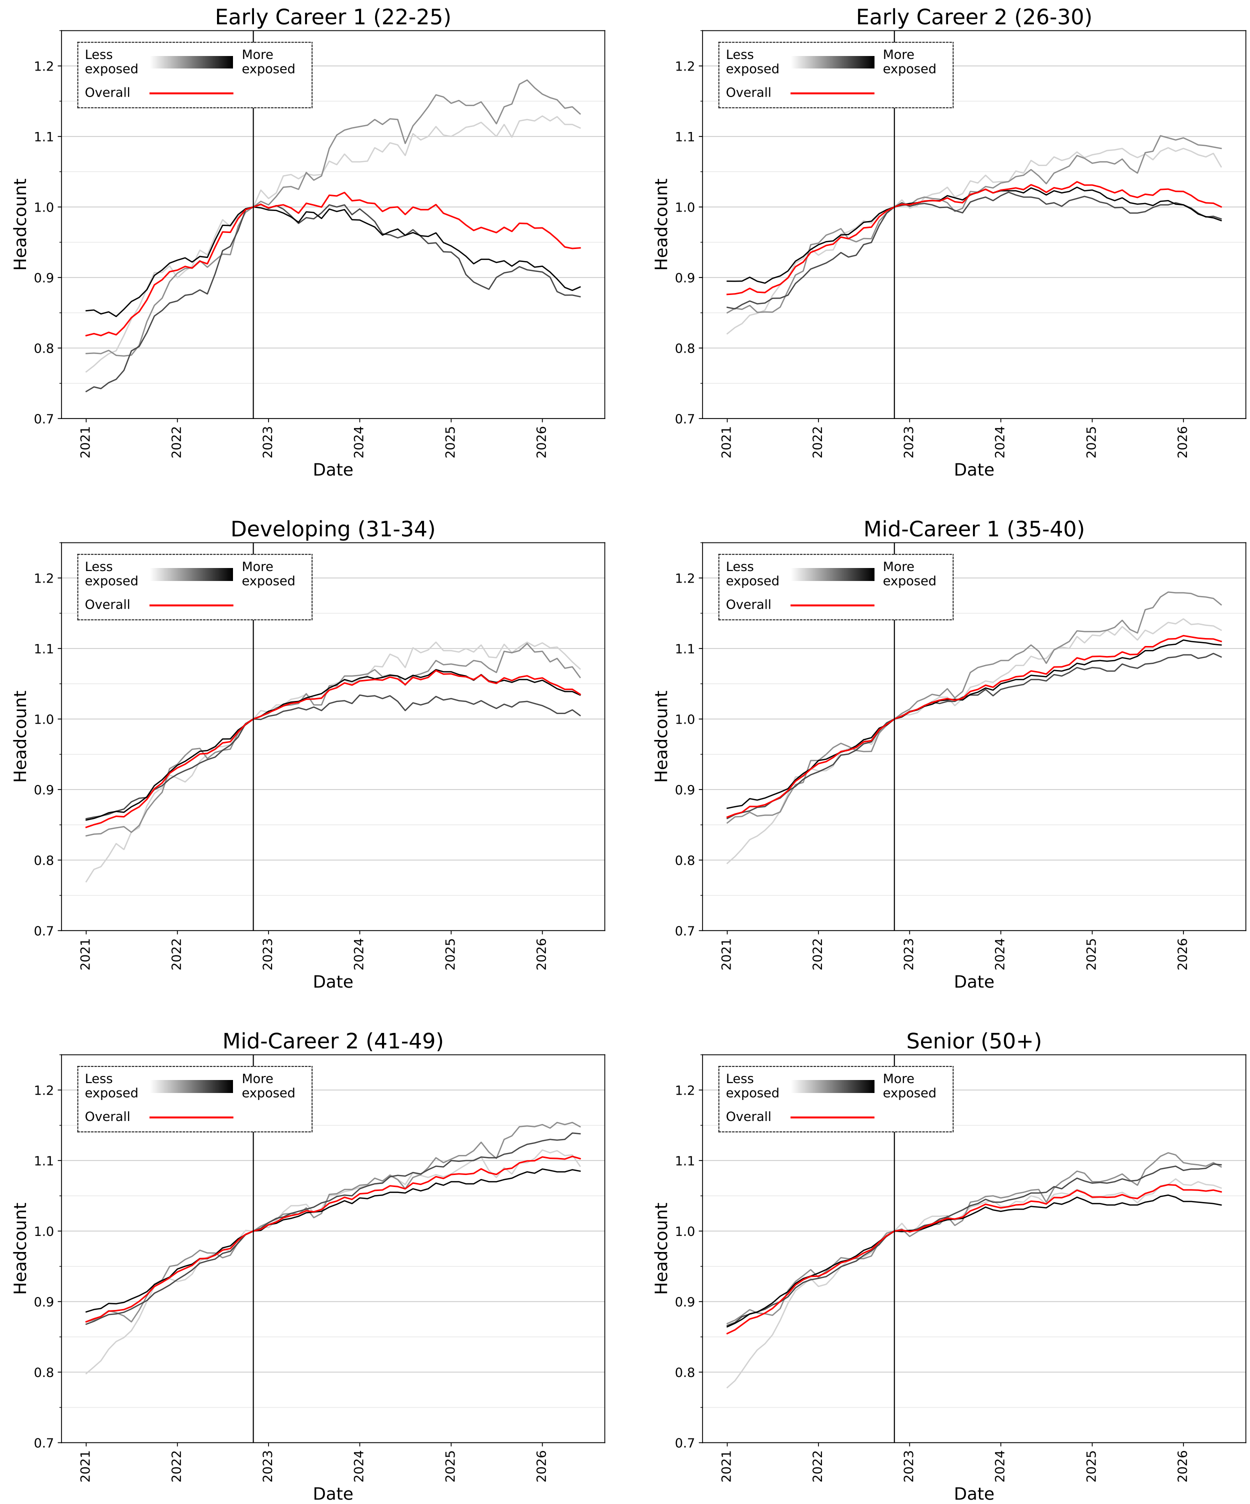

Figure C.2: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only occupations with below median interest rate exposure using data from Zens et al. (2020). Only a few occupations with low interest rate exposure also have low AI exposure. All occupations in the first and second quintile of AI exposure are consequently grouped together in level 1. The remaining quintiles are coded as 2, 3, and 4.

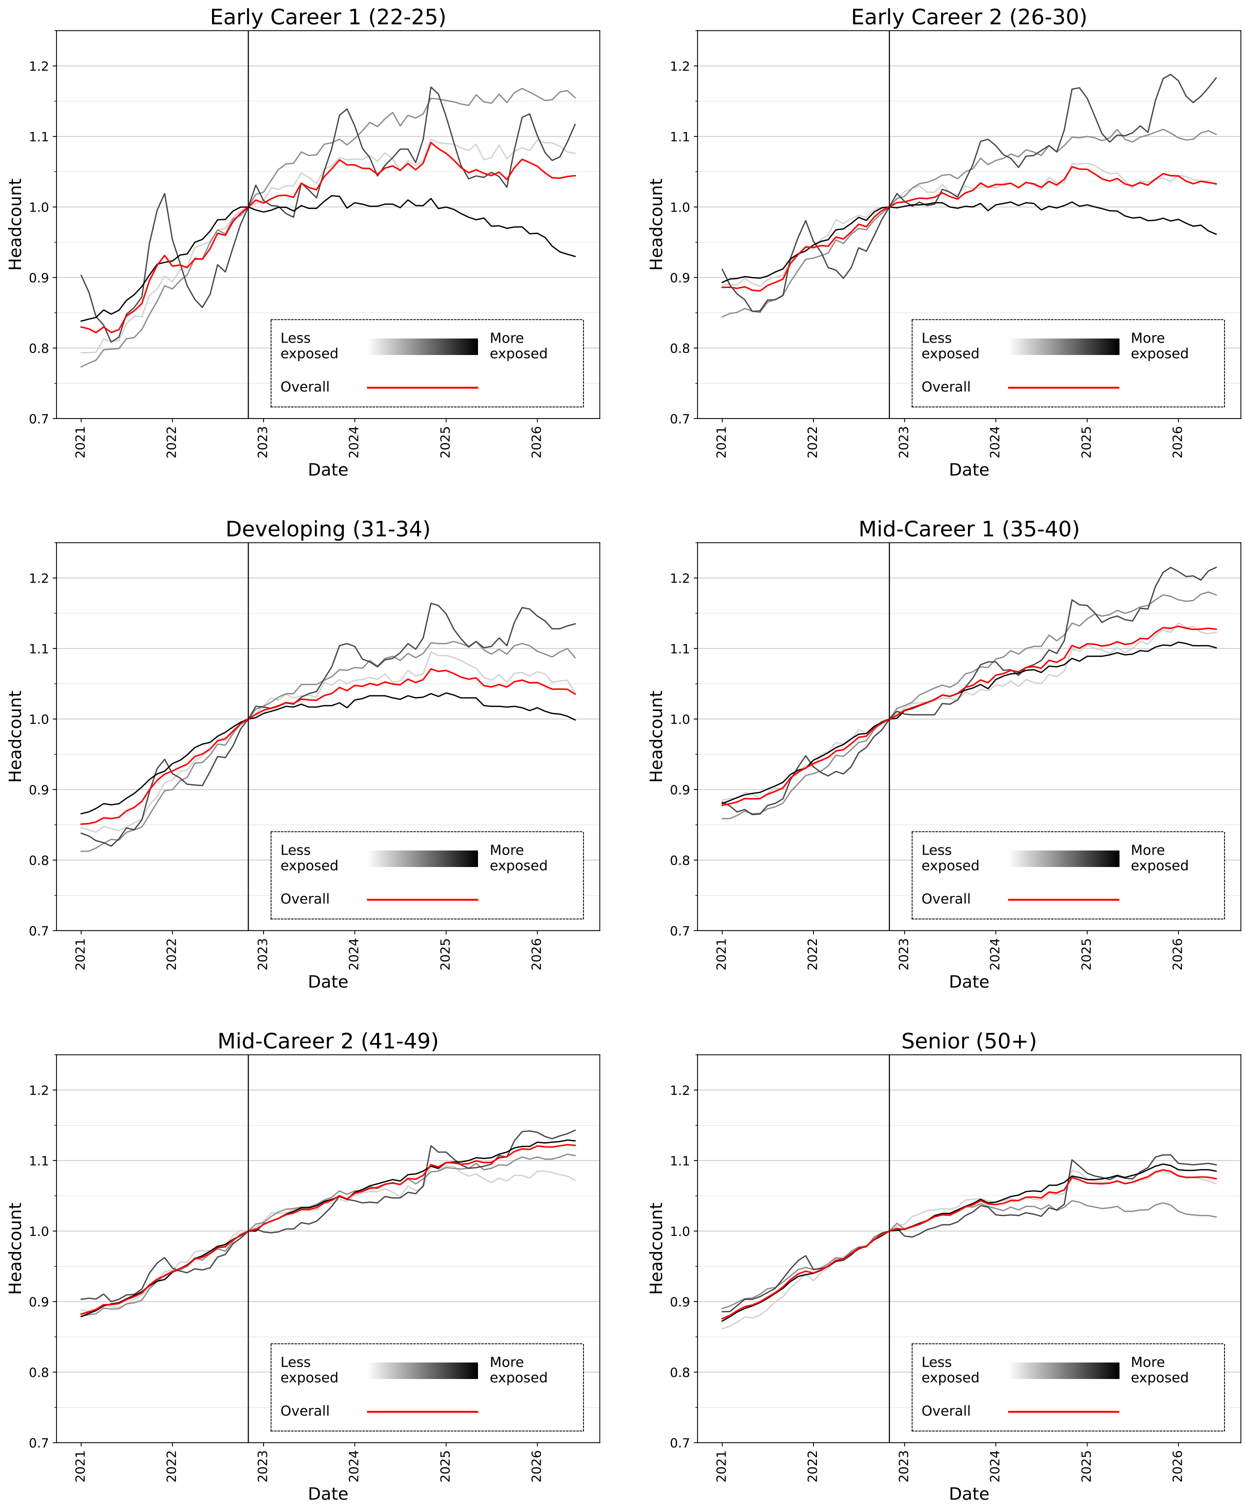

Figure C.3: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only occupations with above median interest rate exposure using data from Zens et al. (2020). Only a few occupations with high interest rate exposure also have high AI exposure. All occupations in the fourth and fifth quintile are consequently grouped together in level 4. The remaining quintiles are coded as 1, 2, and 3.

C.2 Excluding Technology Occupations and Firms

Employment changes by age and exposure quintile, excluding computer occupations (Fig. C.4) and excluding firms in information technology or computer systems design (Fig. C.5).

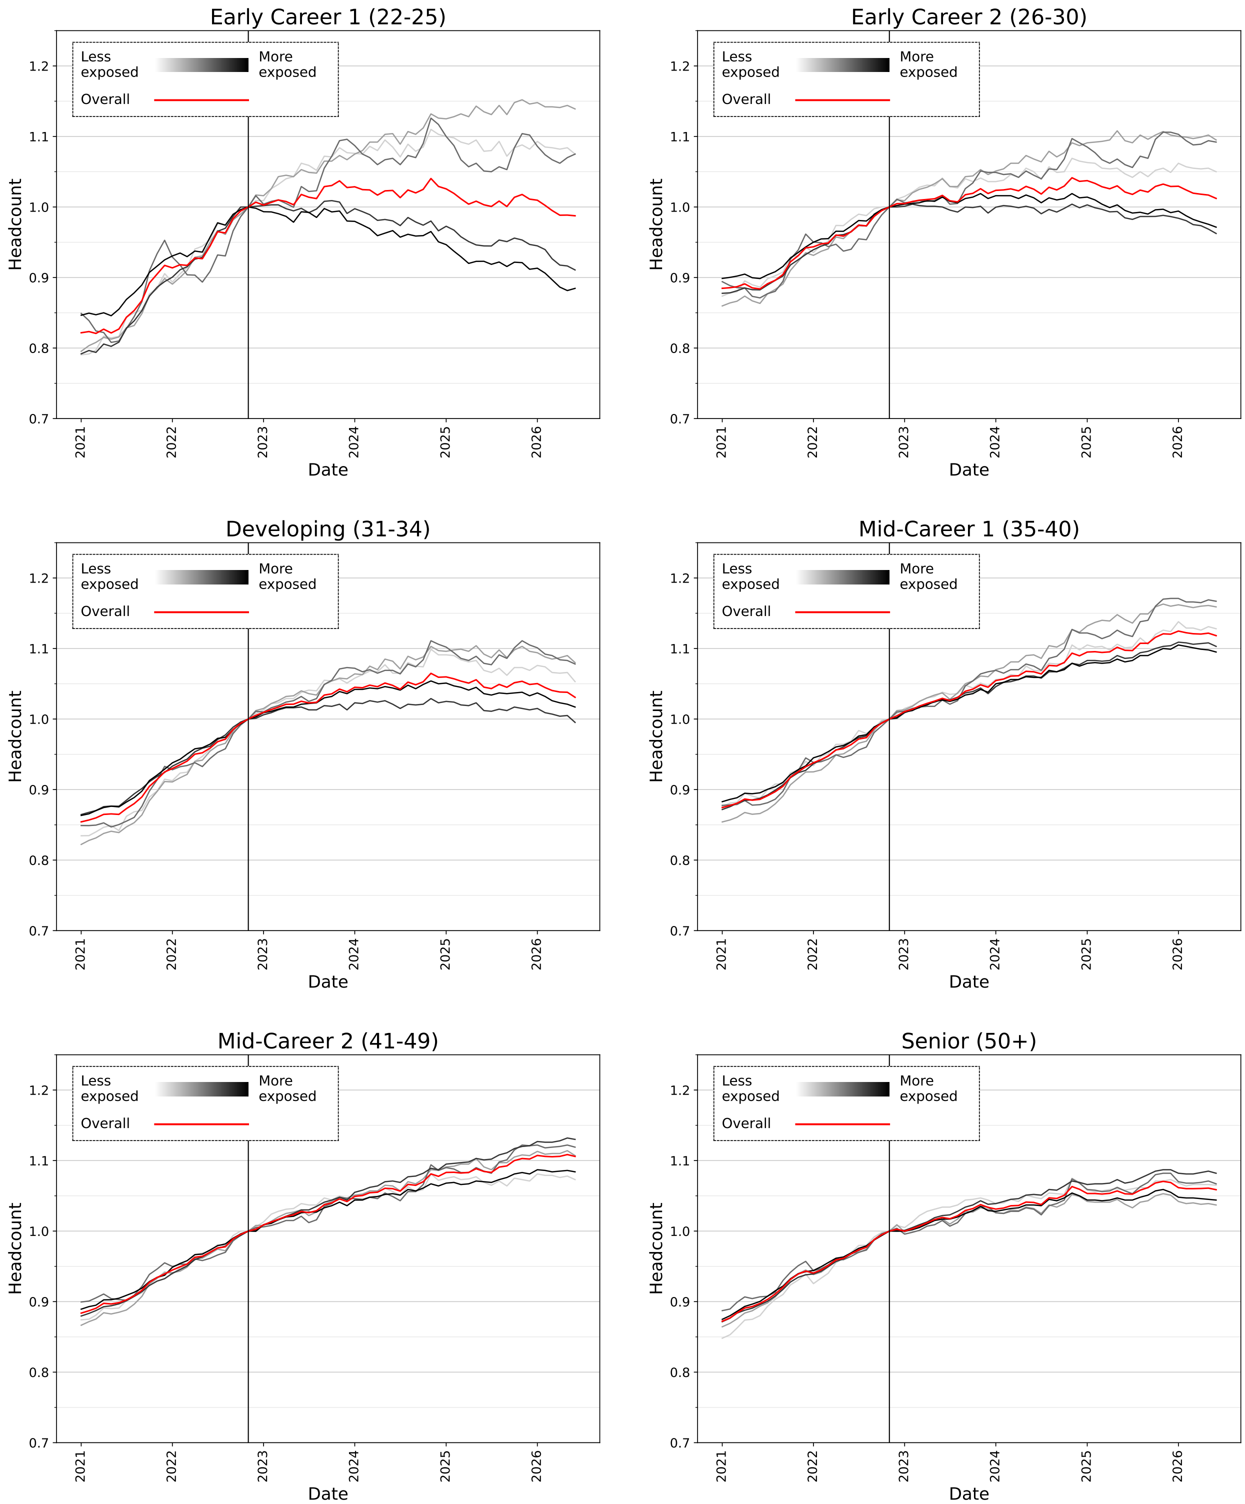

Figure C.4: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Excluding computer occupations (2010 SOC codes starting with 15-1).

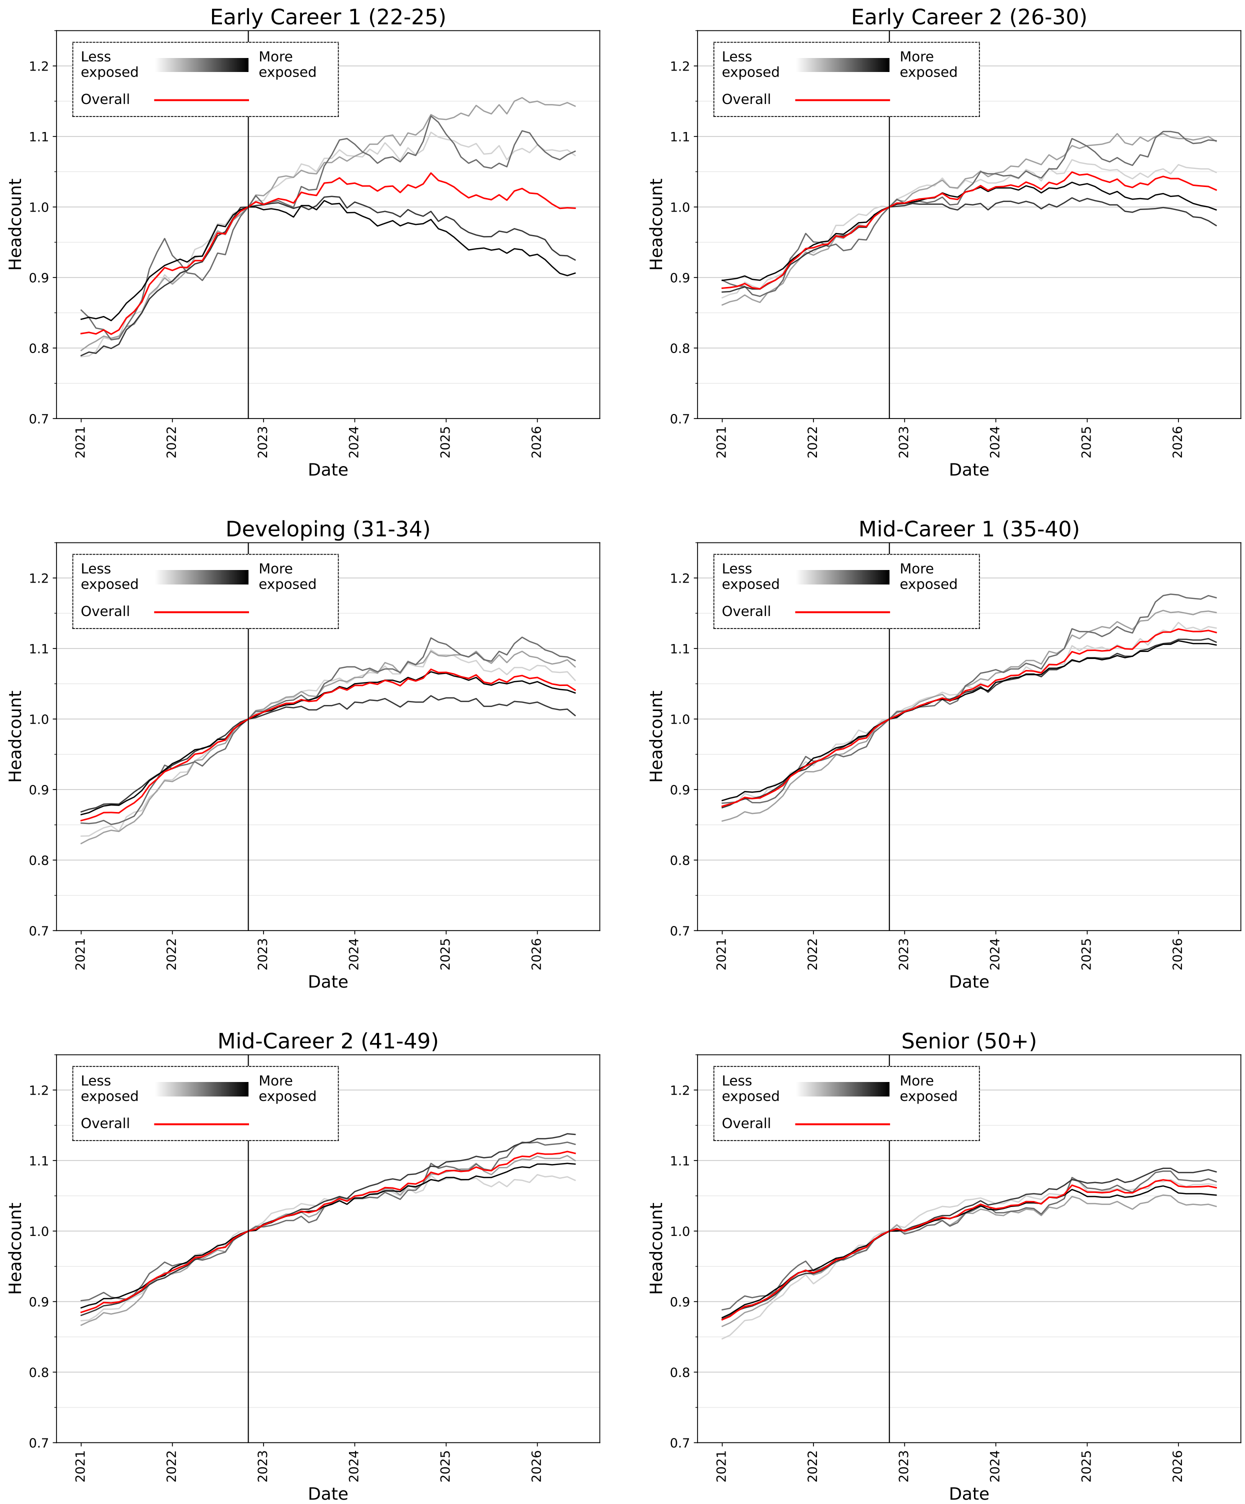

Figure C.5: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Excluding technology firms (NAICS codes 51 and 5415).

Table C.1 reports the corresponding occupation-level long-difference regression estimates for workers aged 22–25. These subsamples are estimated on the panel of firms balanced since January 2021, for which the corresponding full-sample estimate for the most exposed quintile is −0.193 (Table 2).

Table C.1: Occupation-level long-difference estimates for workers aged 22–25 excluding computer occupations (2010 SOC codes beginning 15-1) and technology firms (NAICS 51 and 5415). Heteroskedasticity-robust standard errors in parentheses.

Exposure quintile	Excl. computer occupations	Excl. technology firms
Quintile 2	0.064 (0.046)	0.071 (0.045)
Quintile 3	0.000 (0.040)	0.007 (0.040)
Quintile 4	-0.164^{***} (0.038)	-0.147^{***} (0.039)
Quintile 5	-0.190^{***} (0.034)	-0.166^{***} (0.033)
Occupations	616	621

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Coefficients relative to quintile 1.

C.3 Changes in Education

Employment changes by age and exposure quintile, computed separately for college-intensive occupations, in which at least 70% of workers hold a college degree (Fig. C.6), and for occupations in which at most 30% do (Fig. C.7).

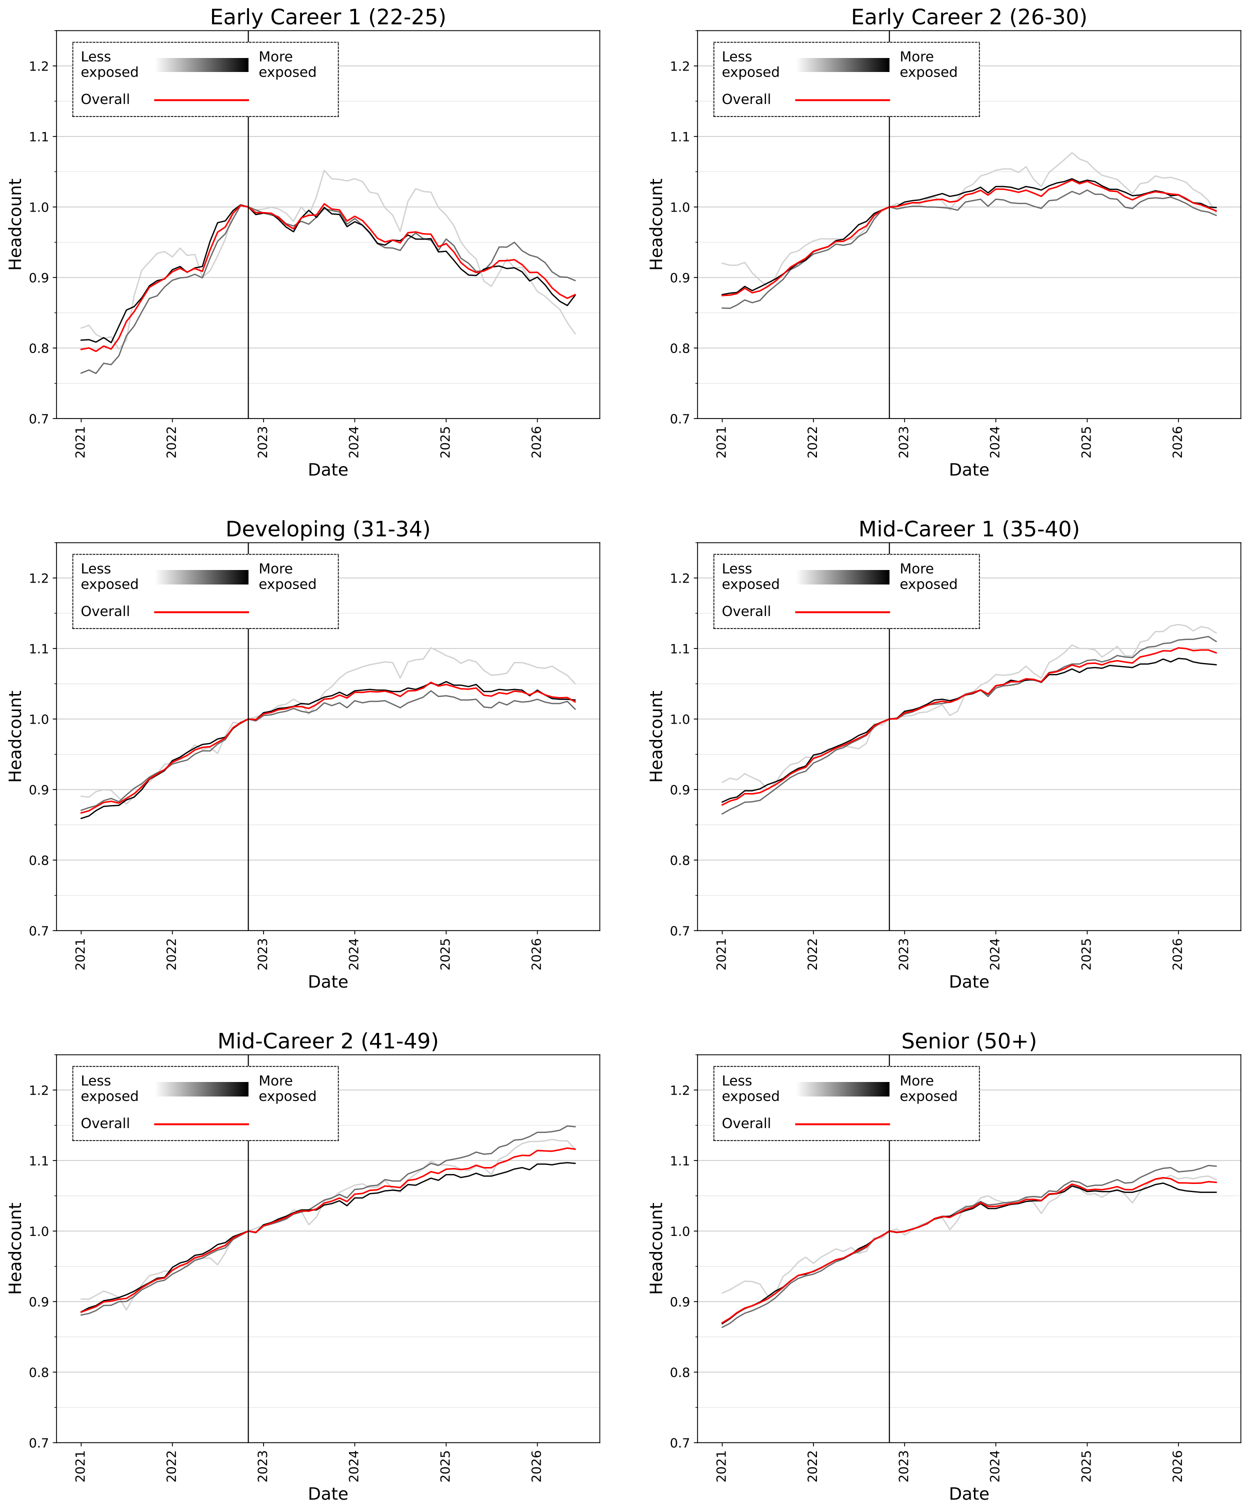

Figure C.6: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only occupations in which at least 70% of workers have a college degree in the 2017 ACS. Note that no such occupations lie in quintile 1 of the GPT-4 β based exposure measure.

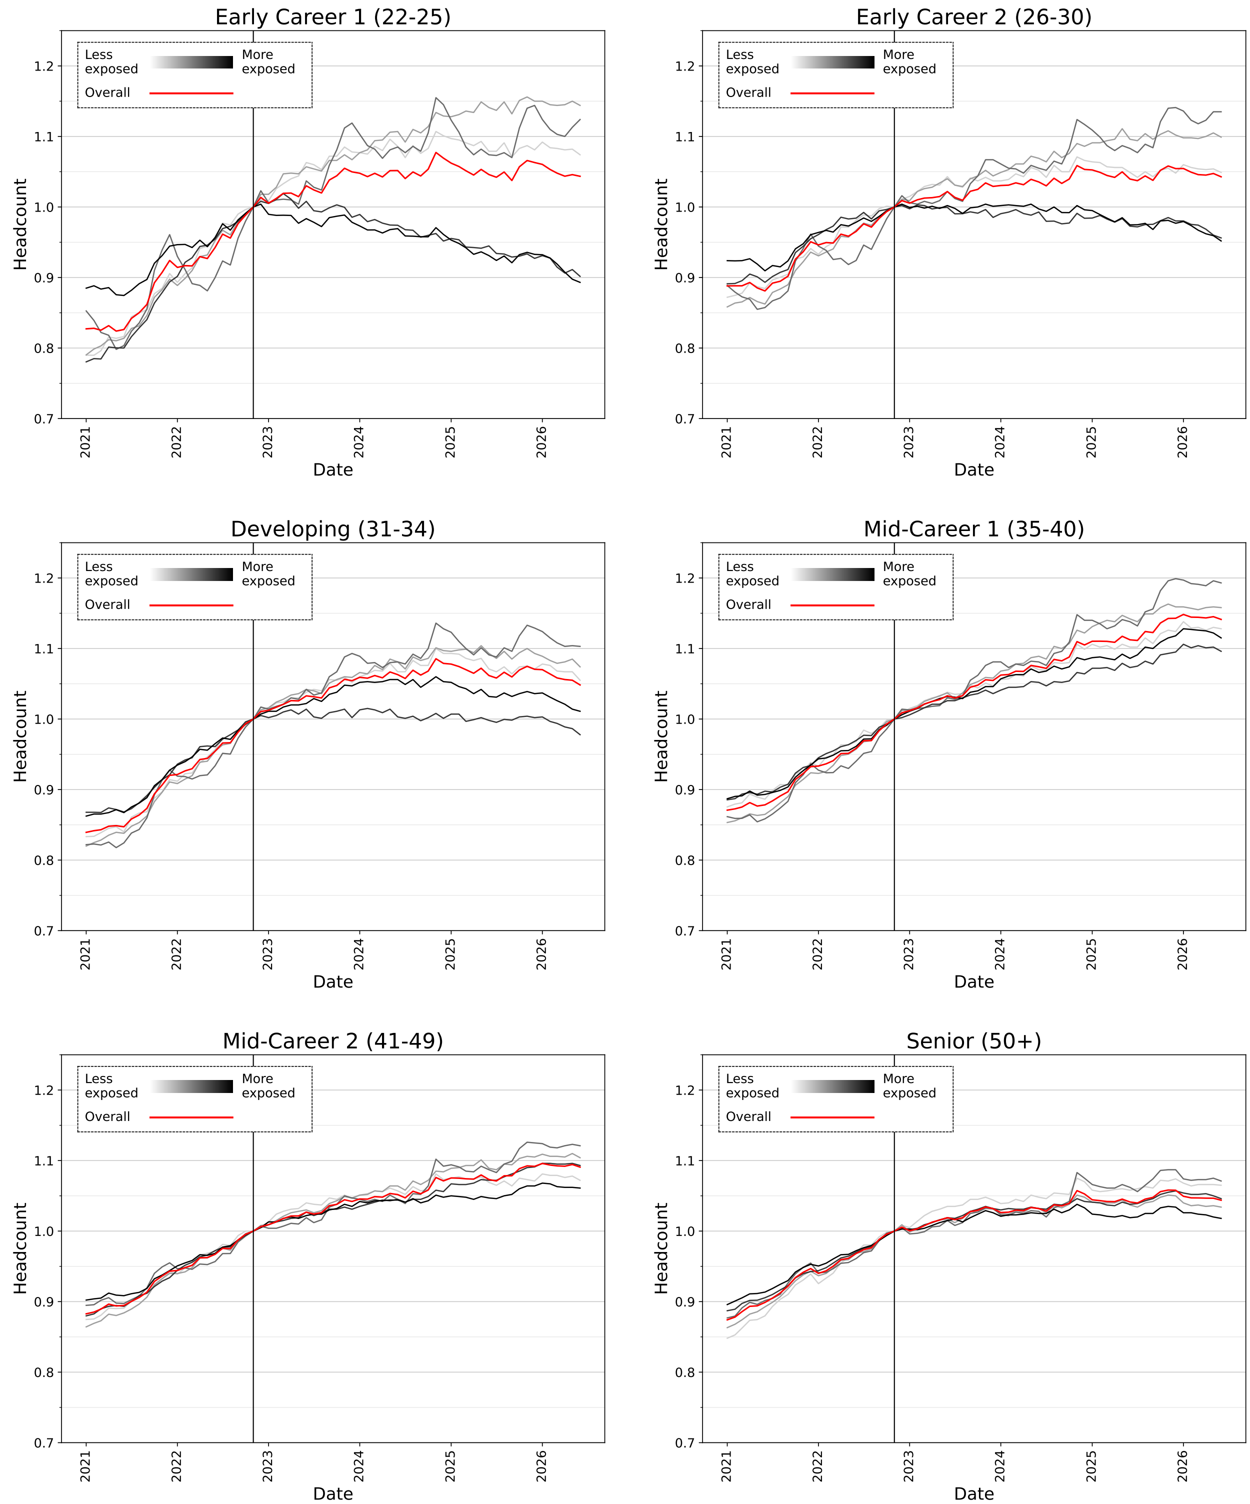

Figure C.7: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only occupations in which at most 30% of workers have a college degree in the 2017 ACS.

C.4 Remote Work

Employment changes by age and exposure group, computed separately for teleworkable occupations (Fig. C.8) and non-teleworkable occupations (Fig. C.9), classified following Dingel and Neiman (2020).

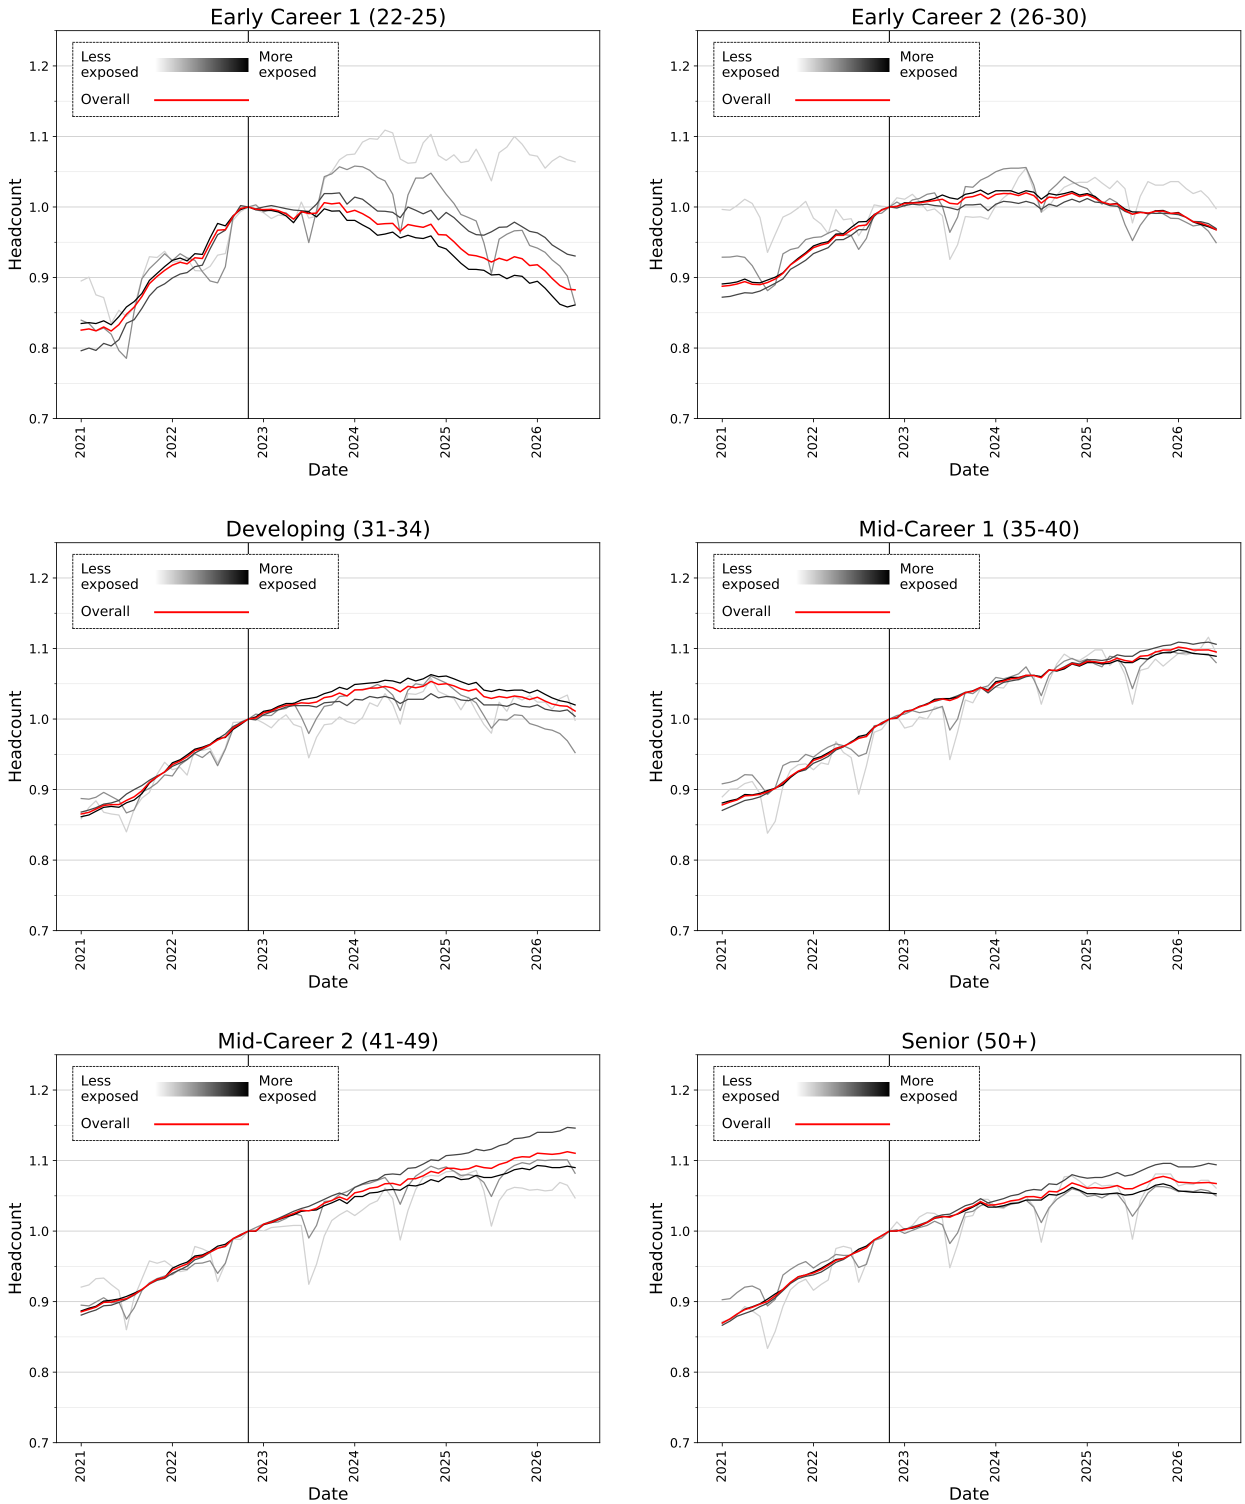

Figure C.8: Employment changes by age and exposure group using measures from Eloundou et al. (2024). Including only teleworkable occupations according to Dingel and Neiman (2020). Note that very few teleworkable occupations fall in the lowest exposure quintile. All occupations in the first and second quintile are consequently grouped together in level 1. The remaining quintiles are coded as 2, 3, and 4.

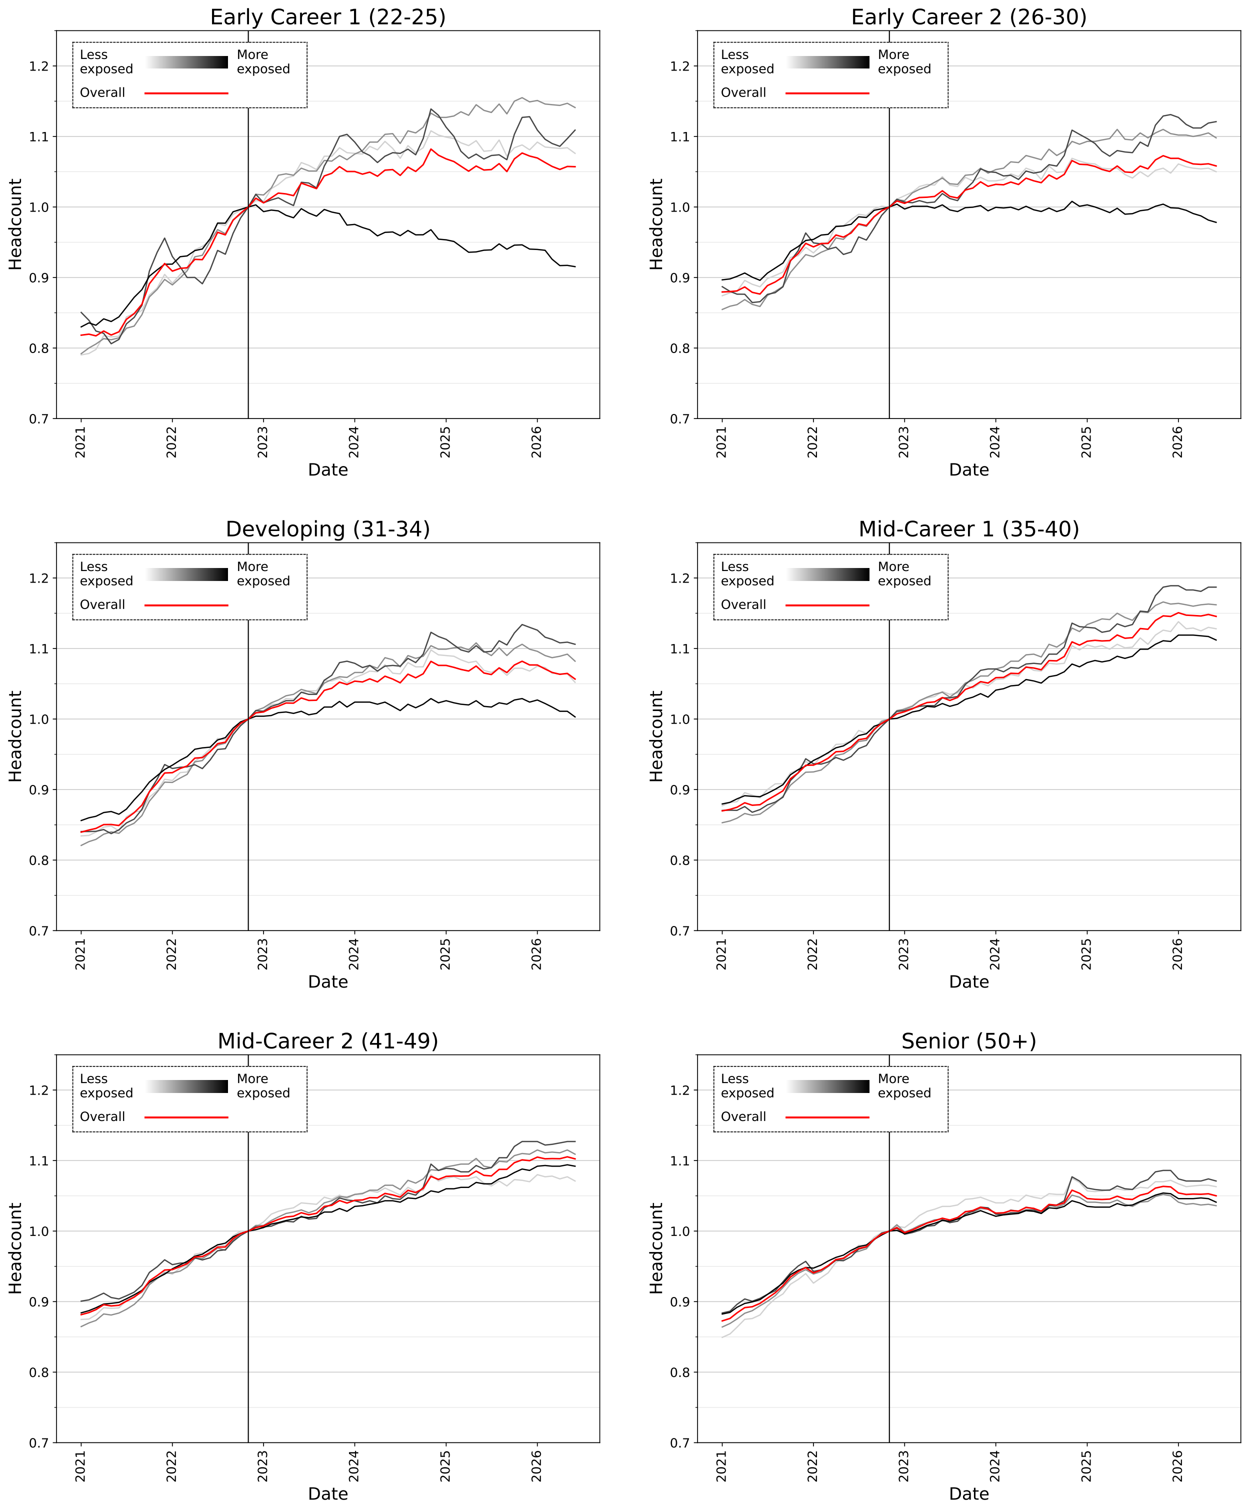

Figure C.9: Employment changes by age and exposure group using measures from Eloundou et al. (2024). Including only non-teleworkable occupations according to Dingel and Neiman (2020). Note that very few non-teleworkable occupations fall in the highest exposure quintile. All occupations in the fourth and fifth quintile are consequently grouped together in level 4. The remaining quintiles are coded as 1, 2, and 3.

C.5 Unbalanced Sample and Part-Time Workers

Employment changes by age and exposure quintile using the full, unbalanced sample of firms (Fig. C.10) and including part-time and temporary workers (Fig. C.11).

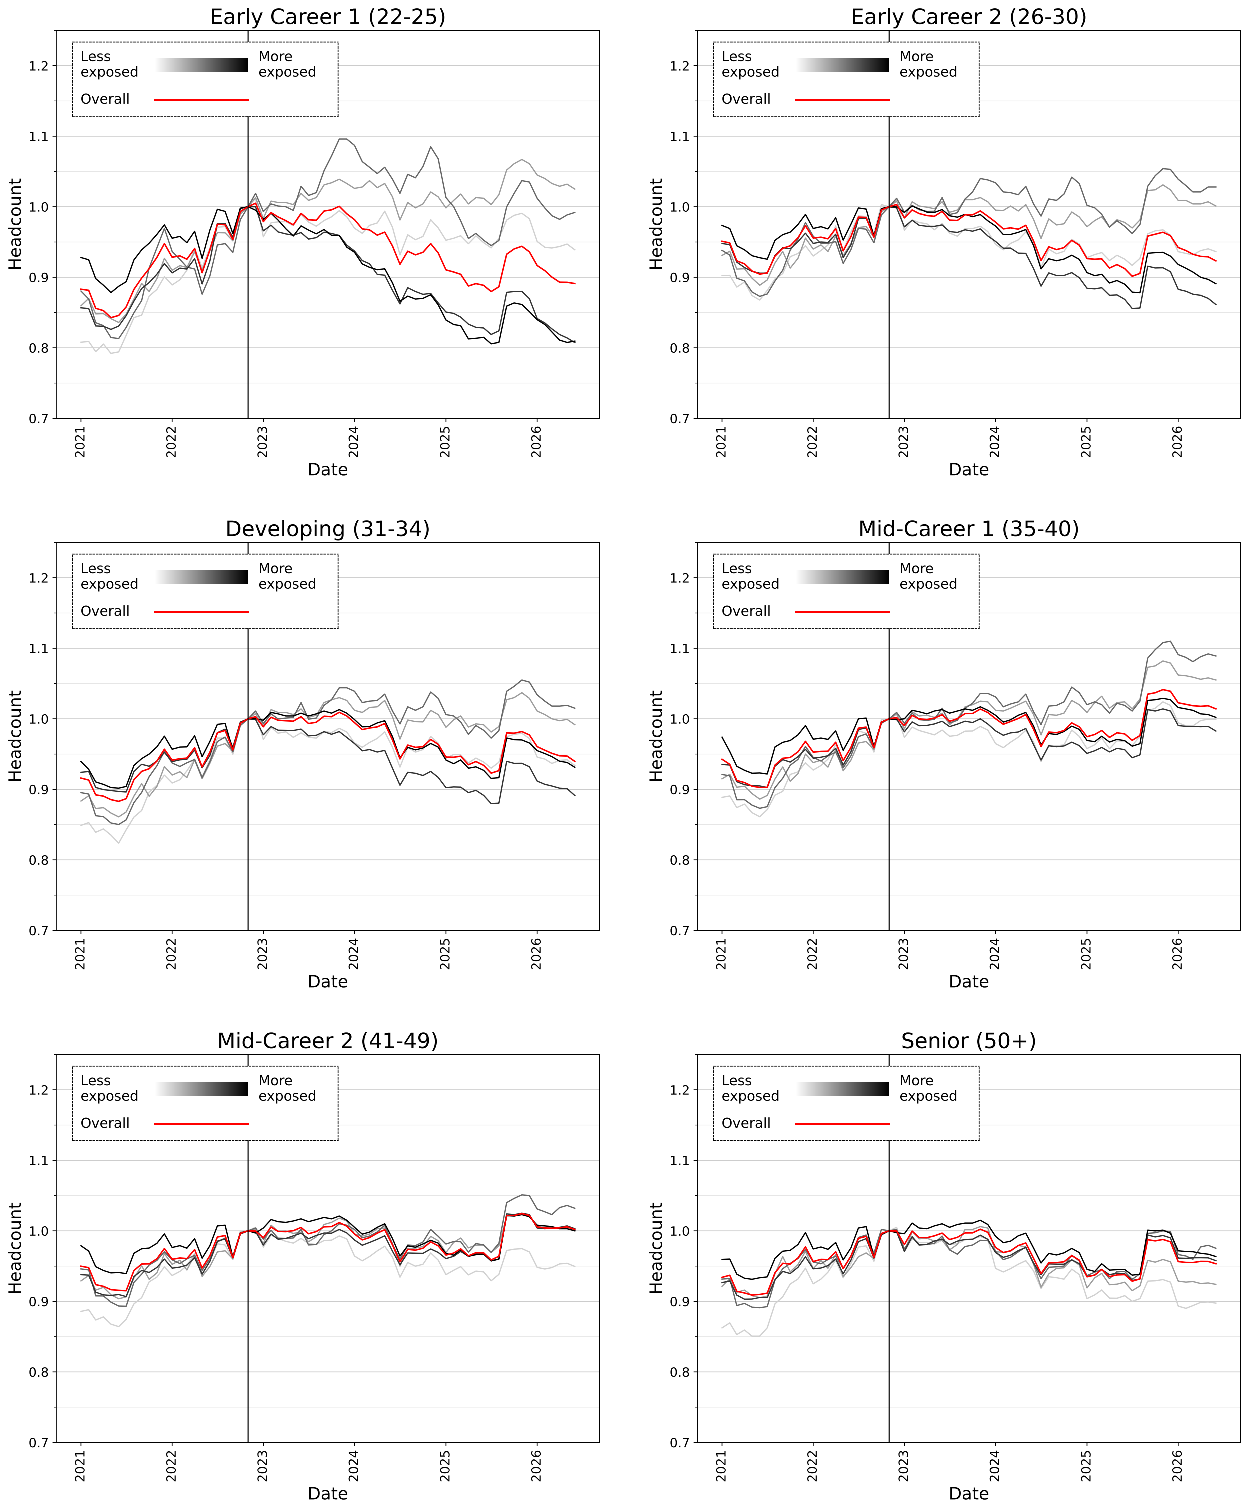

Figure C.10: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Using the full sample of firms.

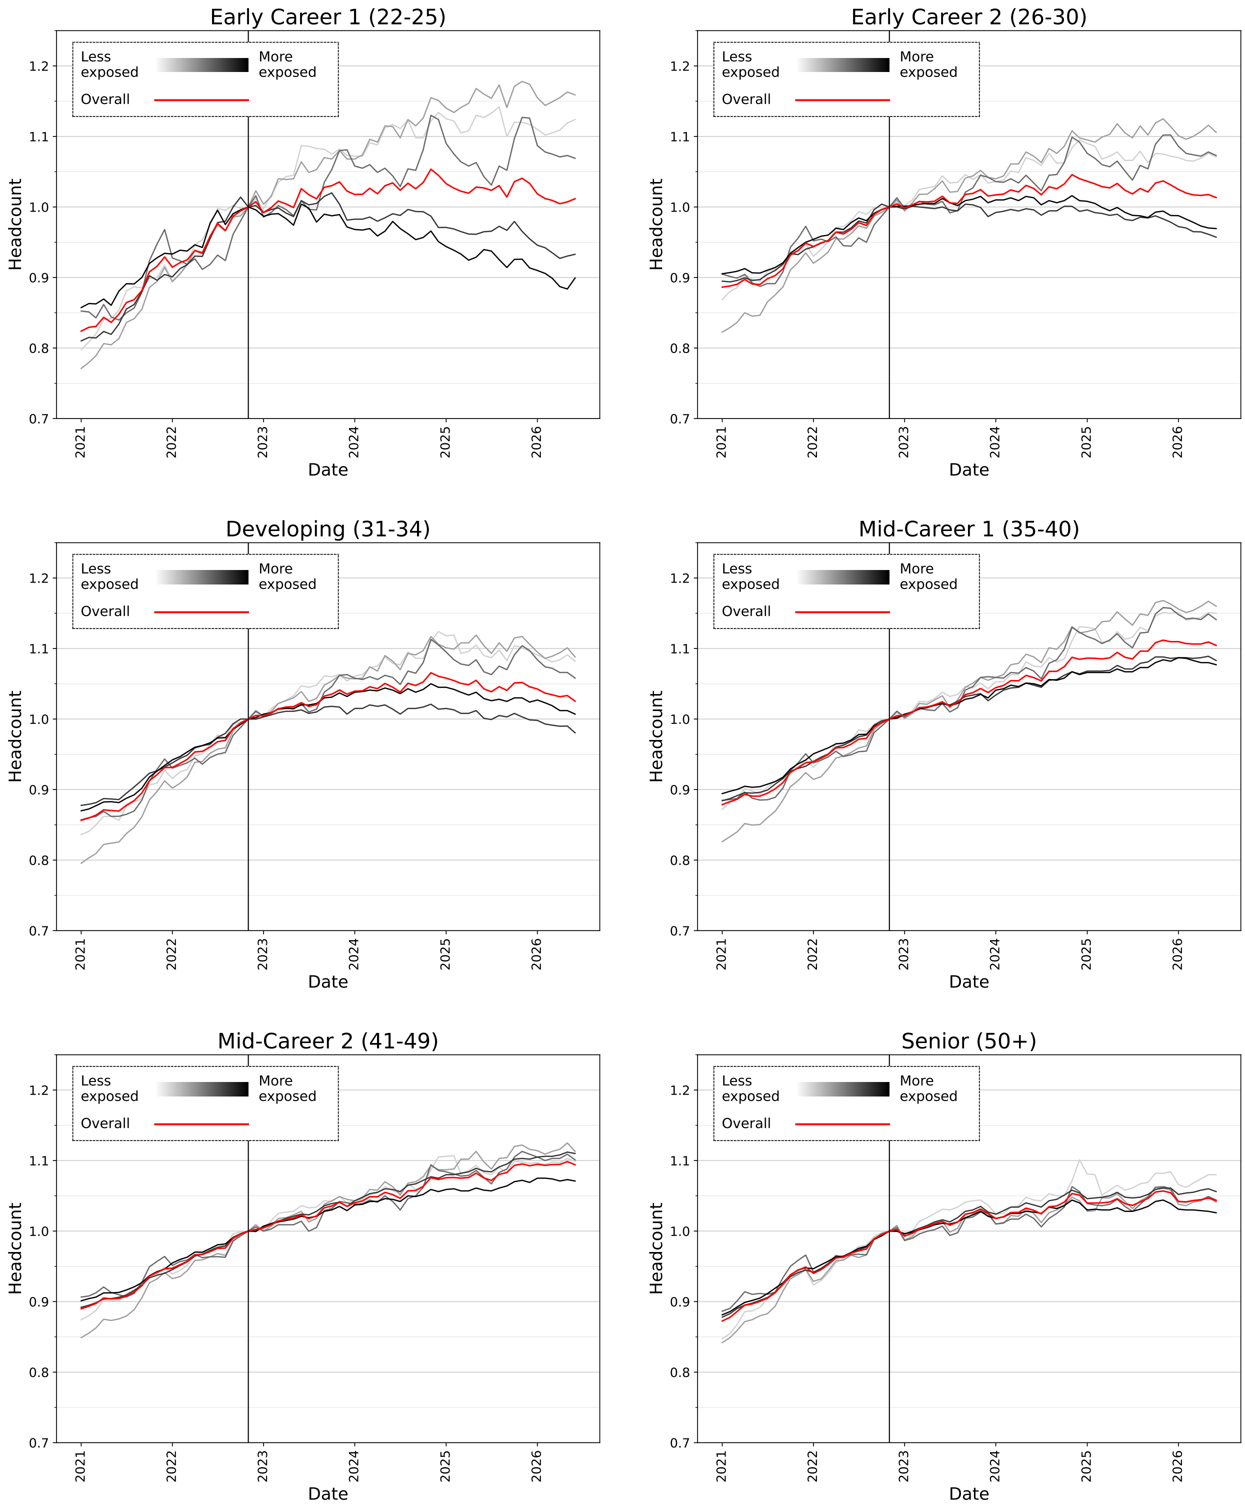

Figure C.11: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Including part-time and temporary workers.

C.6 Within-Firm Event Study Estimates

This section reports Poisson event study estimates of within-firm employment by AI exposure quintile. For each age group, we estimate

$$\begin{aligned}\log(E[y_{f,q,t}]) = \sum_{q' \neq 1} \sum_{j \neq -1} \gamma_{q',j} 1\{t = j\} 1\{q' = q\} \\ + \alpha_{f,q} + \beta_{f,t} + \epsilon_{f,q,t}\end{aligned} \tag{C.1}$$

where f indexes firms, q indexes exposure quintiles (Eloundou et al., 2024), and t indexes months, with $t = -1$ corresponding to November 2022. The outcome variable $y_{f,q,t}$ is employment in f, q, t. The firm-time effects $\beta_{f,t}$ absorb aggregate firm shocks that impact each exposure quintile equally. The firm-quintile effects $\alpha_{f,q}$ adjust for baseline differences in hiring across quintiles within the firm. The coefficients of interest, $\gamma_{q,t}$, measure differential changes in employment growth across quintiles after accounting for both sets of fixed effects.[41]

We run this regression separately for each age group. We consider different specifications. Our first specification applies minimal sample filters: we require only that the Poisson estimator be well defined, keeping firms with at least one worker in each firm-month and each firm-quintile. We also report more conservative specifications that restrict to firms hiring at least 10 workers within the age group in every period and with $\sum_t y_{f,q,t} \geq 10$, 50, or 100 for each q. These stricter filters shrink the sample and skew it toward larger firms. Standard errors are clustered by firm in both cases.

Results under the minimal-filter specification are in Fig. C.12. For workers aged 22–25, employment in the two most exposed quintiles declines relative to quintile 1 by June 2026: about 7 log points for quintile 5 ($p = 0.03$) and about 11 log points for quintile 4 ($p = 0.003$). Estimates for other age groups are smaller in magnitude and are either statistically insignificant or, where significant, much smaller than for 22–25 or of the opposite sign. Fig. C.13 reports the more conservative min-10-filter specification: quintile 5 declines -0.07 and quintile 4 -0.11, though results are less precisely estimated. Fig. C.14 raises the requirement to at least 50 workers per exposure-quintile cell, finding an insignificant decline of -0.06 for quintile 5 and a significant decline of -0.12 for quintile 4. Fig. C.15 raises the requirement to at least 100 workers per exposure-quintile

[41] Because of zero counts in the outcome variable, we estimate a Poisson regression instead of an OLS regression in logs (Chen and Roth, 2024).

cell, finding an insignificant decline of −0.05 for quintile 5 and a significant decline of −0.13 for quintile 4.

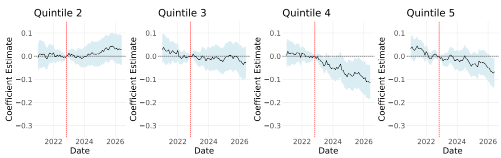

(a) Age 22–25

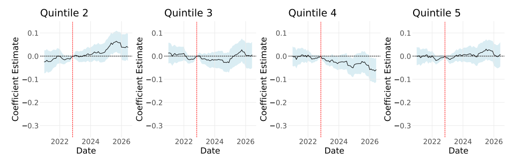

(b) Age 26–30

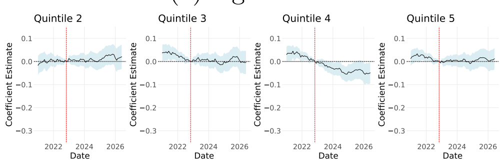

(c) Age 31–34

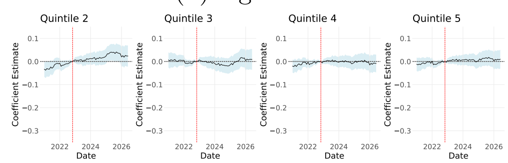

(d) Age 35–40

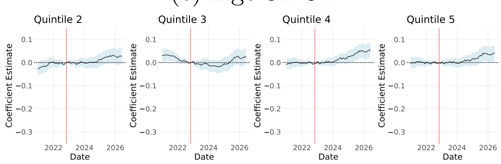

(e) Age 41–49

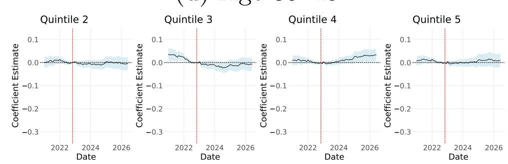

(f) Age 50+

Figure C.12: Within-firm Poisson event study estimates for employment changes by age and AI exposure, under the minimal-filter sample (firms need only one worker per firm-month and firm-quintile). Estimates are all relative to occupational exposure quintile 1. Exposure quintiles use Eloundou et al. (2024) GPT-4 β measures. Estimates control for firm-time and firm-quintile fixed effects. Shaded regions are 95% confidence intervals. Standard errors are clustered by firm.

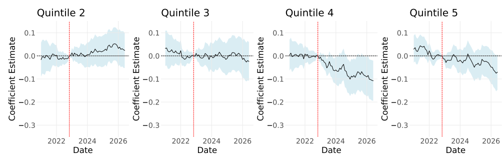

(a) Age 22–25

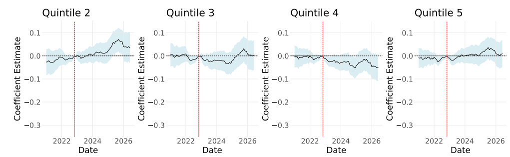

(b) Age 26–30

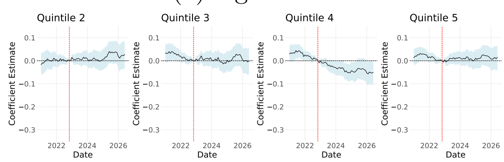

(c) Age 31–34

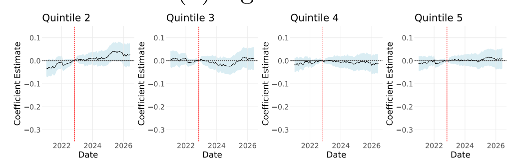

(d) Age 35–40

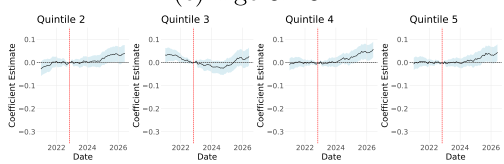

(e) Age 41–49

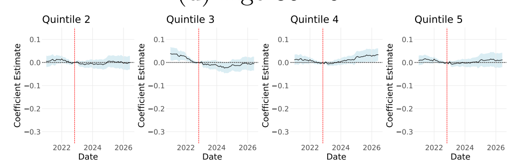

(f) Age 50+

Figure C.13: Within-firm Poisson event study estimates restricting to firms with at least 10 workers in the age group in every period and $\sum_t y_{f,q,t} \geq 10$ for each quintile. Estimates are relative to exposure quintile 1 (Eloundou et al., 2024) and control for firm-time and firm-quintile fixed effects. Shaded regions are 95% confidence intervals; standard errors clustered by firm.

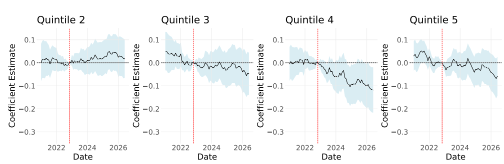

(a) Age 22–25

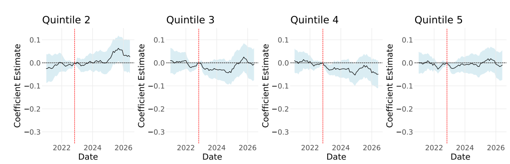

(b) Age 26–30

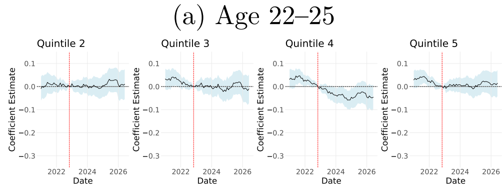

(c) Age 31–34

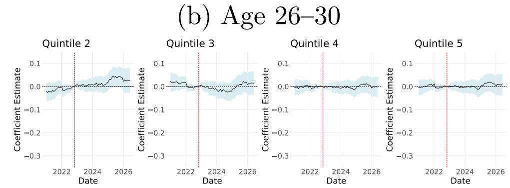

(d) Age 35–40

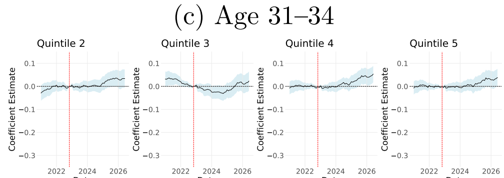

(e) Age 41–49

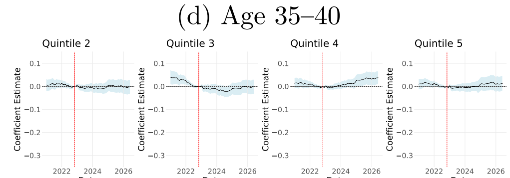

(f) Age 50+

Figure C.14: Within-firm Poisson event study estimates restricting to firms with at least 50 workers per exposure-quintile cell. Estimates are relative to exposure quintile 1 (Eloundou et al., 2024) and control for firm-time and firm-quintile fixed effects. Shaded regions are 95% confidence intervals; standard errors clustered by firm.

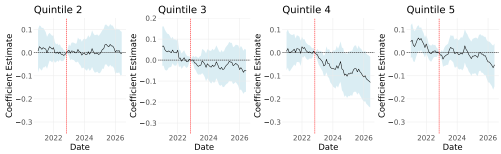

(a) Age 22–25

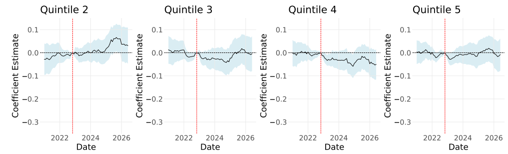

(b) Age 26–30

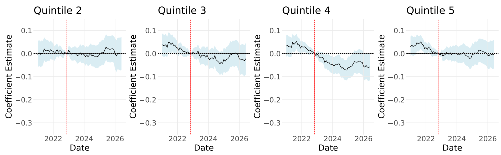

(c) Age 31–34

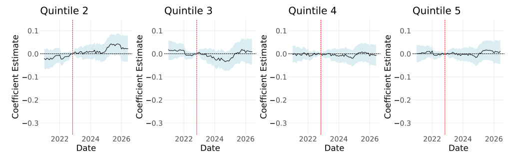

(d) Age 35–40

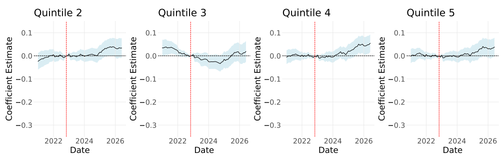

(e) Age 41–49

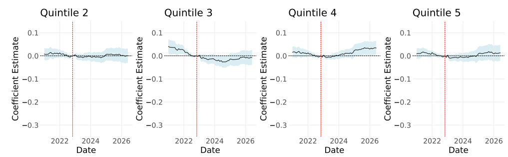

(f) Age 50+

Figure C.15: Within-firm Poisson event study estimates restricting to firms with at least 100 workers per exposure-quintile cell. Estimates are relative to exposure quintile 1 (Eloundou et al., 2024) and control for firm-time and firm-quintile fixed effects. Shaded regions are 95% confidence intervals; standard errors clustered by firm.

C.7 Occupation Event Study Estimates, by Age

Occupation event study estimates for the most exposed quintile (relative to quintile 1), by age group, under the four control sets from Fig. 4 in the main text. Each series is the occupation-level long-difference coefficient with the difference endpoint rolled forward over time. Estimated on the extended balanced panel of firms observed from January 2018.

D Additional Heterogeneity

D.1 Men and Women

Employment changes by age and exposure quintile, computed separately for men (Fig. D.1) and women (Fig. D.2).

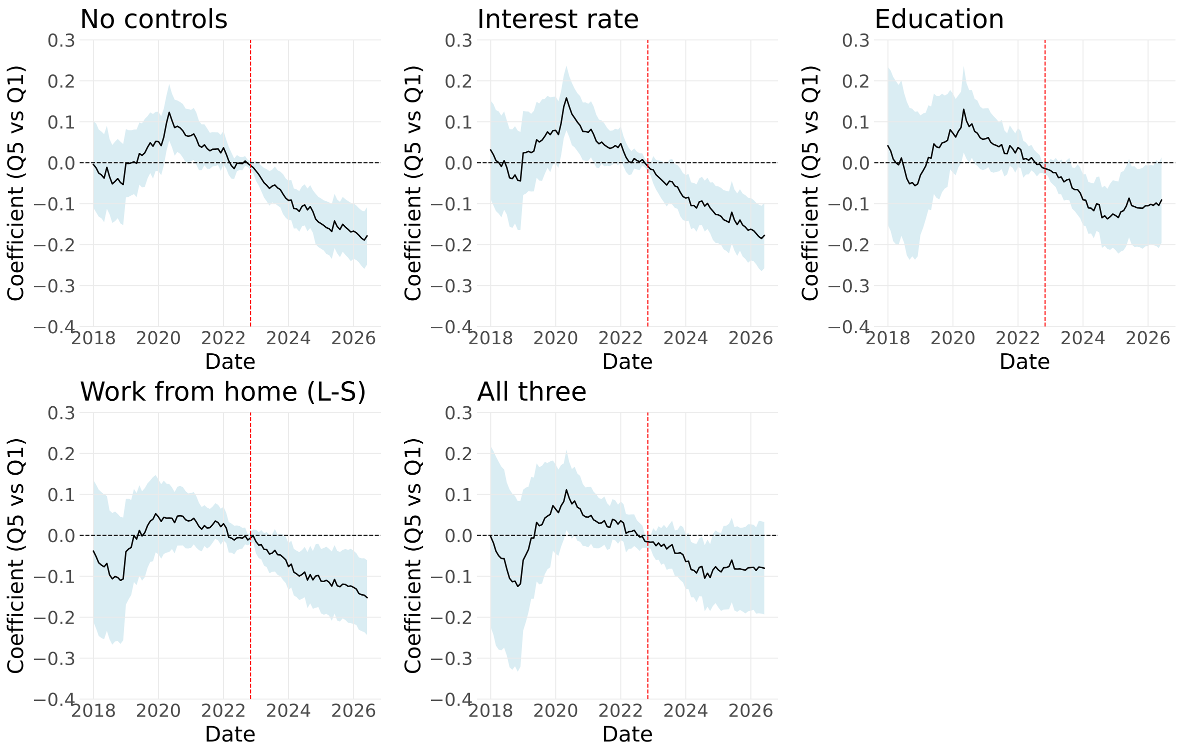

Figure C.16: Occupation event study estimates for workers aged 22–25 (reproduces main-text Fig. 4 for comparison with the other age groups).

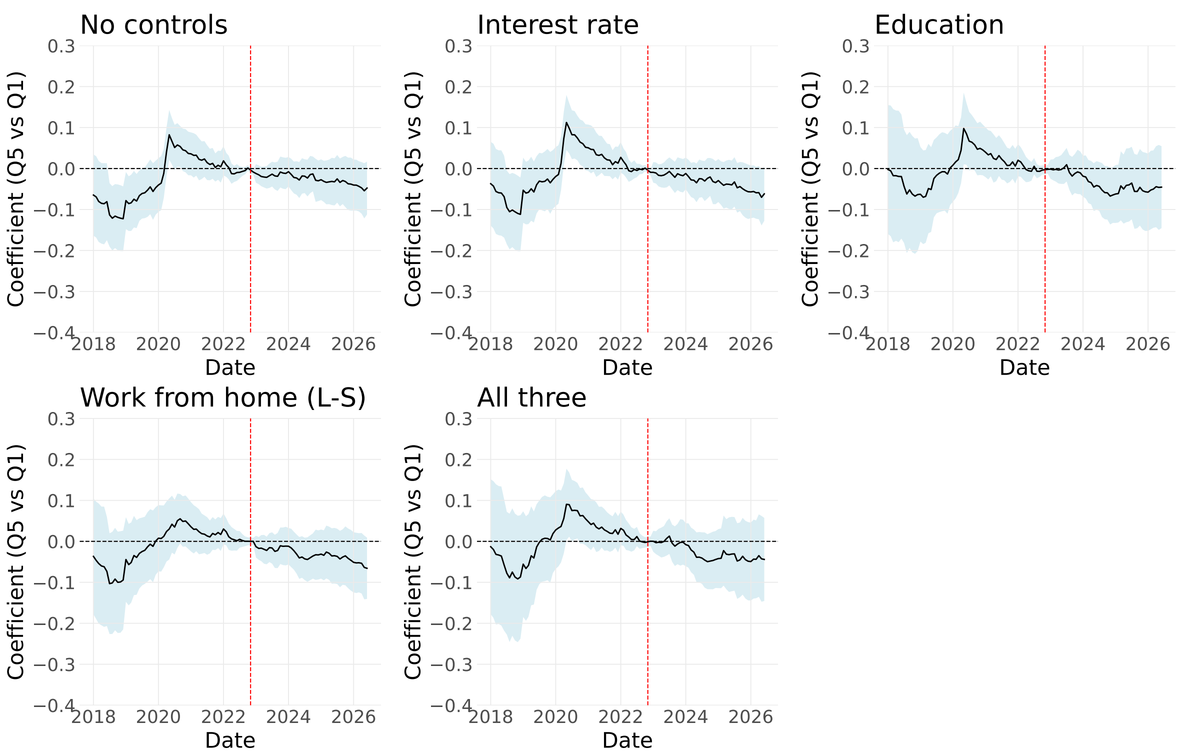

Figure C.17: Occupation event study estimates for workers aged 26–30.

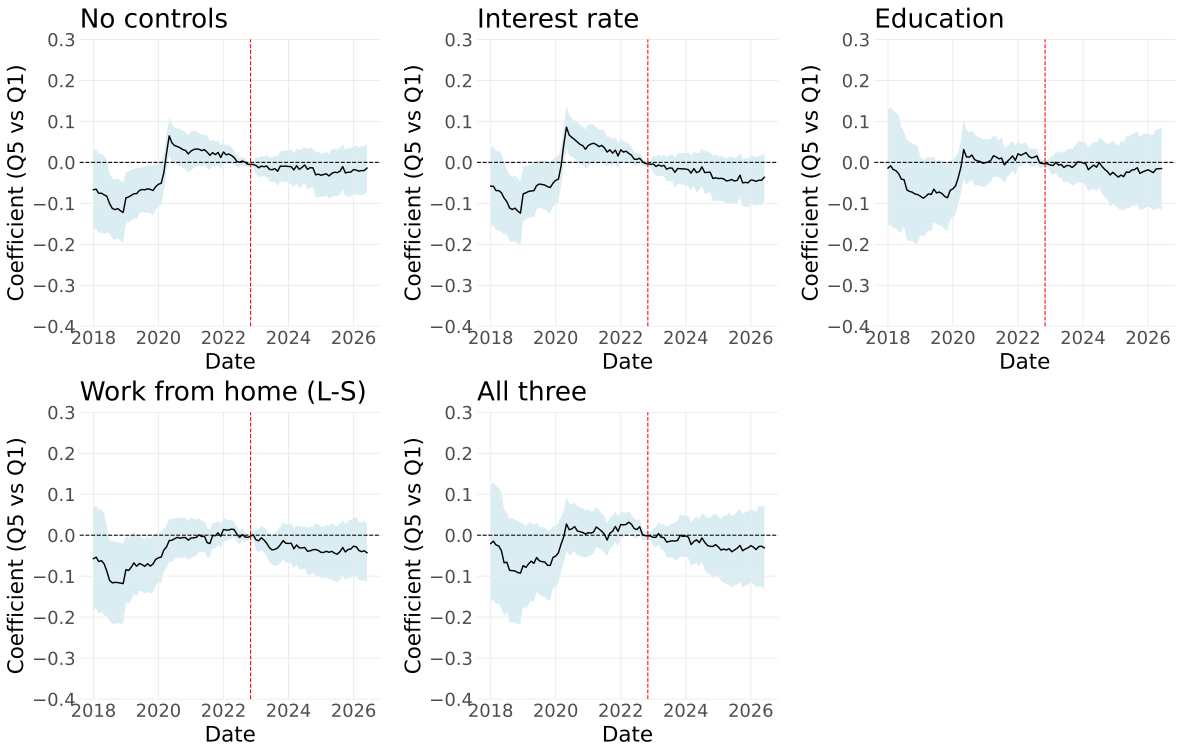

Figure C.18: Occupation event study estimates for workers aged 31–34.

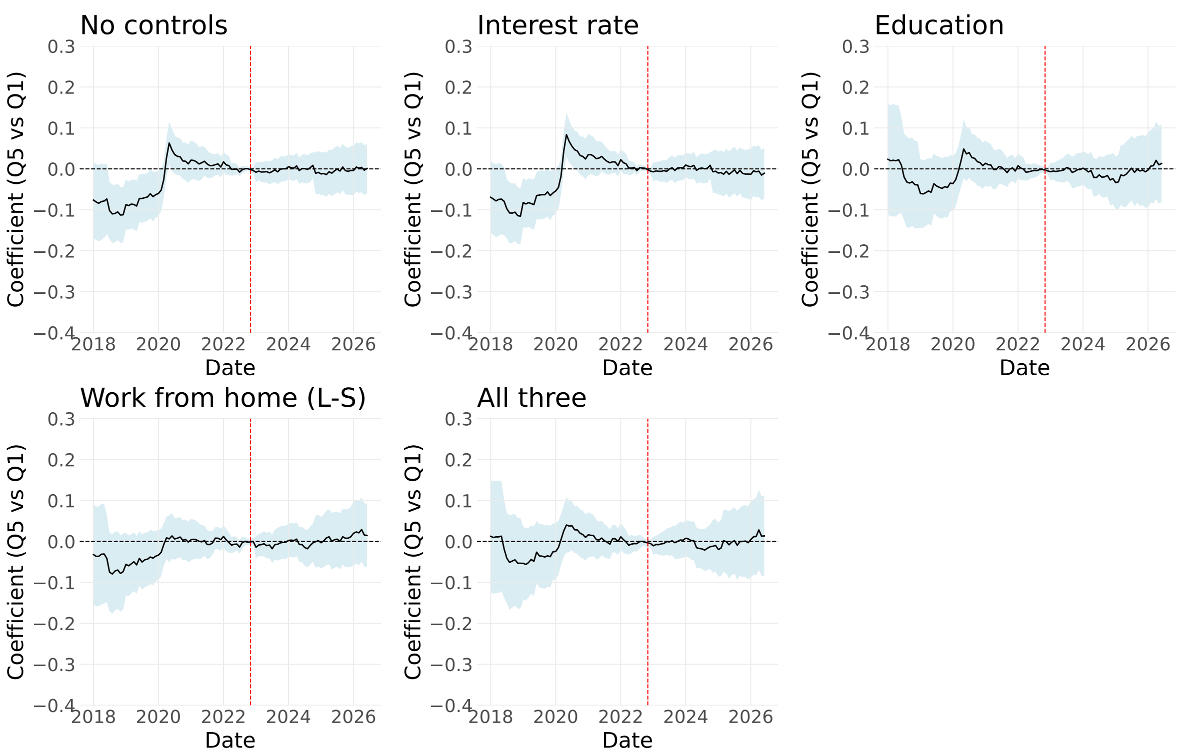

Figure C.19: Occupation event study estimates for workers aged 35–40.

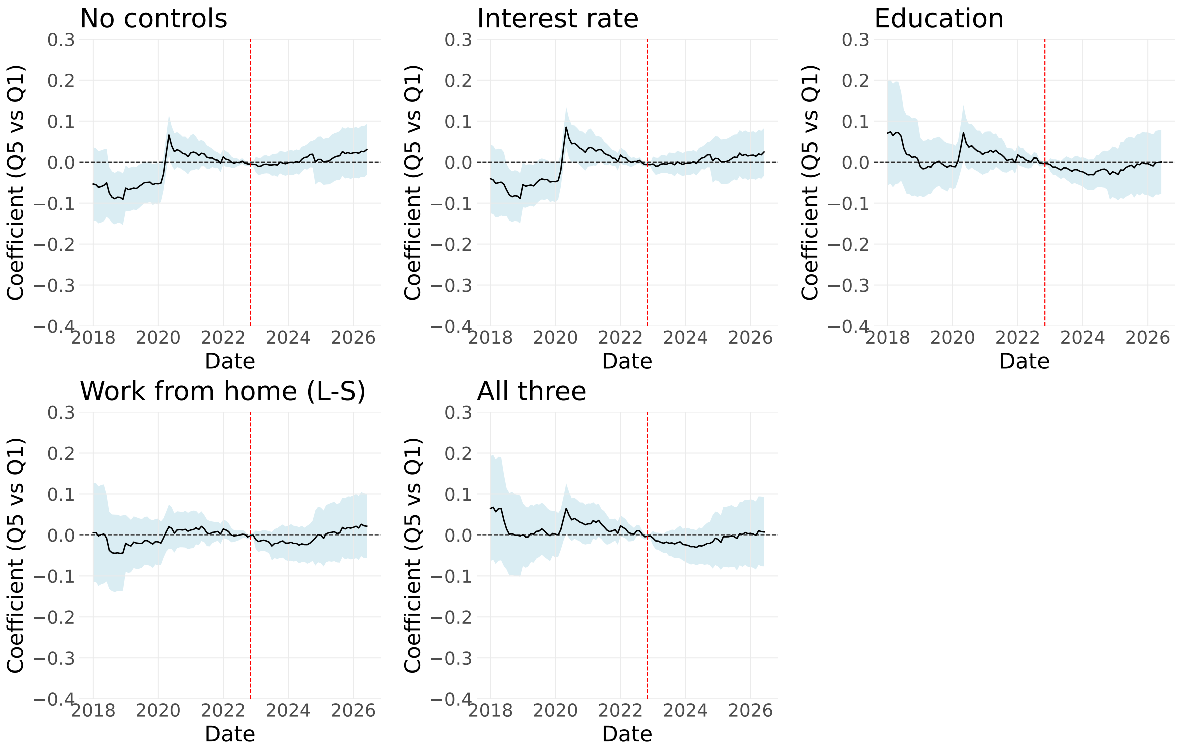

Figure C.20: Occupation event study estimates for workers aged 41–49.

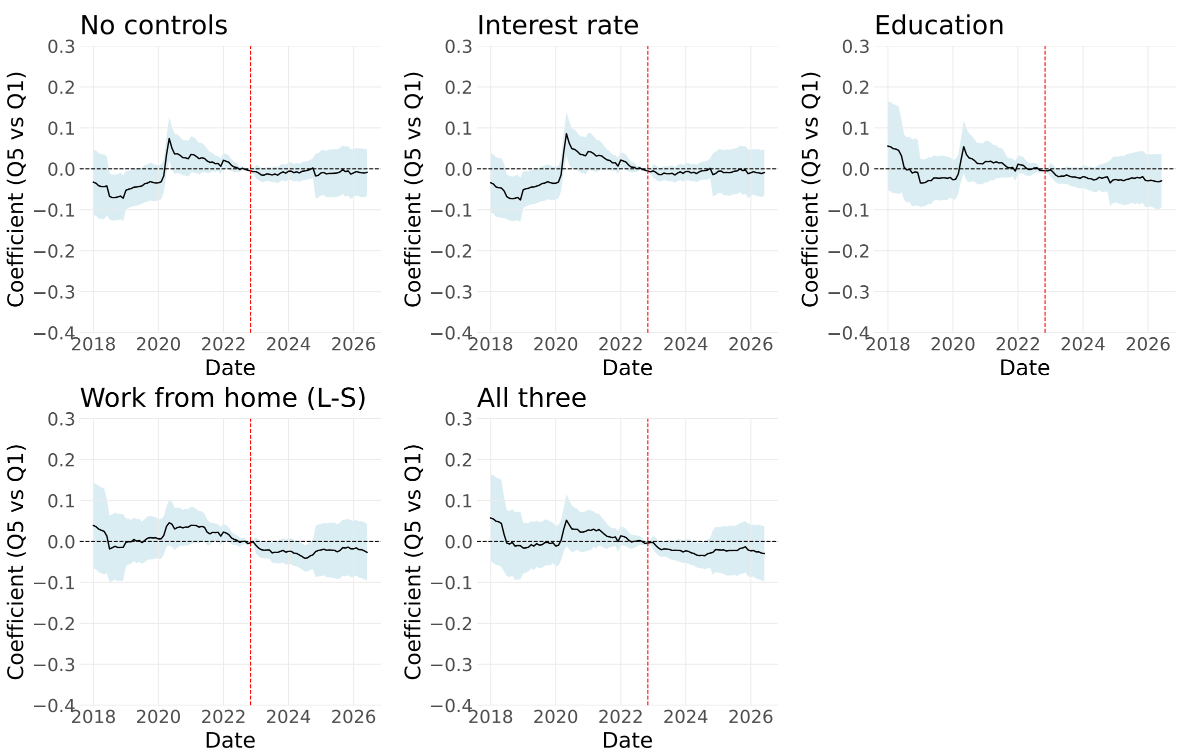

Figure C.21: Occupation event study estimates for workers aged 50+.

Table D.1: Share of each gender's employment in the top two AI-exposure quintiles (Q4 and Q5) by age band, 2022. Using GPT-4 β exposure measure from Eloundou et al. (2024).

	ADP		CPS		ACS	
Age band	Men	Women	Men	Women	Men	Women
22–25	50.8	65.8	31.4	37.8	34.0	41.4
26–30	59.3	71.1	40.2	44.7	42.8	49.6
31–34	62.7	73.3	41.6	48.0	47.1	51.7
35–40	65.1	74.0	42.6	48.2	48.2	52.4
41–49	65.0	73.3	42.4	48.8	48.2	52.6
50+	62.1	71.7	42.5	49.5	47.5	52.9

Table D.2: Five largest occupations by ADP employment within each AI exposure quintile in 2022, separately for men and women. Exposure quintiles are defined based on the GPT-4 β measures from Eloundou et al. (2024).

Quintile	Rank	Men	Women
Q1	1	Laborers and Freight, Stock, and Material Movers, Hand	Maids and Housekeeping Cleaners
	2	Industrial Truck and Tractor Operators	Laborers and Freight, Stock, and Material Movers, Hand
	3	Landscaping and Groundskeeping Workers	Packaging and Filling Machine Operators and Tenders
	4	Packaging and Filling Machine Operators and Tenders	Home Health Aides
	5	Maids and Housekeeping Cleaners	Janitors and Cleaners, Except Maids and Housekeeping Cleaners
Q2	1	Maintenance and Repair Workers, General	Nursing Assistants
	2	Helpers–Production Workers	Helpers–Production Workers
	3	Stock Clerks- Stockroom, Warehouse, or Storage Yard	Maintenance and Repair Workers, General
	4	Welders, Cutters, and Welder Fitters	Stock Clerks- Stockroom, Warehouse, or Storage Yard
	5	Machinists	Waiters and Waitresses
Q3	1	Taxi Drivers and Chauffeurs	Registered Nurses
	2	Heavy and Tractor-Trailer Truck Drivers	Medical Assistants
	3	Retail Salespersons	Social and Community Service Managers
	4	Computer Numerically Controlled Machine Tool Programmers, Metal and Plastic	First-Line Supervisors of Housekeeping and Janitorial Workers
	5	First-Line Supervisors of Housekeeping and Janitorial Workers	Licensed Practical and Licensed Vocational Nurses
Q4	1	Sales Managers	Sales Managers
	2	General and Operations Managers	Medical and Health Services Managers
	3	First-Line Supervisors of Production and Operating Workers	General and Operations Managers
	4	Industrial Production Managers	First-Line Supervisors of Office and Administrative Support Workers
	5	Marketing Managers	Marketing Managers
Q5	1	Chief Executives	Customer Service Representatives
	2	Customer Service Representatives	Chief Executives
	3	Managers, All Other	Accountants
	4	Software Developers, Systems Software	Managers, All Other
	5	Computer and Information Systems Managers	Human Resources Specialists

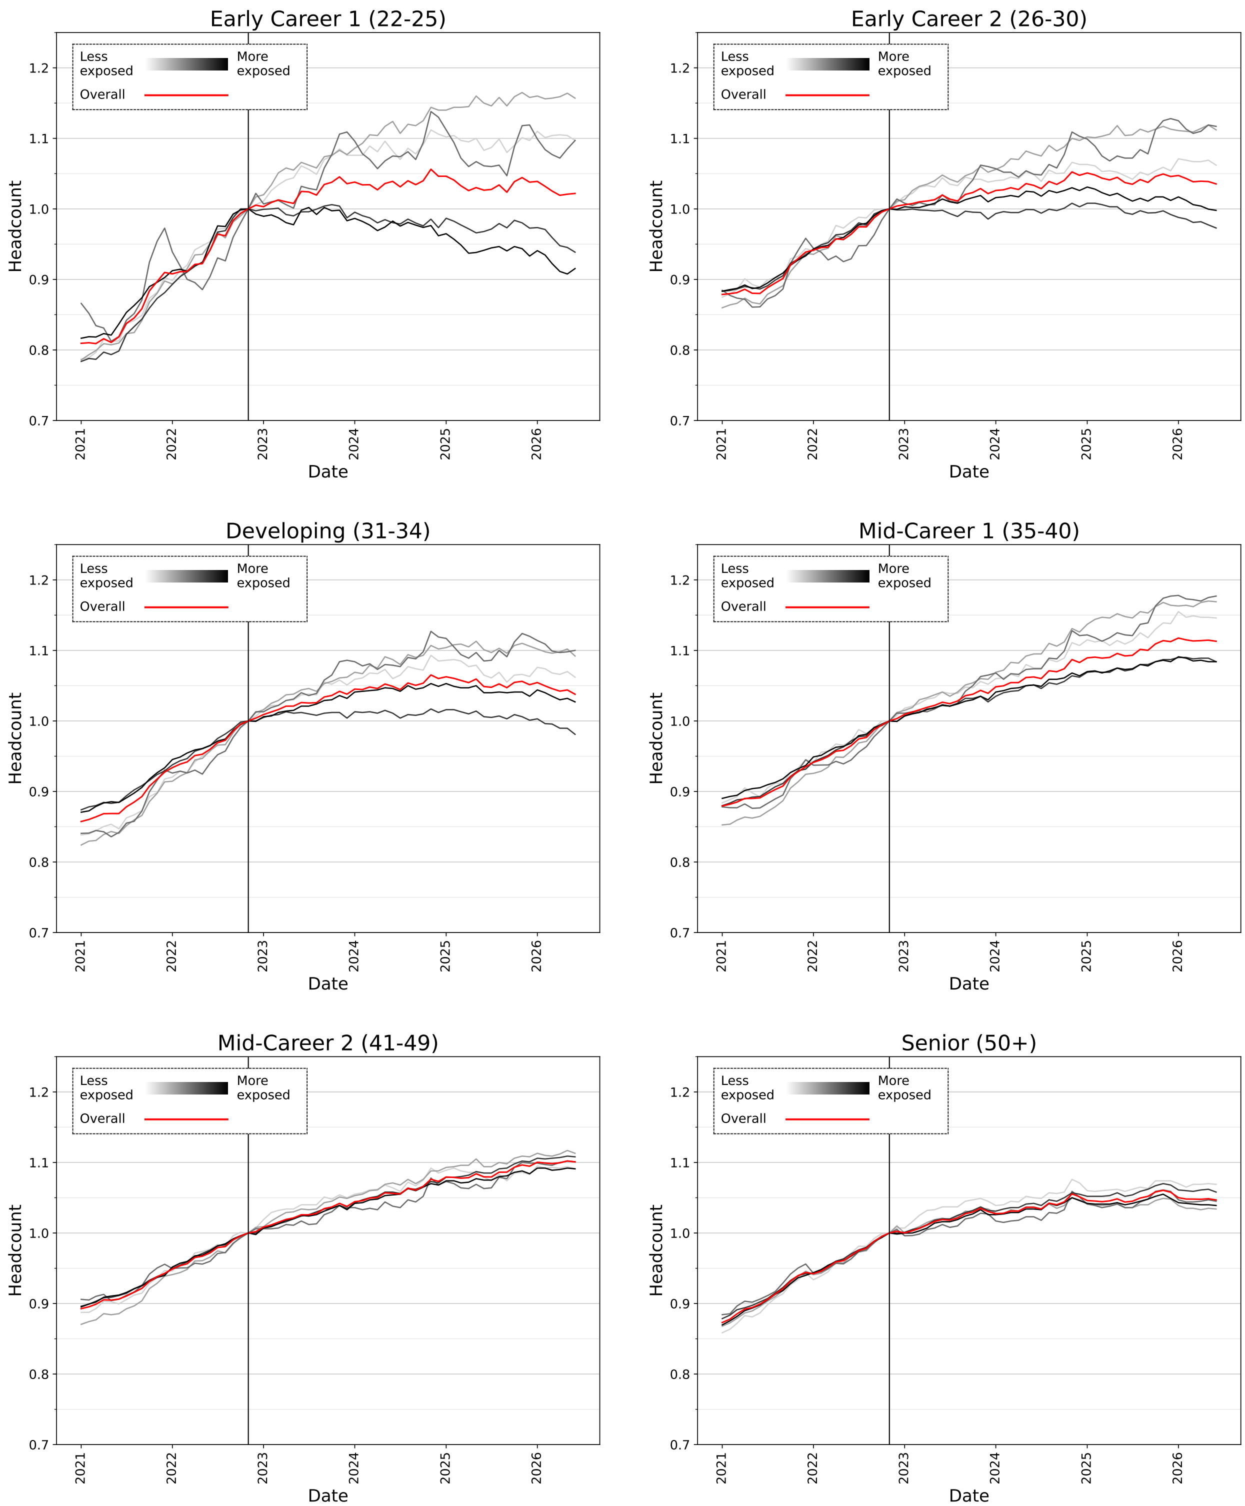

Figure D.1: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only men.

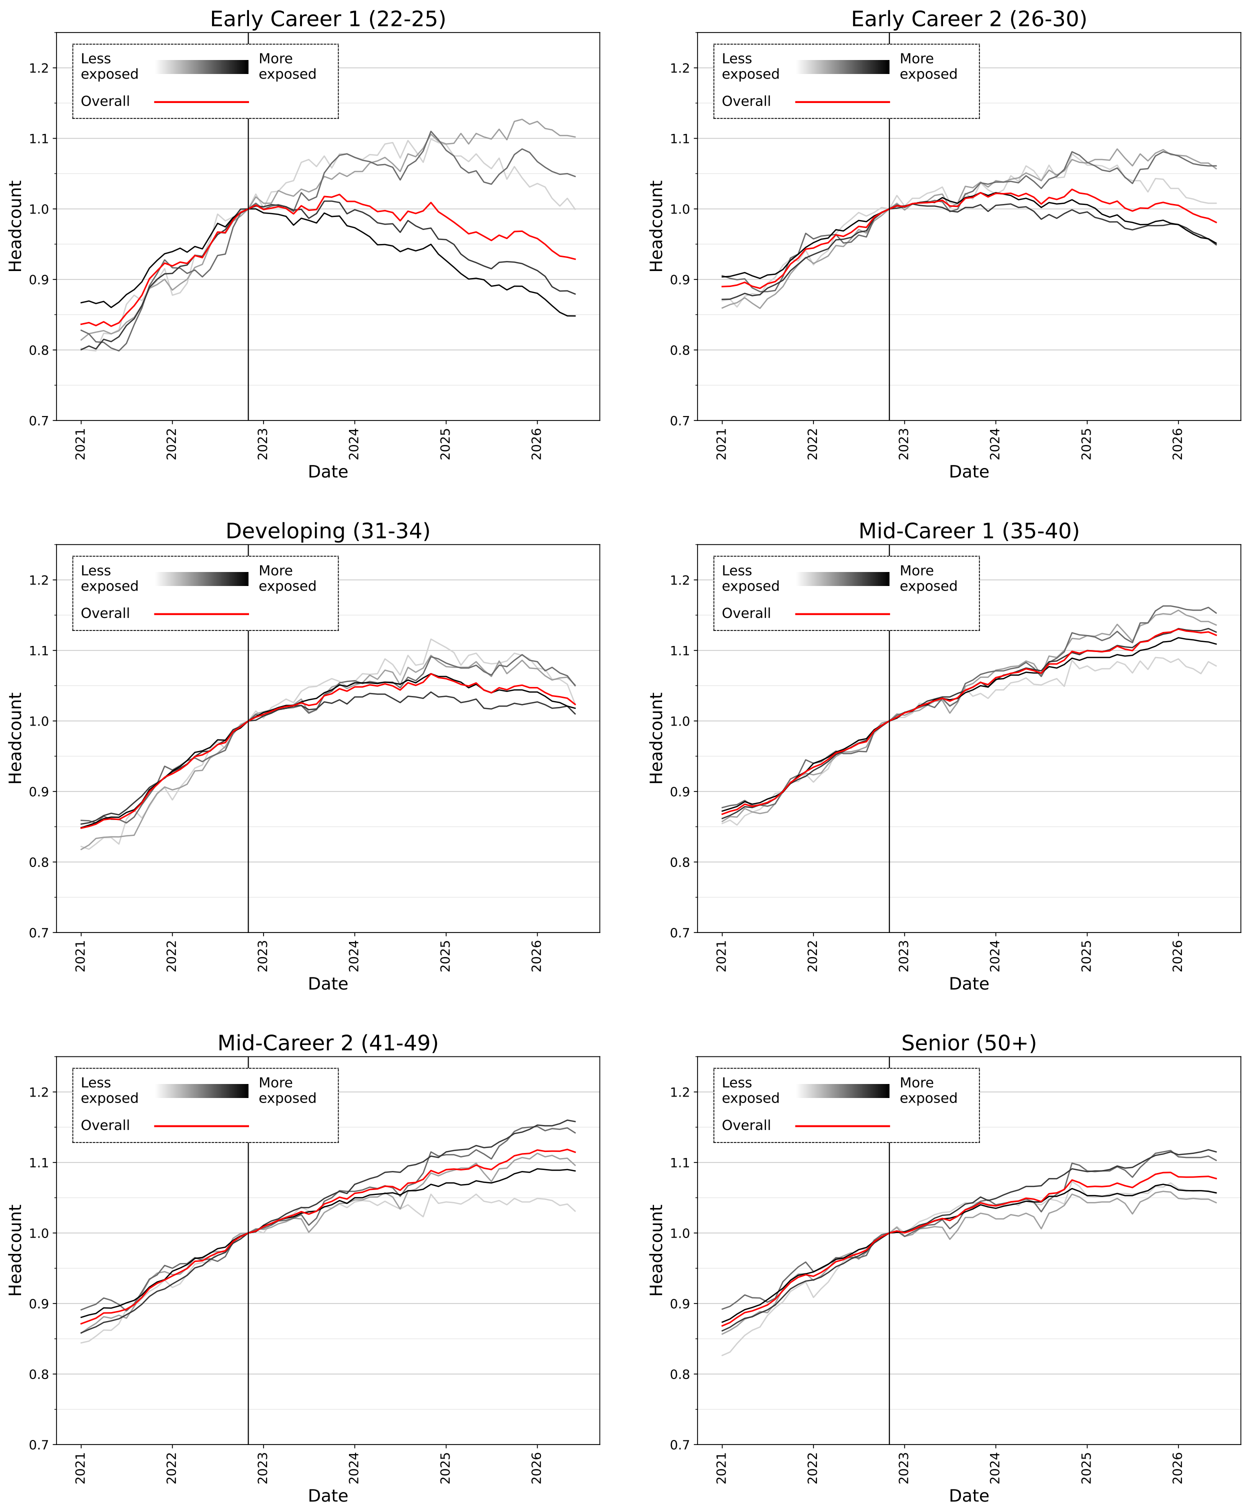

Figure D.2: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only women.

D.2 Geographic Variation: Census Regions

Employment changes by age and GPT-4 β exposure quintile, computed separately within each Census region. The relative decline for the most exposed young workers is greatest in the West and Northeast and smallest in the South and Midwest. Figure D.3 summarizes the endline gap between the most and least exposed quintiles (Q5 $-$ Q1, with no controls) for workers aged 22–25 across regions; Figures D.4–D.7 show the underlying age-by-exposure series. Table D.3 reports the corresponding occupation-level long-difference regression estimates by region, which reproduce the ordering in Figure D.3 closely. As with the other subsample regressions, these are estimated on the panel of firms balanced since January 2021, for which the full-sample estimate for the most exposed quintile is -0.193 (Table 2).

Table D.3: Occupation-level long-difference estimates for workers aged 22–25 by Census region. Heteroskedasticity-robust standard errors in parentheses.

Exposure quintile	West	Midwest	Northeast	South
Quintile 2	0.083 (0.052)	0.115 (0.071)	0.005 (0.067)	0.044 (0.064)
Quintile 3	0.072 (0.056)	0.027 (0.061)	-0.129^{**} (0.062)	0.000 (0.051)
Quintile 4	-0.195^{***} (0.049)	-0.087 (0.058)	-0.236^{***} (0.057)	-0.157^{***} (0.055)
Quintile 5	-0.256^{***} (0.042)	-0.139^{***} (0.053)	-0.254^{***} (0.054)	-0.157^{***} (0.047)
Occupations	546	523	530	569

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. Coefficients relative to quintile 1.

D.3 Geographic Variation: AI Adoption Tiers

Employment changes by age and GPT-4 β exposure quintile, computed separately for states grouped into tiers of AI relative usage from the Anthropic Economic Index (Appel et al., 2025).

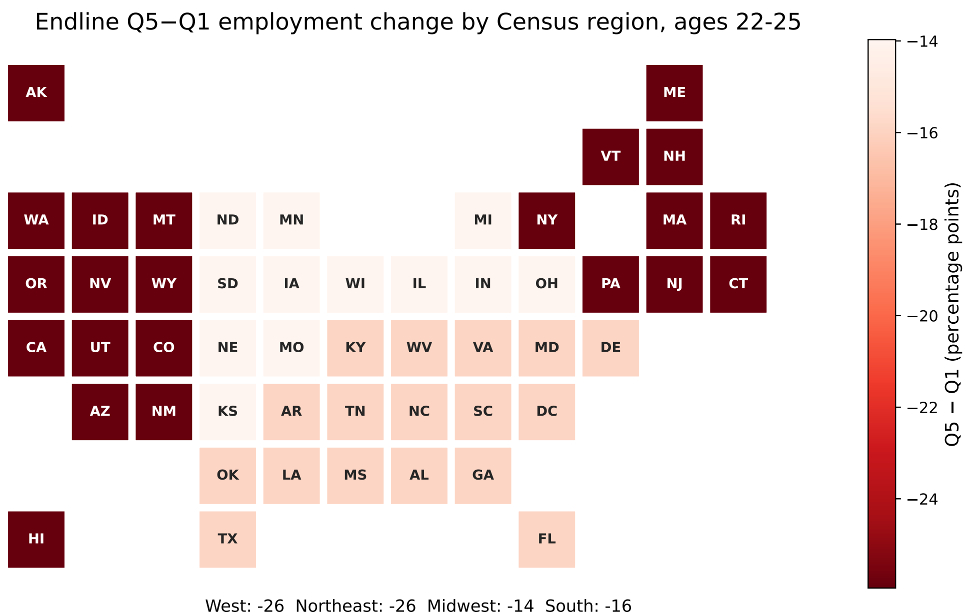

Figure D.3: Endline (June 2026) difference in normalized employment between the most exposed (Q5) and least exposed (Q1) occupation quintiles for workers aged 22–25, by Census region, with no controls. Darker shades indicate a larger relative decline in exposed employment. States are shown in an approximate tile-grid layout and colored by their Census region's value.

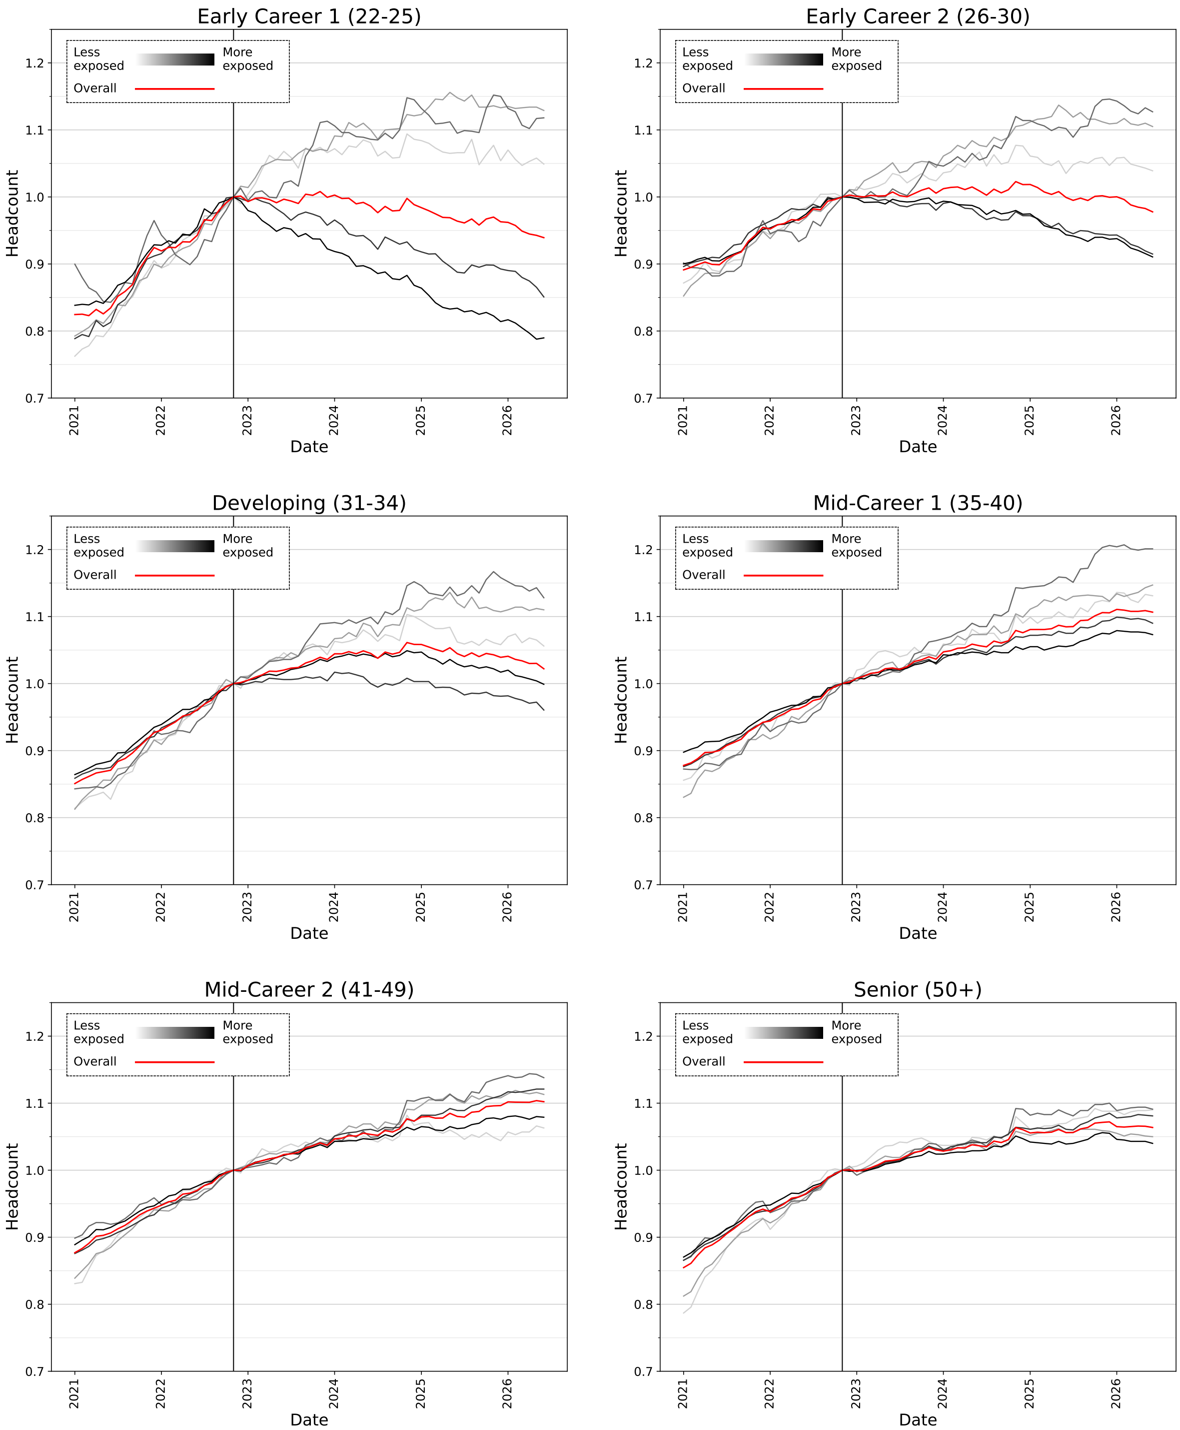

Figure D.4: West region: employment changes by age and AI exposure quintile.

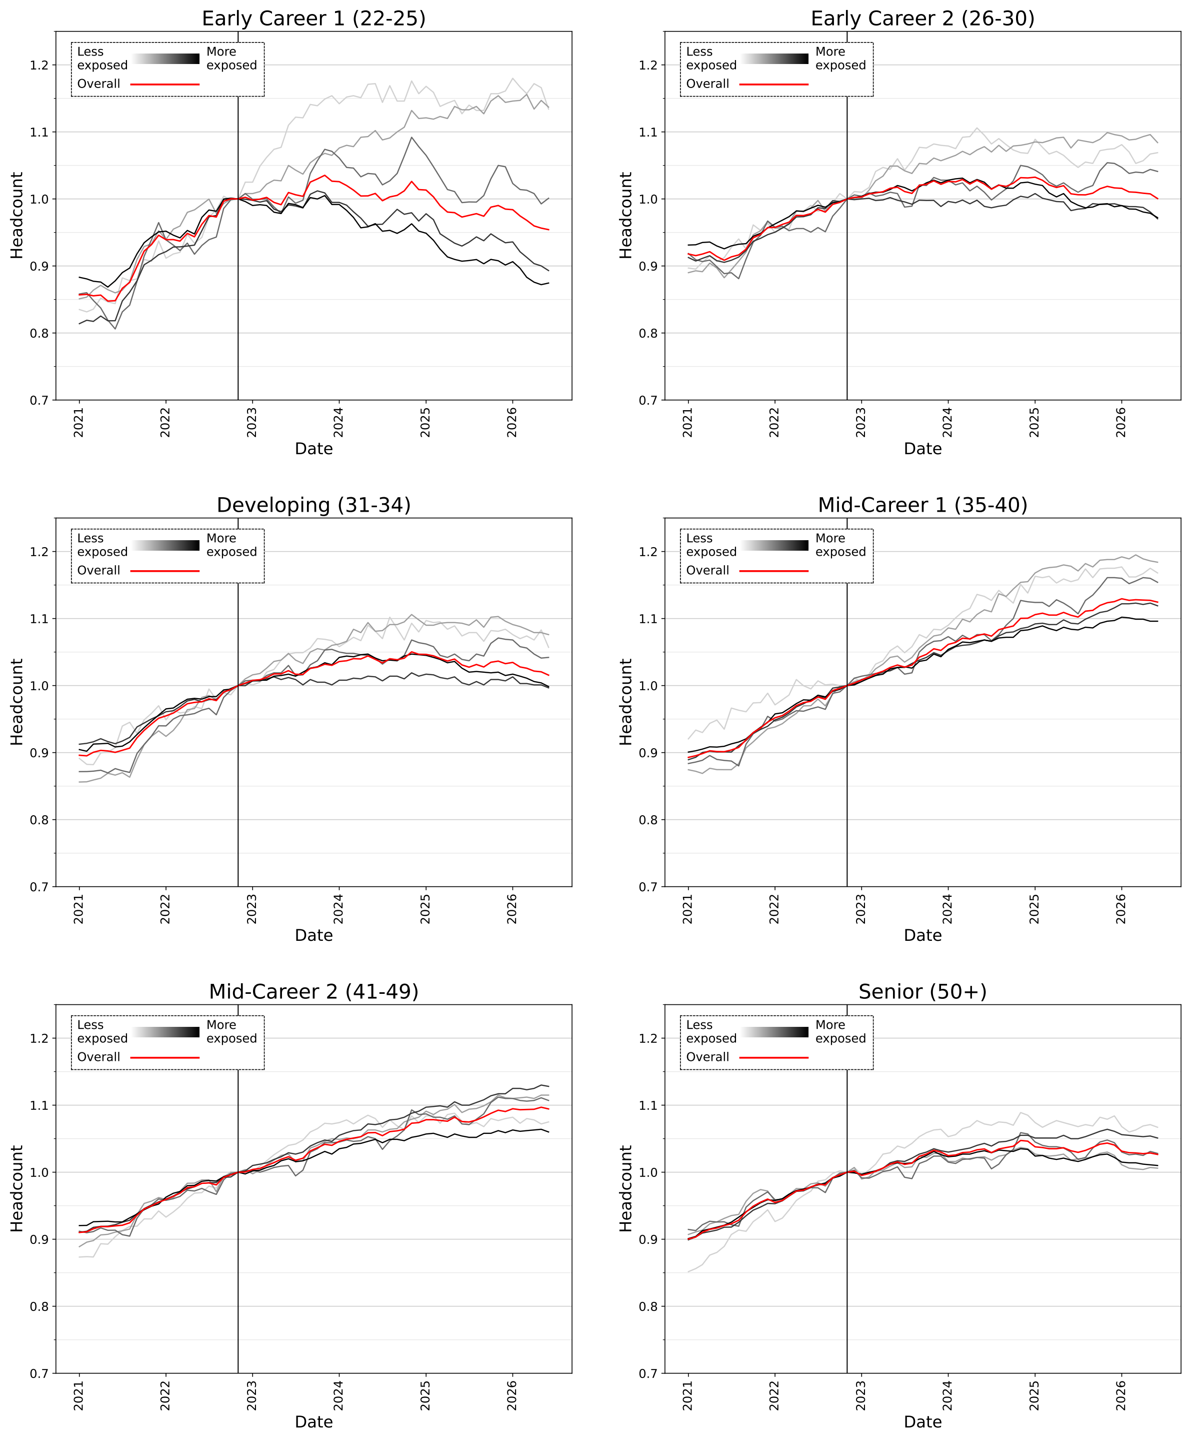

Figure D.5: Northeast region: employment changes by age and AI exposure quintile.

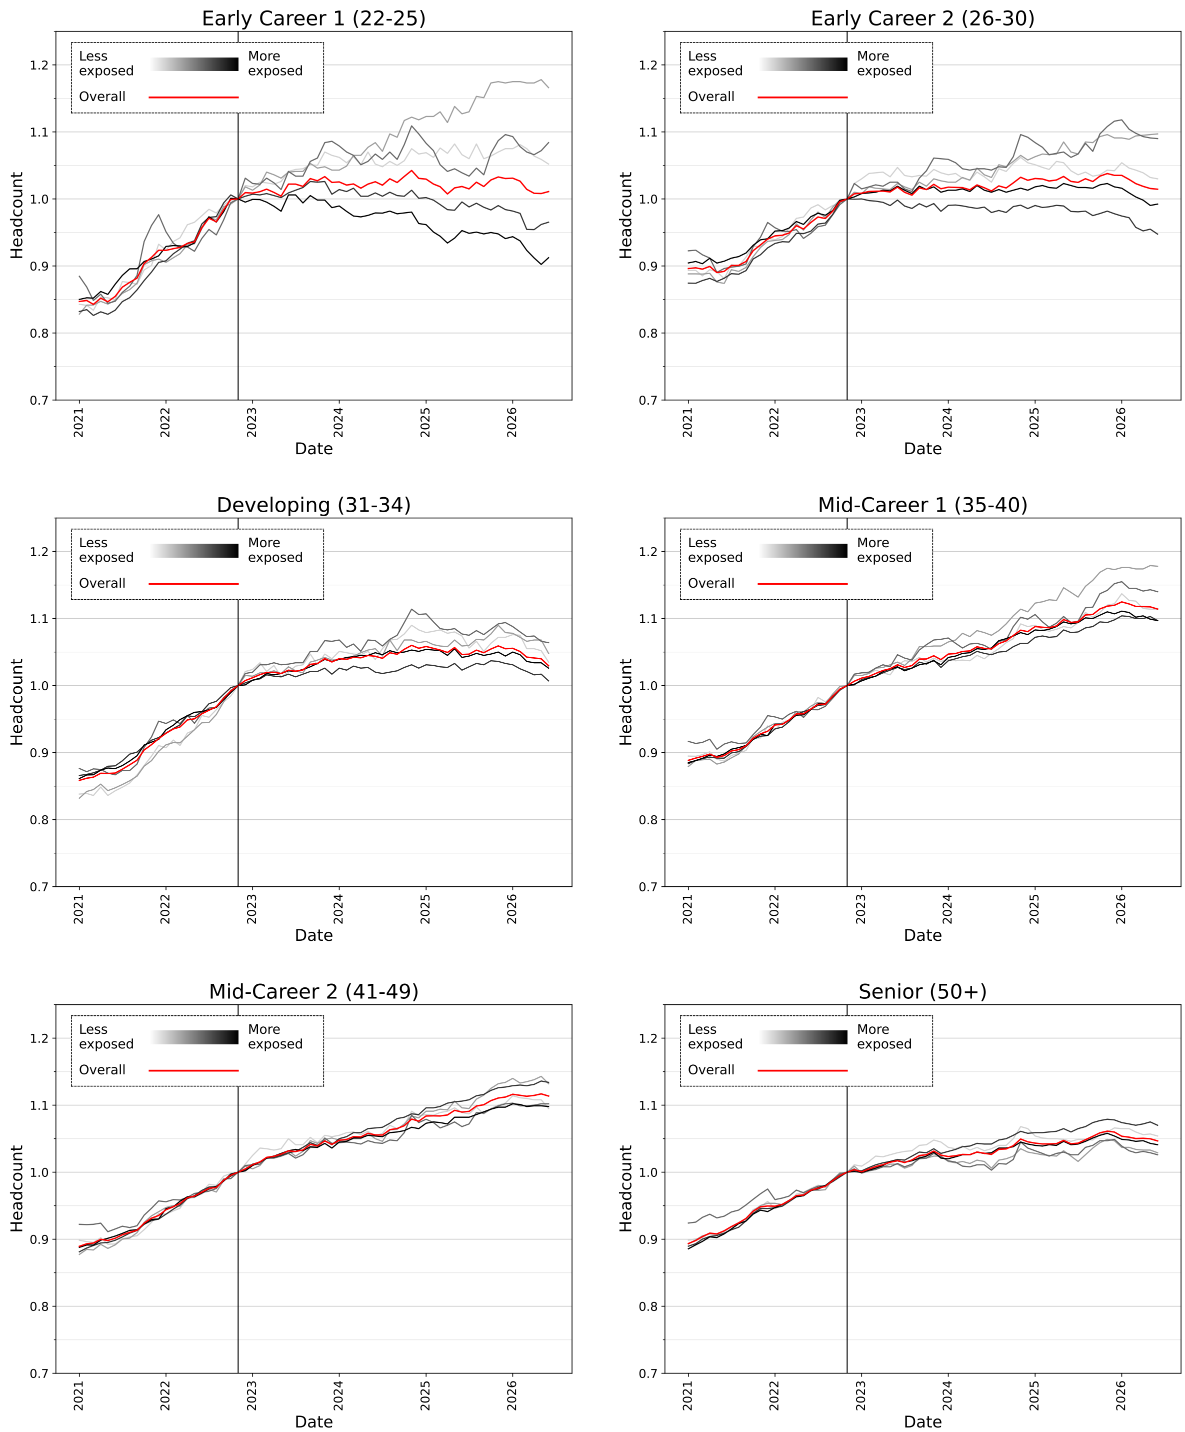

Figure D.6: Midwest region: employment changes by age and AI exposure quintile.

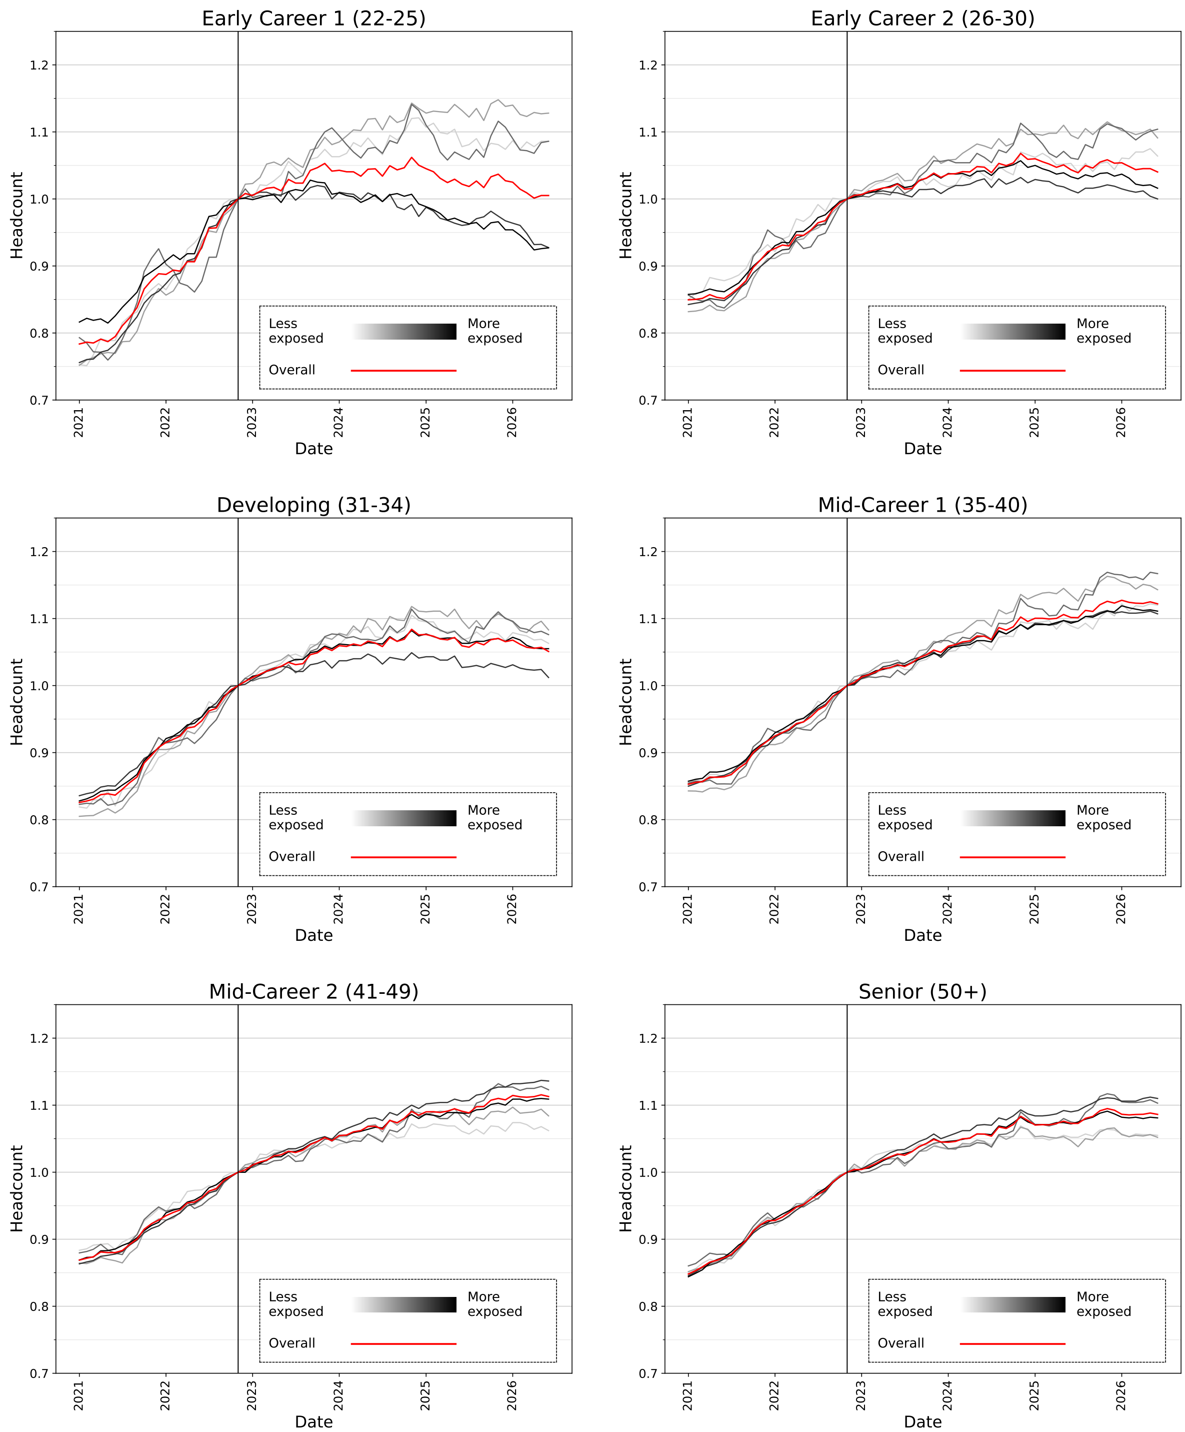

Figure D.7: South region: employment changes by age and AI exposure quintile.

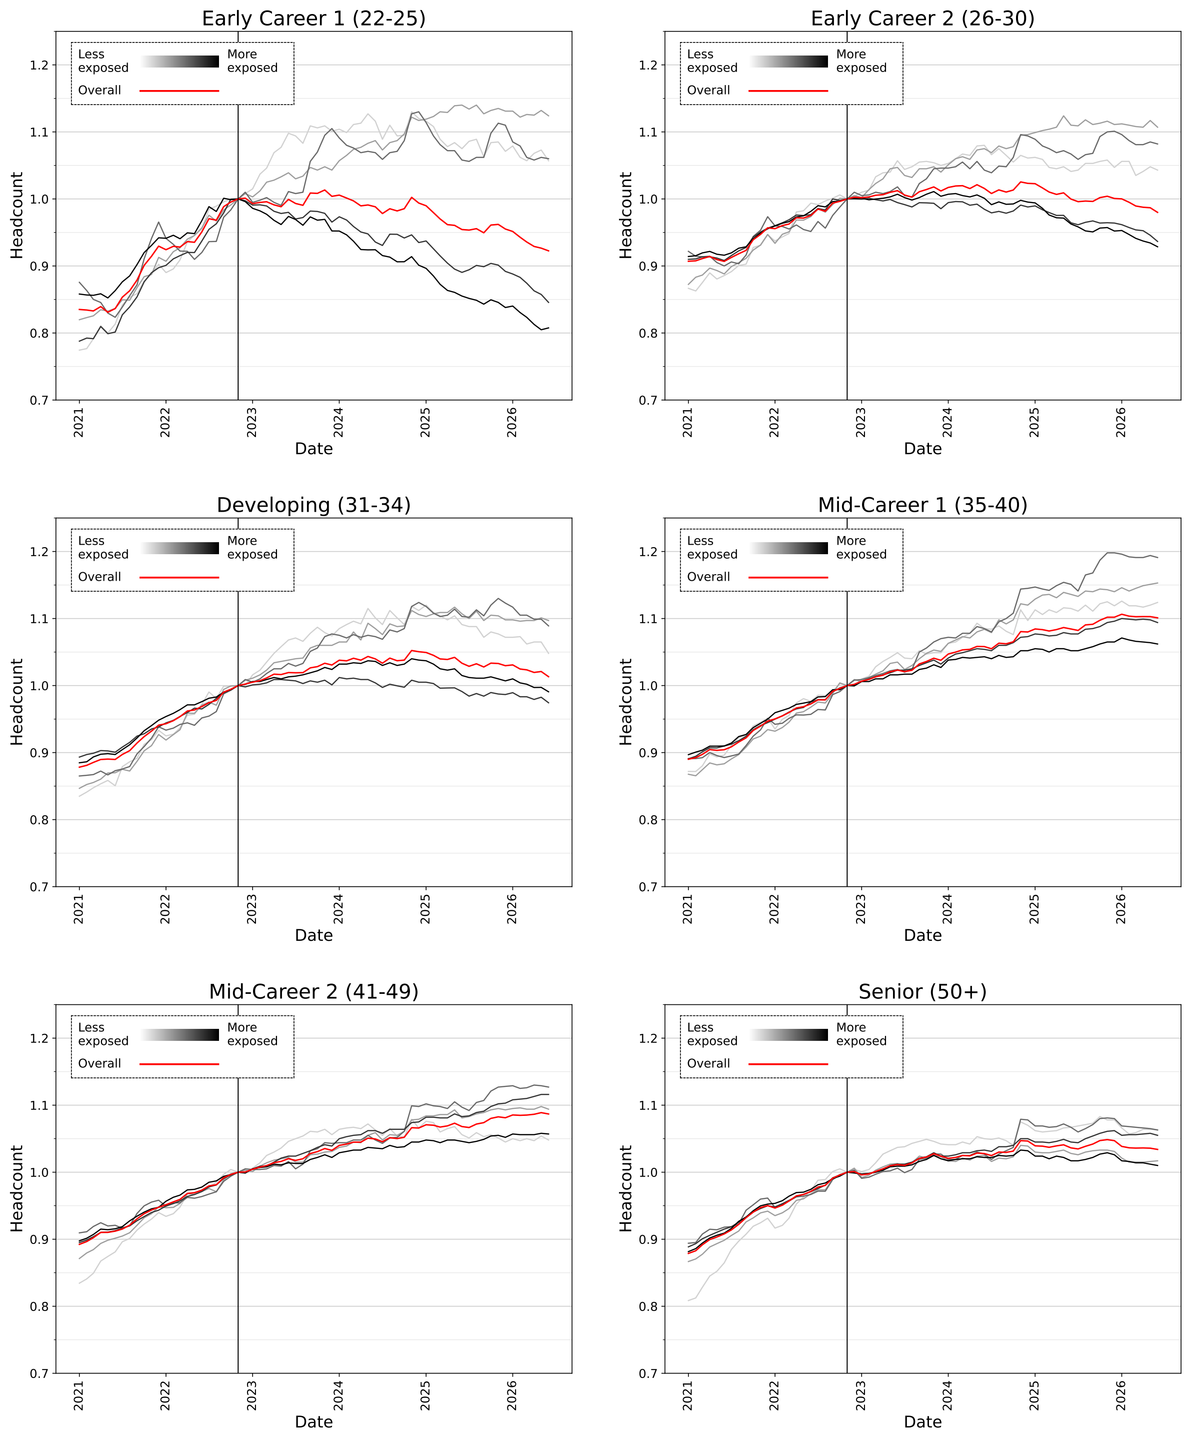

Figure D.8: Employment changes by age and AI exposure for leading AI-adoption states.

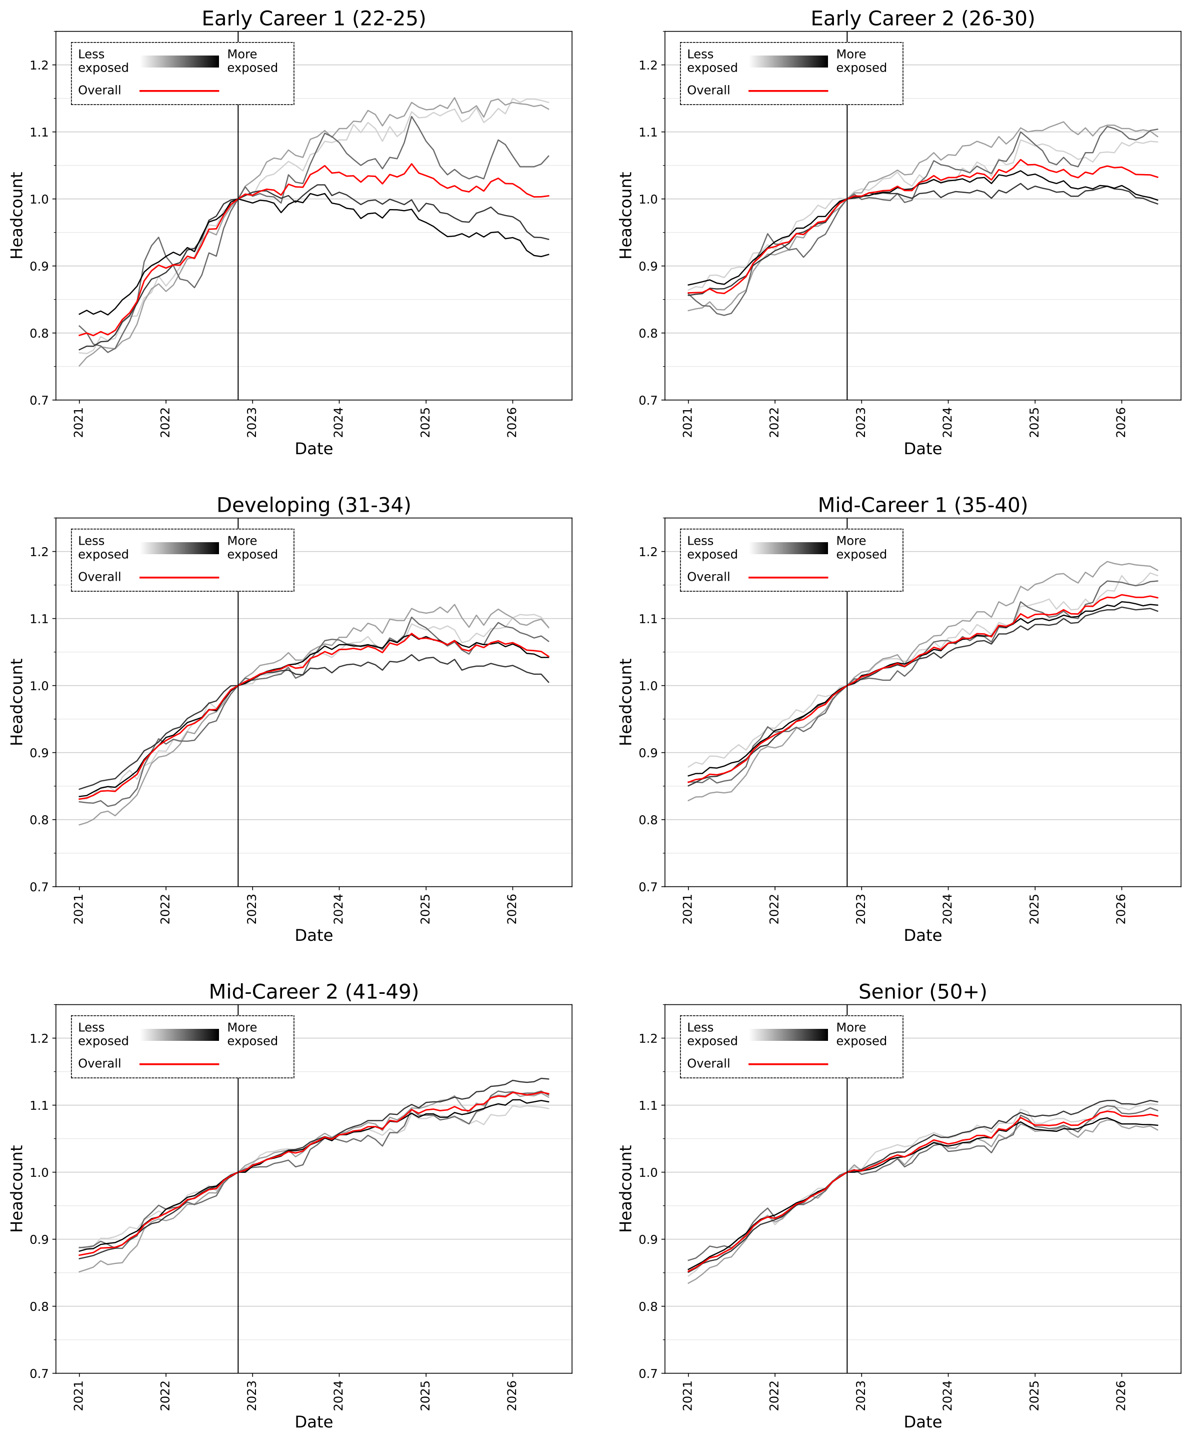

Figure D.9: Employment changes by age and AI exposure for upper-middle AI-adoption states.

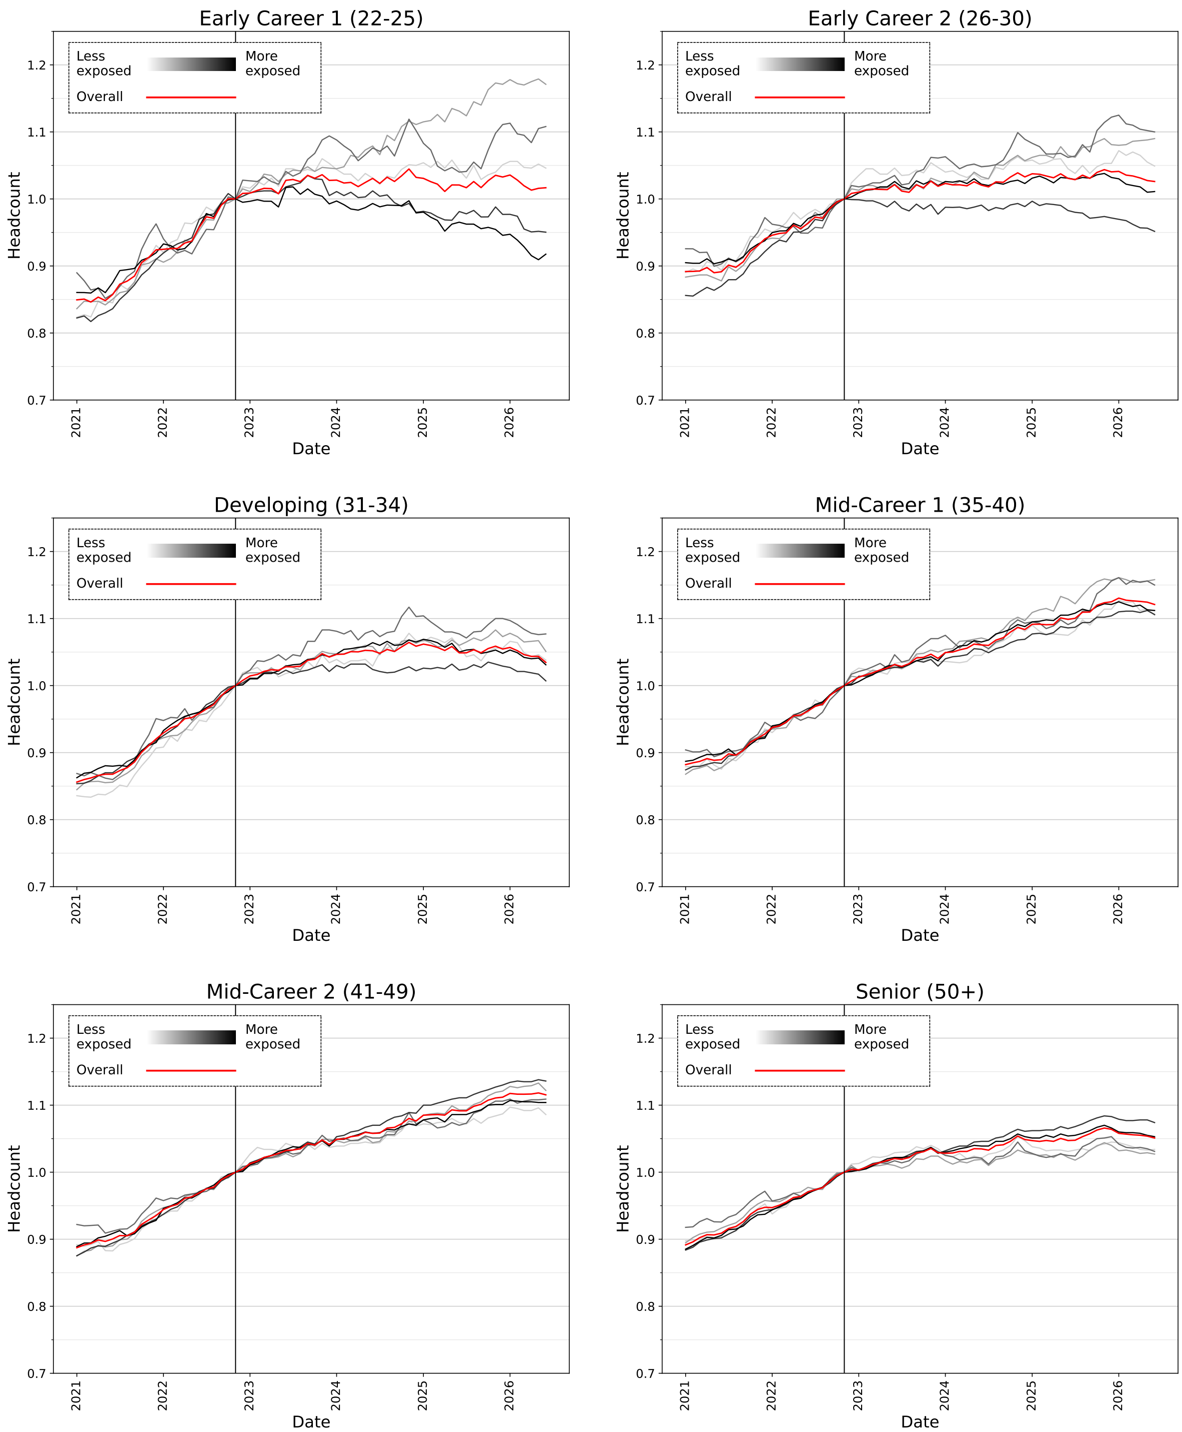

Figure D.10: Employment changes by age and AI exposure for lower-middle AI-adoption states.

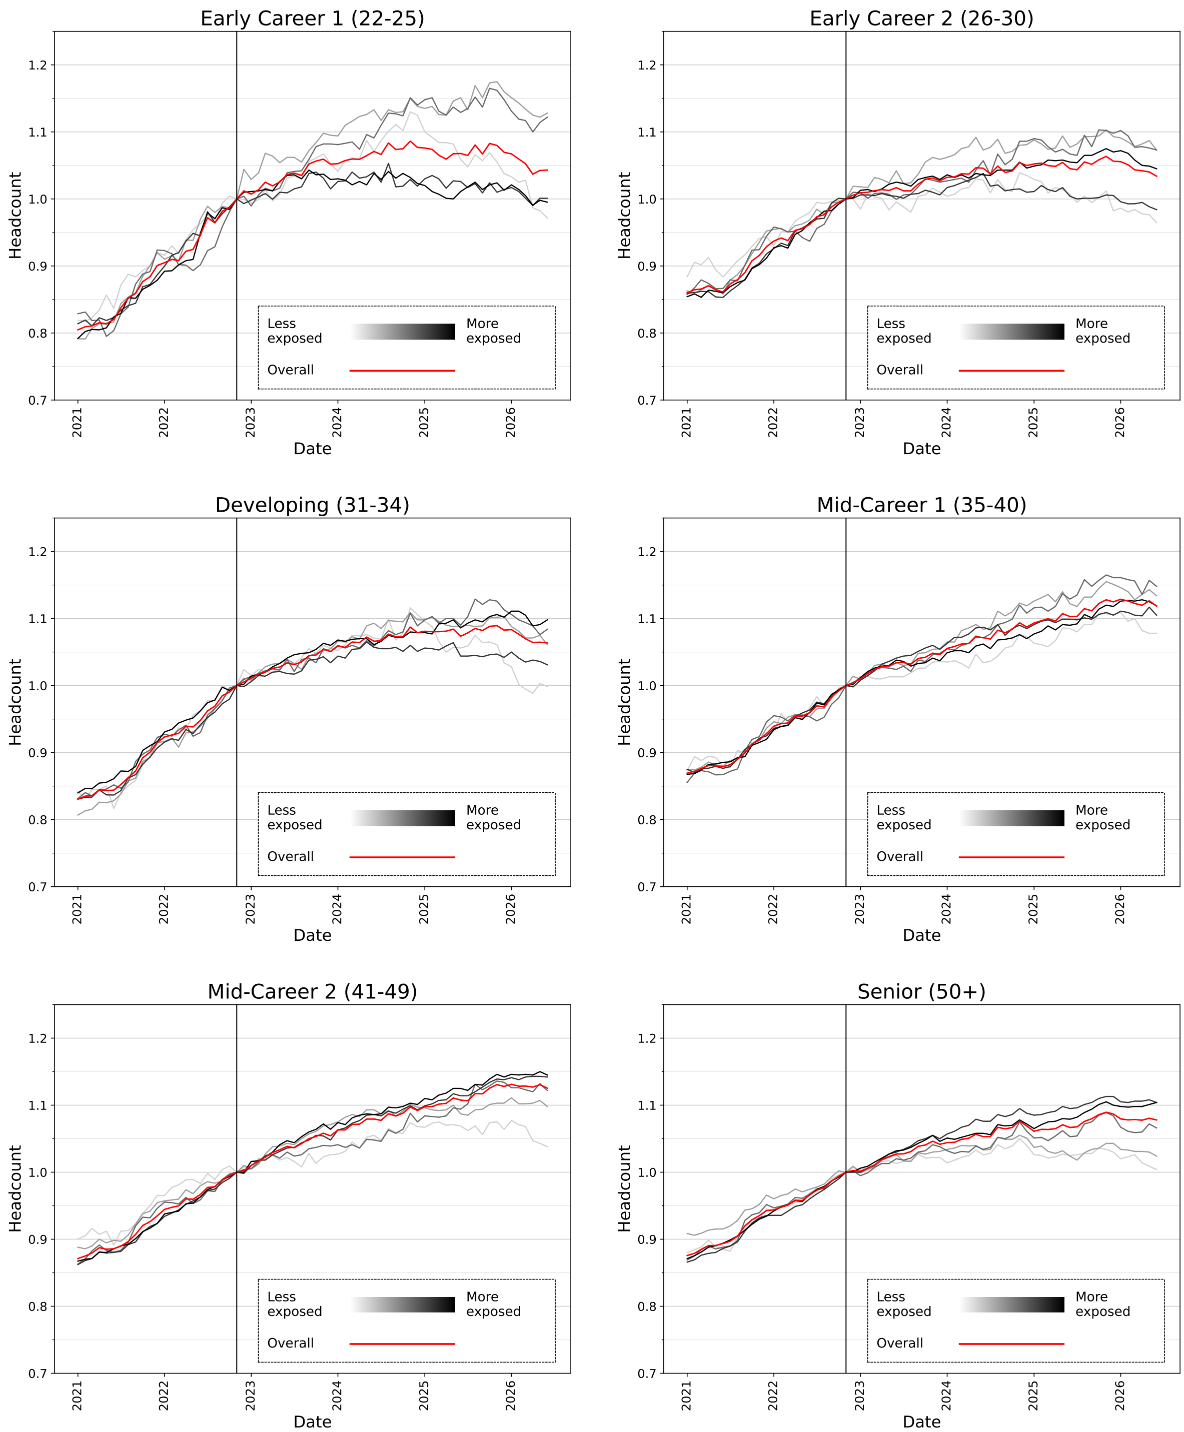

Figure D.11: Employment changes by age and AI exposure for emerging AI-adoption states.

E Alternative AI Exposure Measures

Employment changes by age and exposure quintile using alternative occupational AI-exposure measures in place of the GPT-4 β measure.

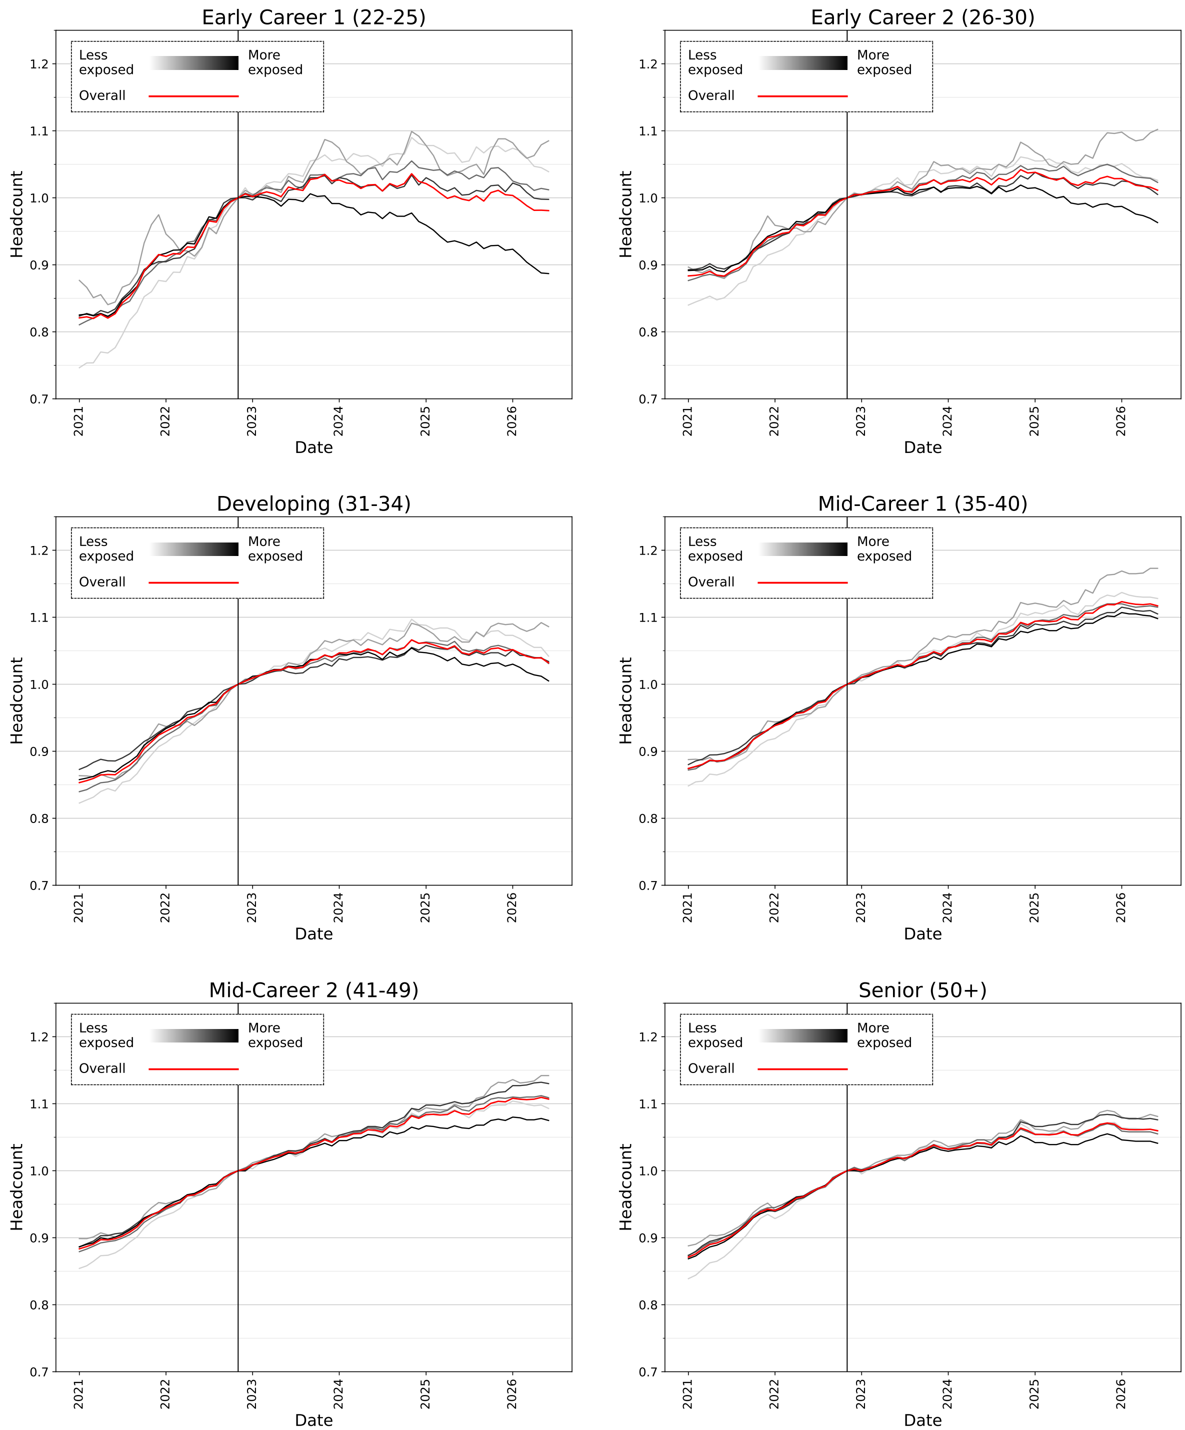

Figure E.1: Employment changes by age and Suitability-for-machine-learning (SML) score from Brynjolfsson et al. (2018).

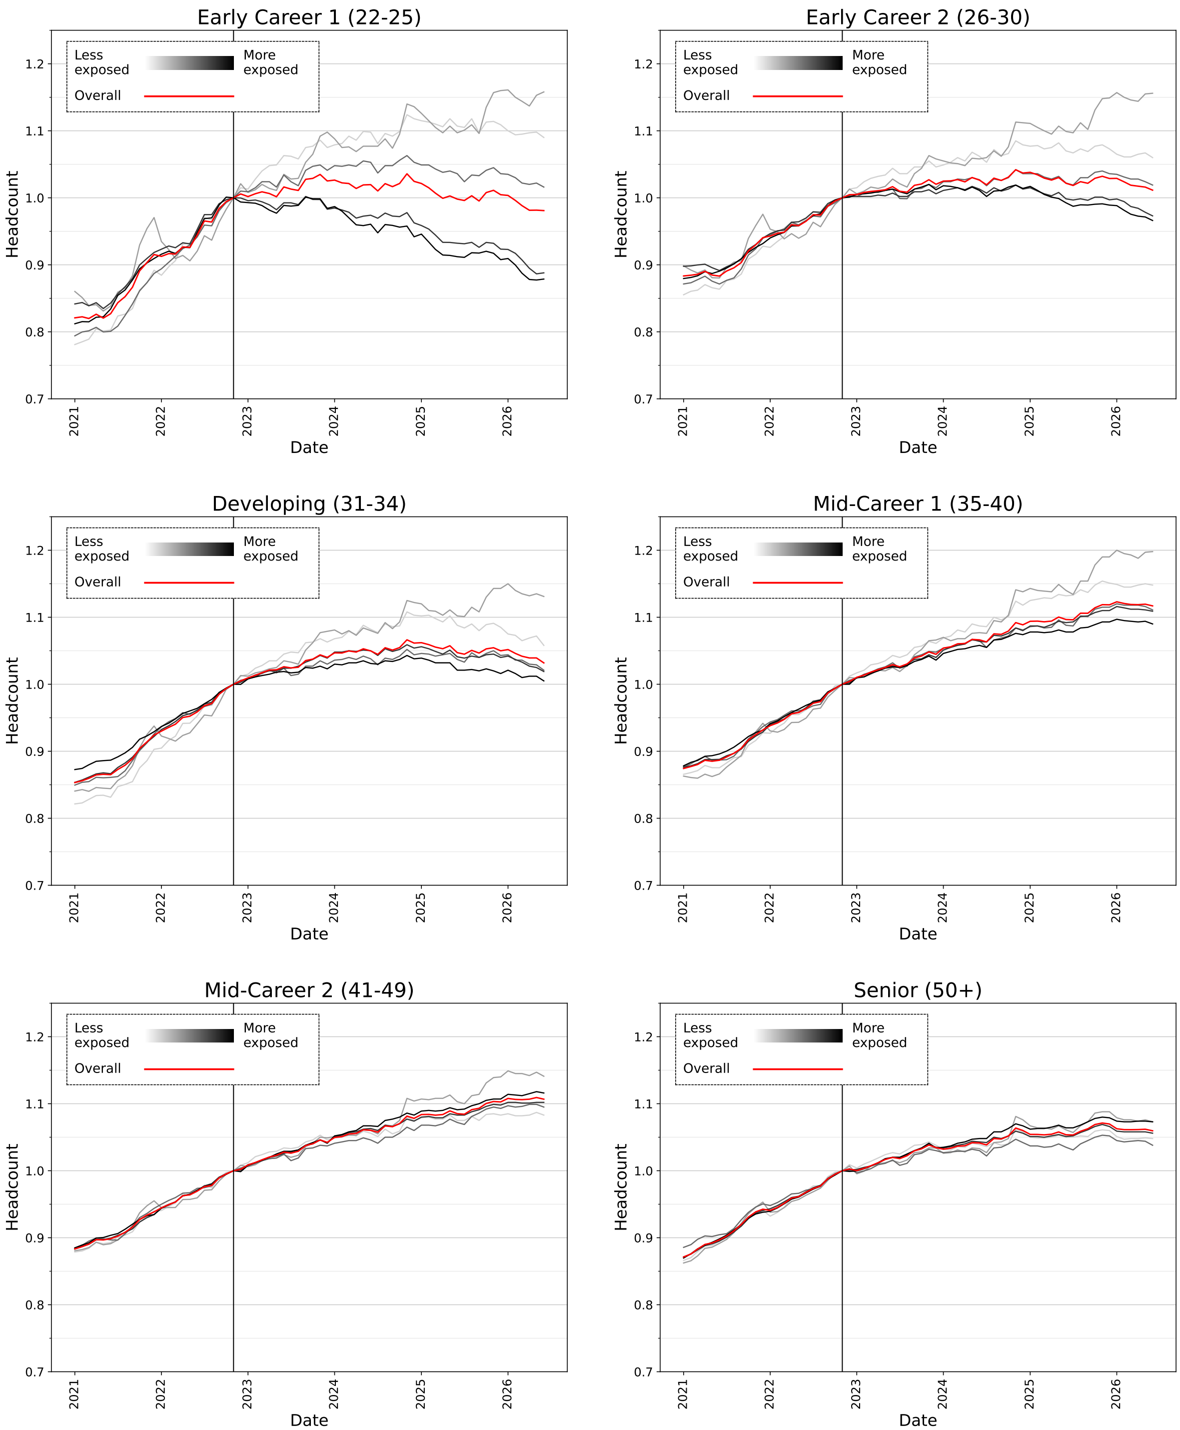

Figure E.2: Employment changes by age and AI Occupational Exposure (AIOE) measure from Felten et al. (2021).

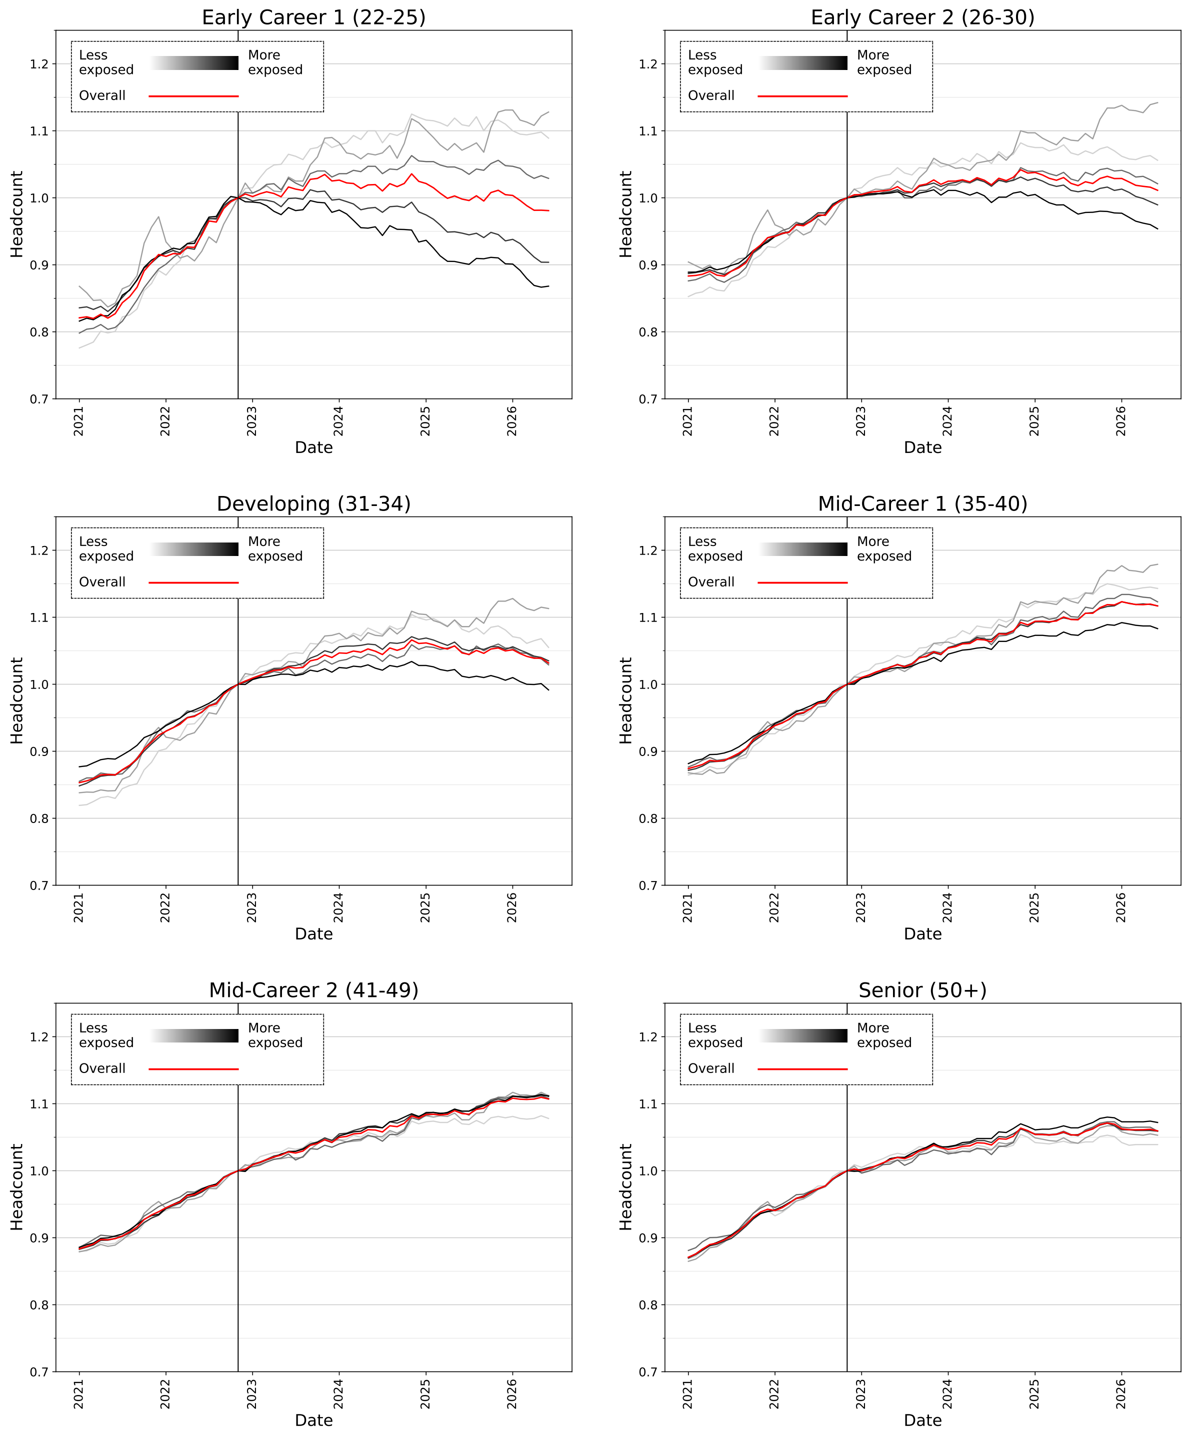

Figure E.3: Employment changes by age and language-modeling AIOE measure from Felten et al. (2023).

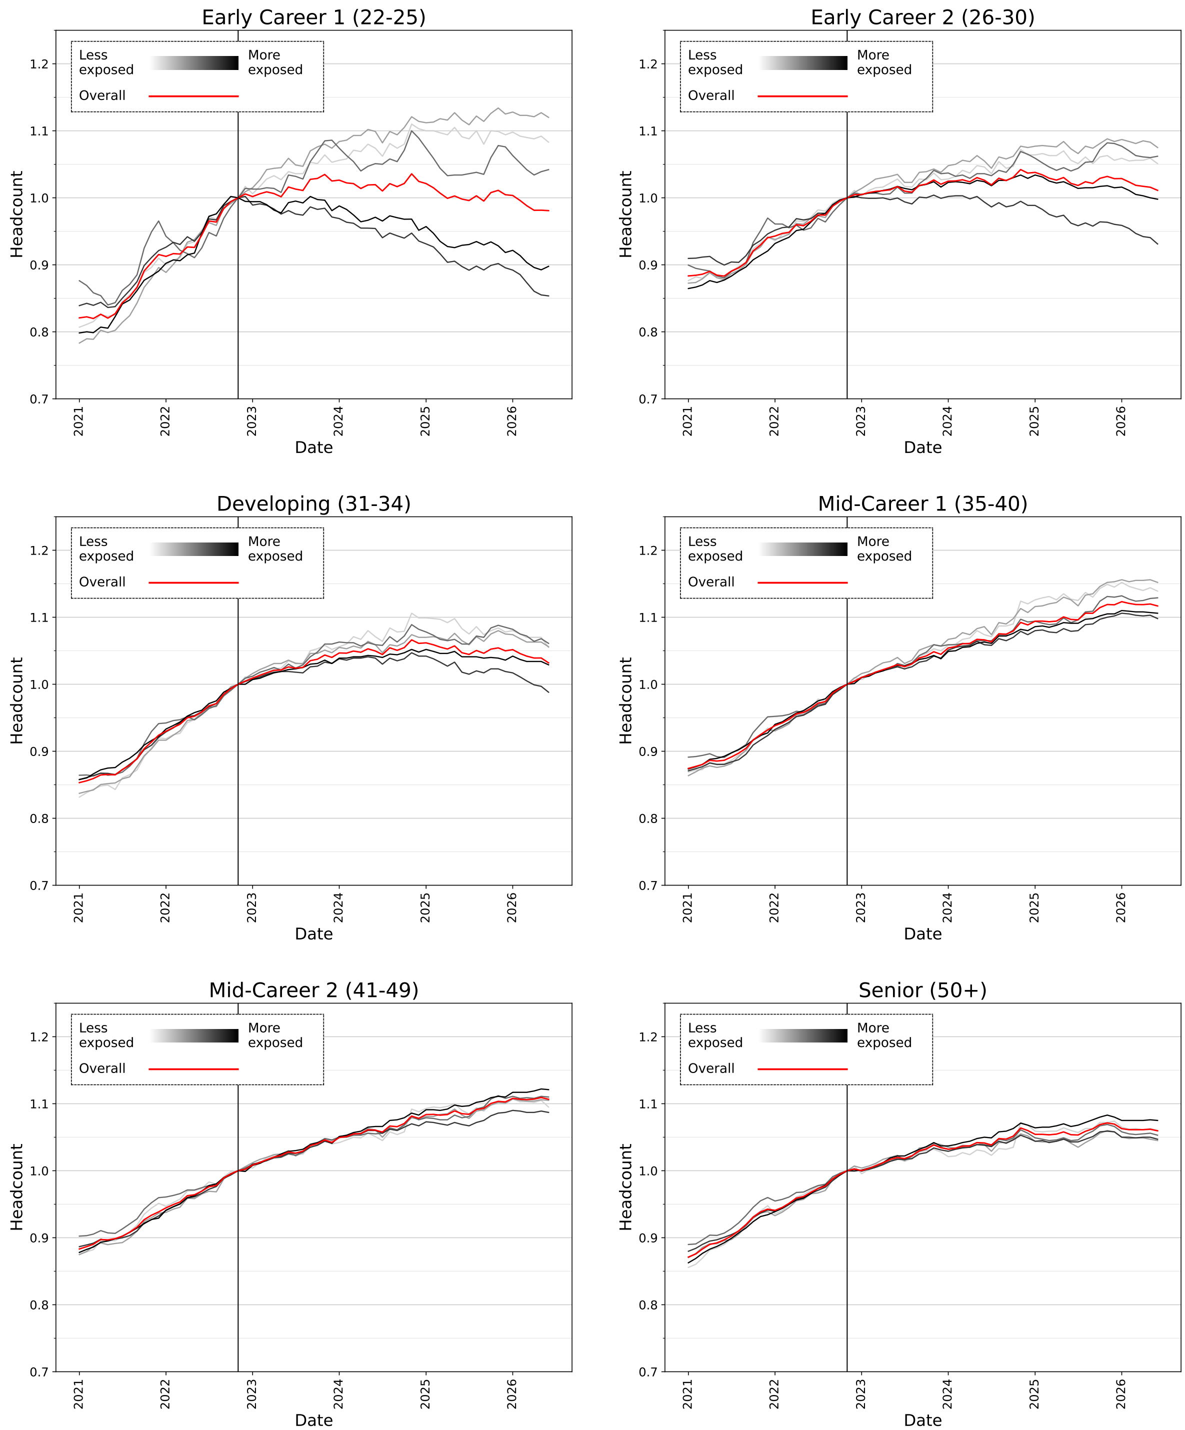

Figure E.4: Employment changes by age and image-generation AIOE measure from Felten et al. (2023).

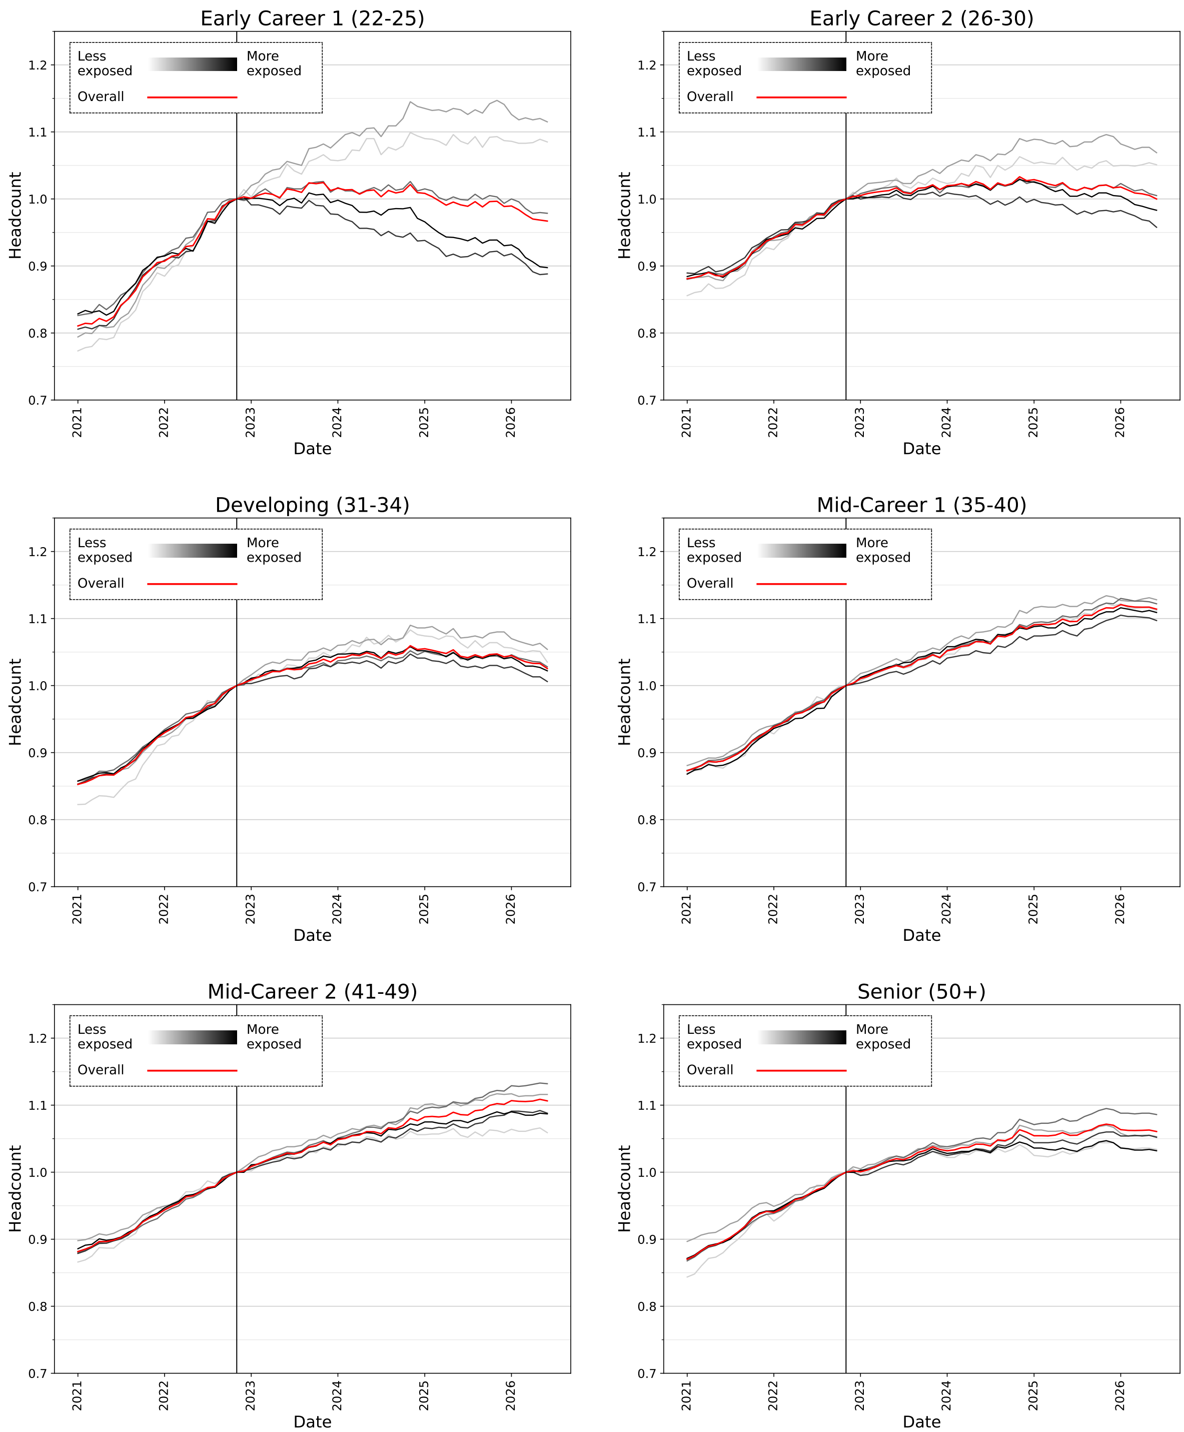

Figure E.5: Employment changes by age and AI applicability measure from Tomlinson et al. (2025).

F Automation versus Complementarity

This section reports employment changes by generative-AI usage exposure from the Anthropic Economic Index, distinguishing automative (substitutive) from “augmentative” (complementary) usage.

F.1 March 2025 AEI Release

Employment changes by age and exposure quintile, measured by overall Claude usage (Fig. F.1), by automation share (Fig. F.2), and by complementarity share (Fig. F.3). Figures F.4 and F.5 repeat the automation and complementarity results after dropping occupations with low overall Claude usage.

F.2 More Recent AEI Releases

Employment changes by age and exposure quintile using the more recent releases of the Anthropic Economic Index, pooling the September 2025, January 2026, and April 2026 data (Appel et al., 2025, 2026; Massenkoff et al., 2026).

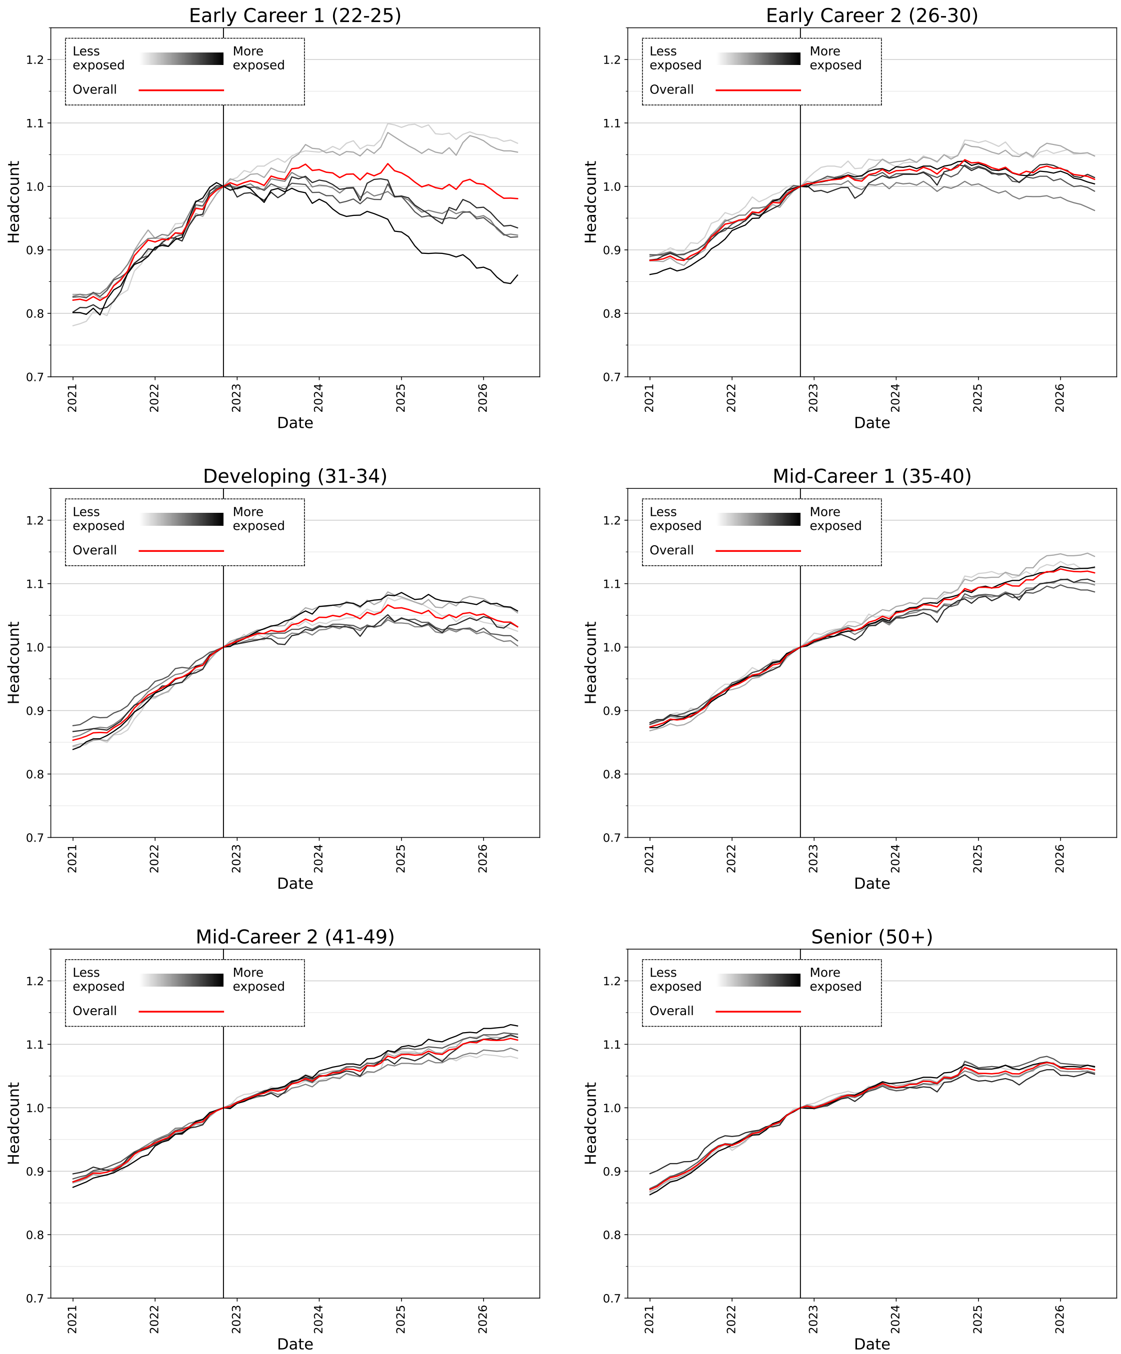

Figure F.1: *Per-capita Claude usage.* Employment changes by age and exposure quintile using Claude usage data from Handa et al. (2025). Exposure quintiles are defined based on the share of queries to Claude that relate to tasks associated with an occupation, scaled by BLS employment in the occupation. Darker lines are more exposed quintiles. The red line shows the overall trend pooling across quintiles. 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.

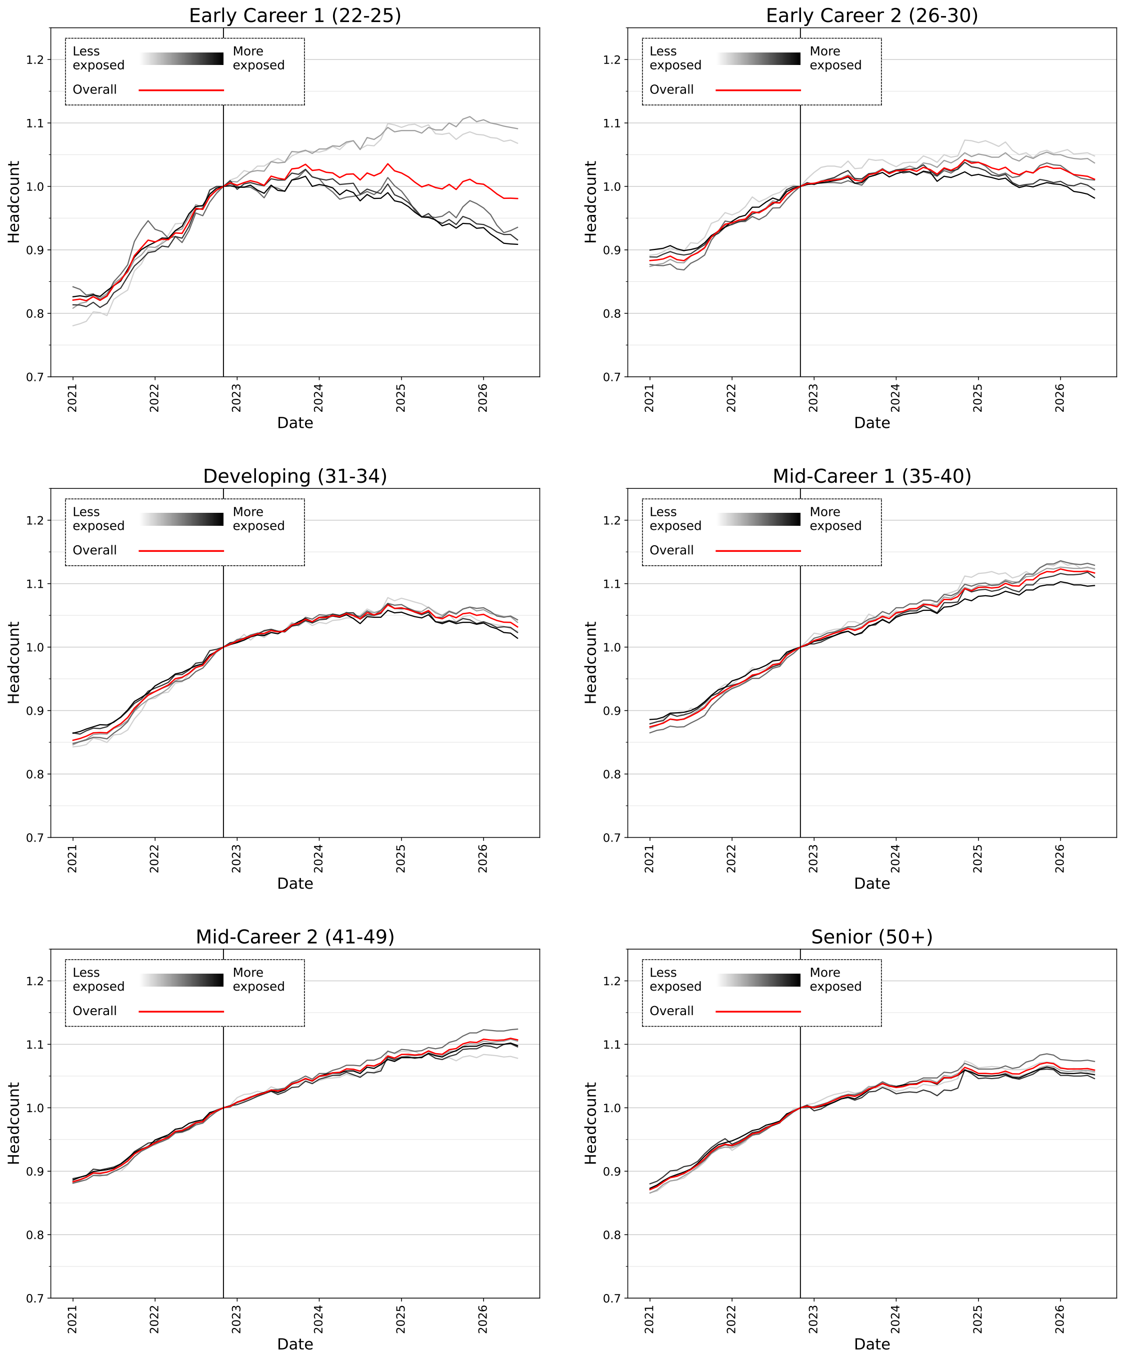

Figure F.2: *Automation.* Employment changes by age and automation quintile using Claude usage data from Handa et al. (2025). Automation quintiles are defined based on the share of queries related to an occupation that are classified by Claude as automative in nature. Darker lines are more automative quintiles. The red line shows the overall trend pooling across quintiles. 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.

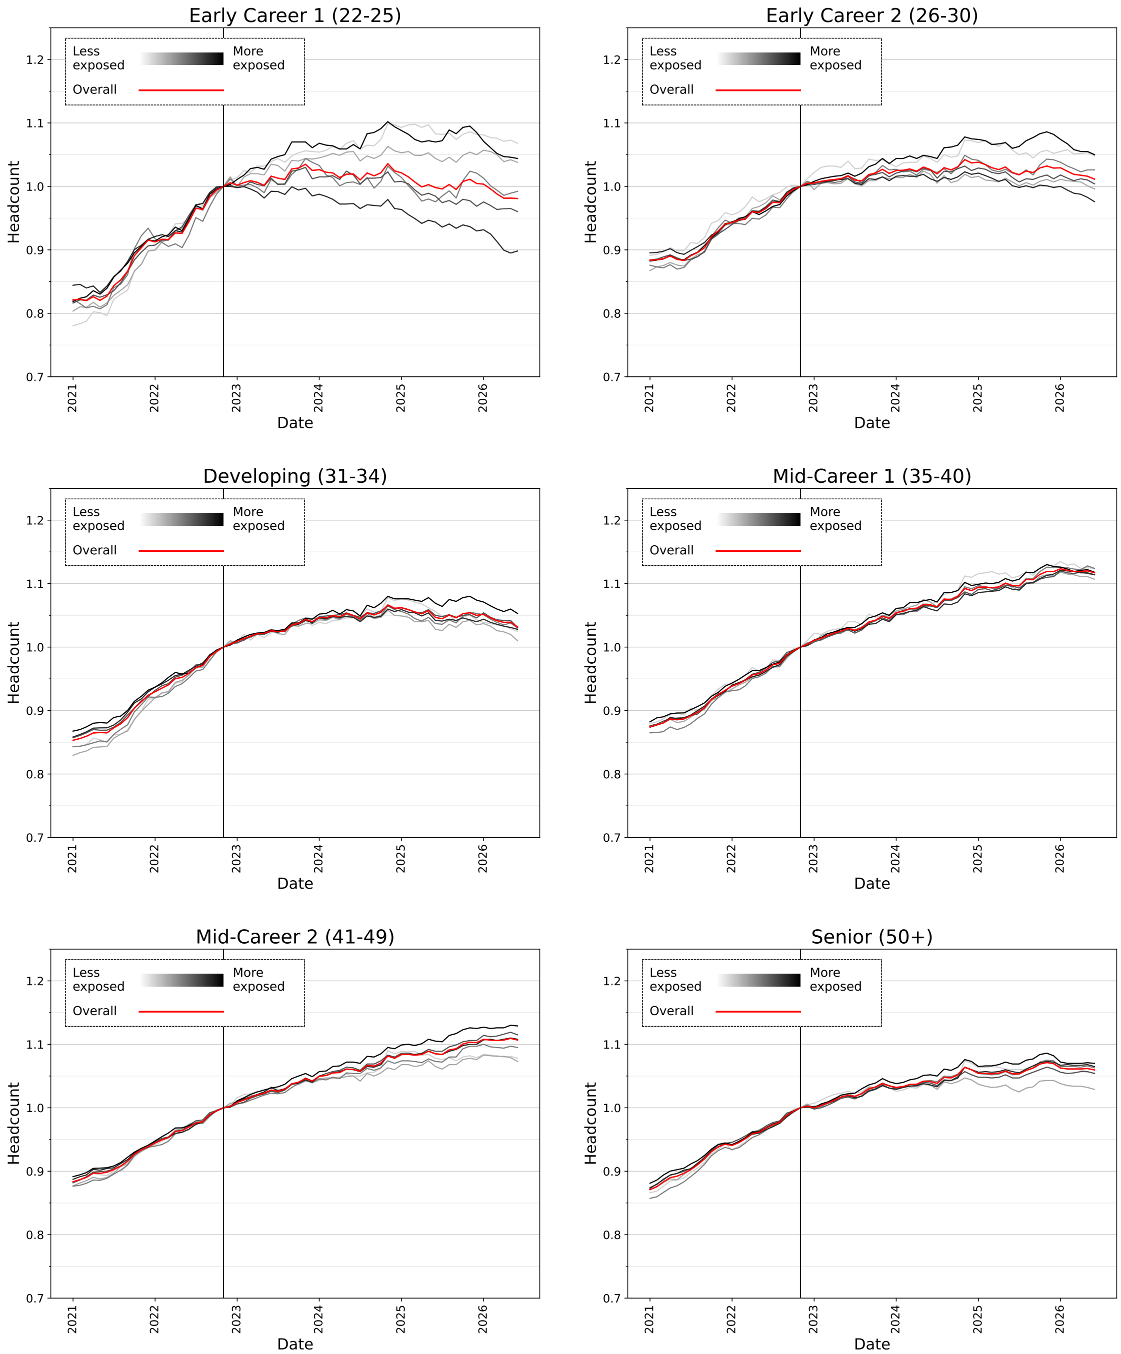

Figure F.3: *Complementarity ("augmentative" usage)*. Employment changes by age and complementarity quintile using Claude usage data from Handa et al. (2025). Augmentation quintiles are defined based on the share of queries related to an occupation that are classified by Claude as augmentative. Darker lines are more augmentative quintiles. The red line shows the overall trend pooling across quintiles. 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.

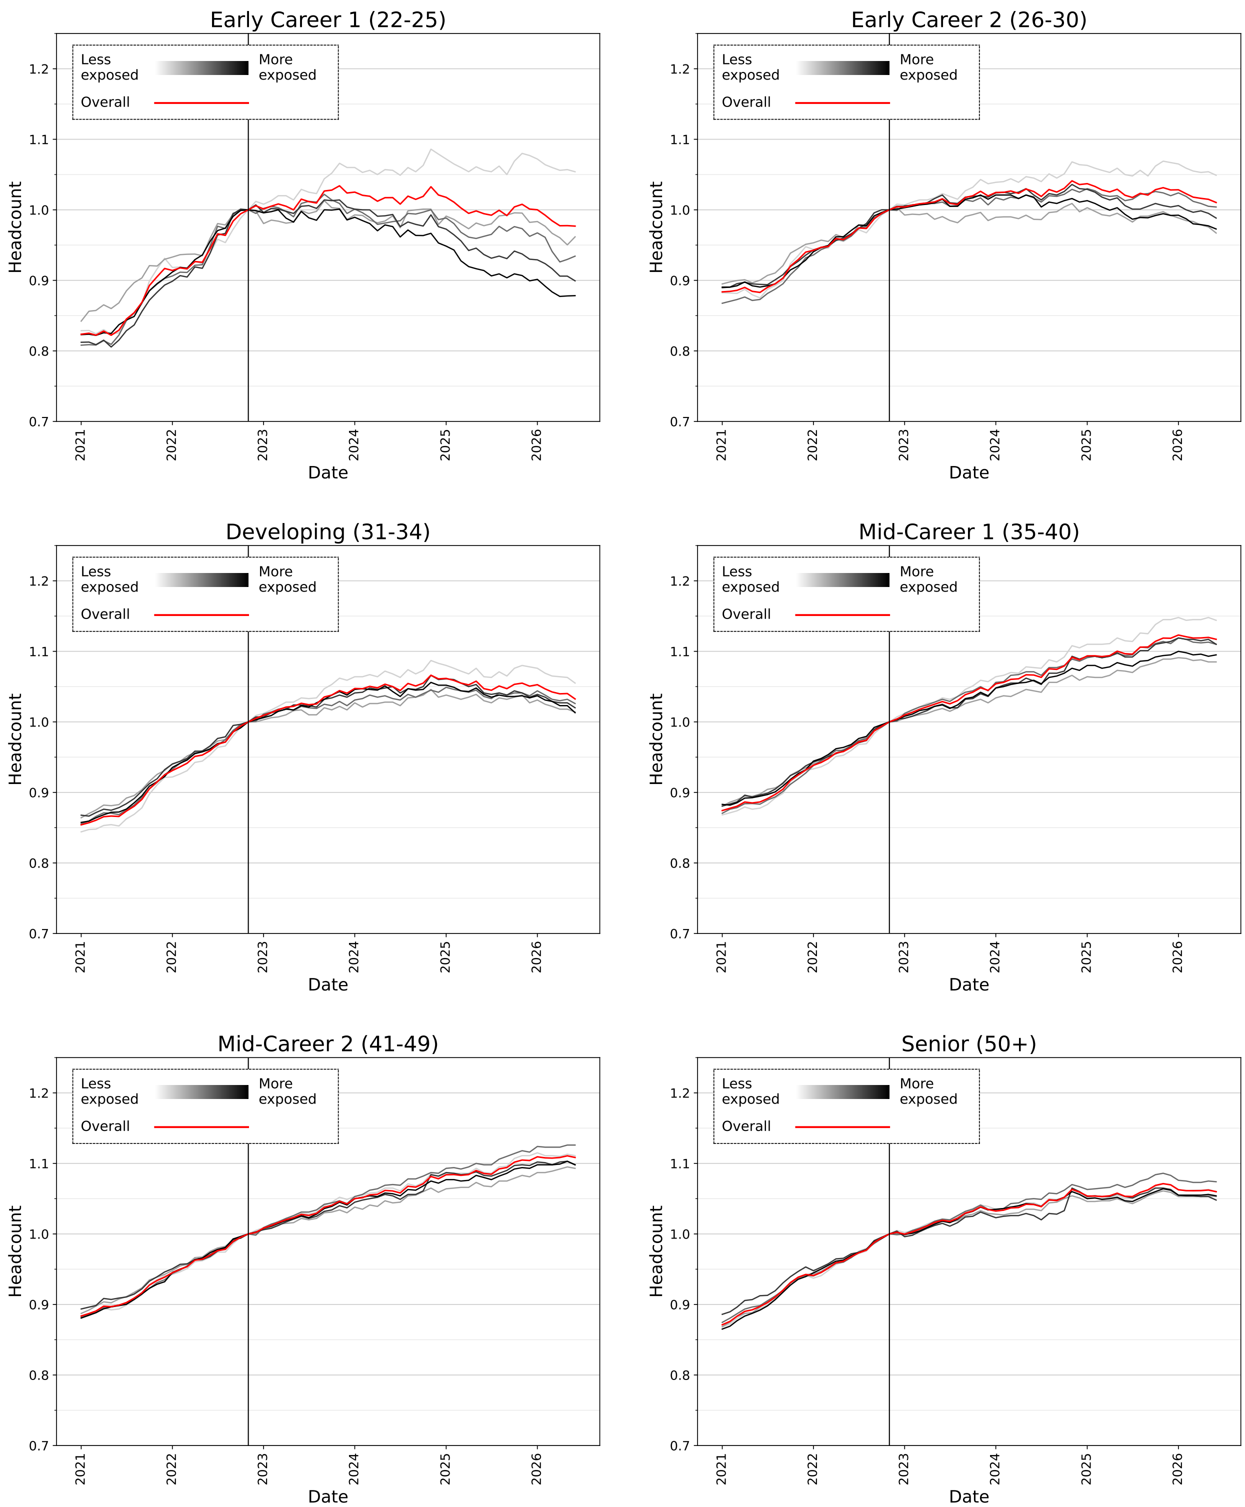

Figure F.4: Employment changes by age and automation quintile using Claude usage data from Handa et al. (2025). Excluding occupations that have no Claude usage or are in the lowest quintile of overall Claude usage conditional on some usage.

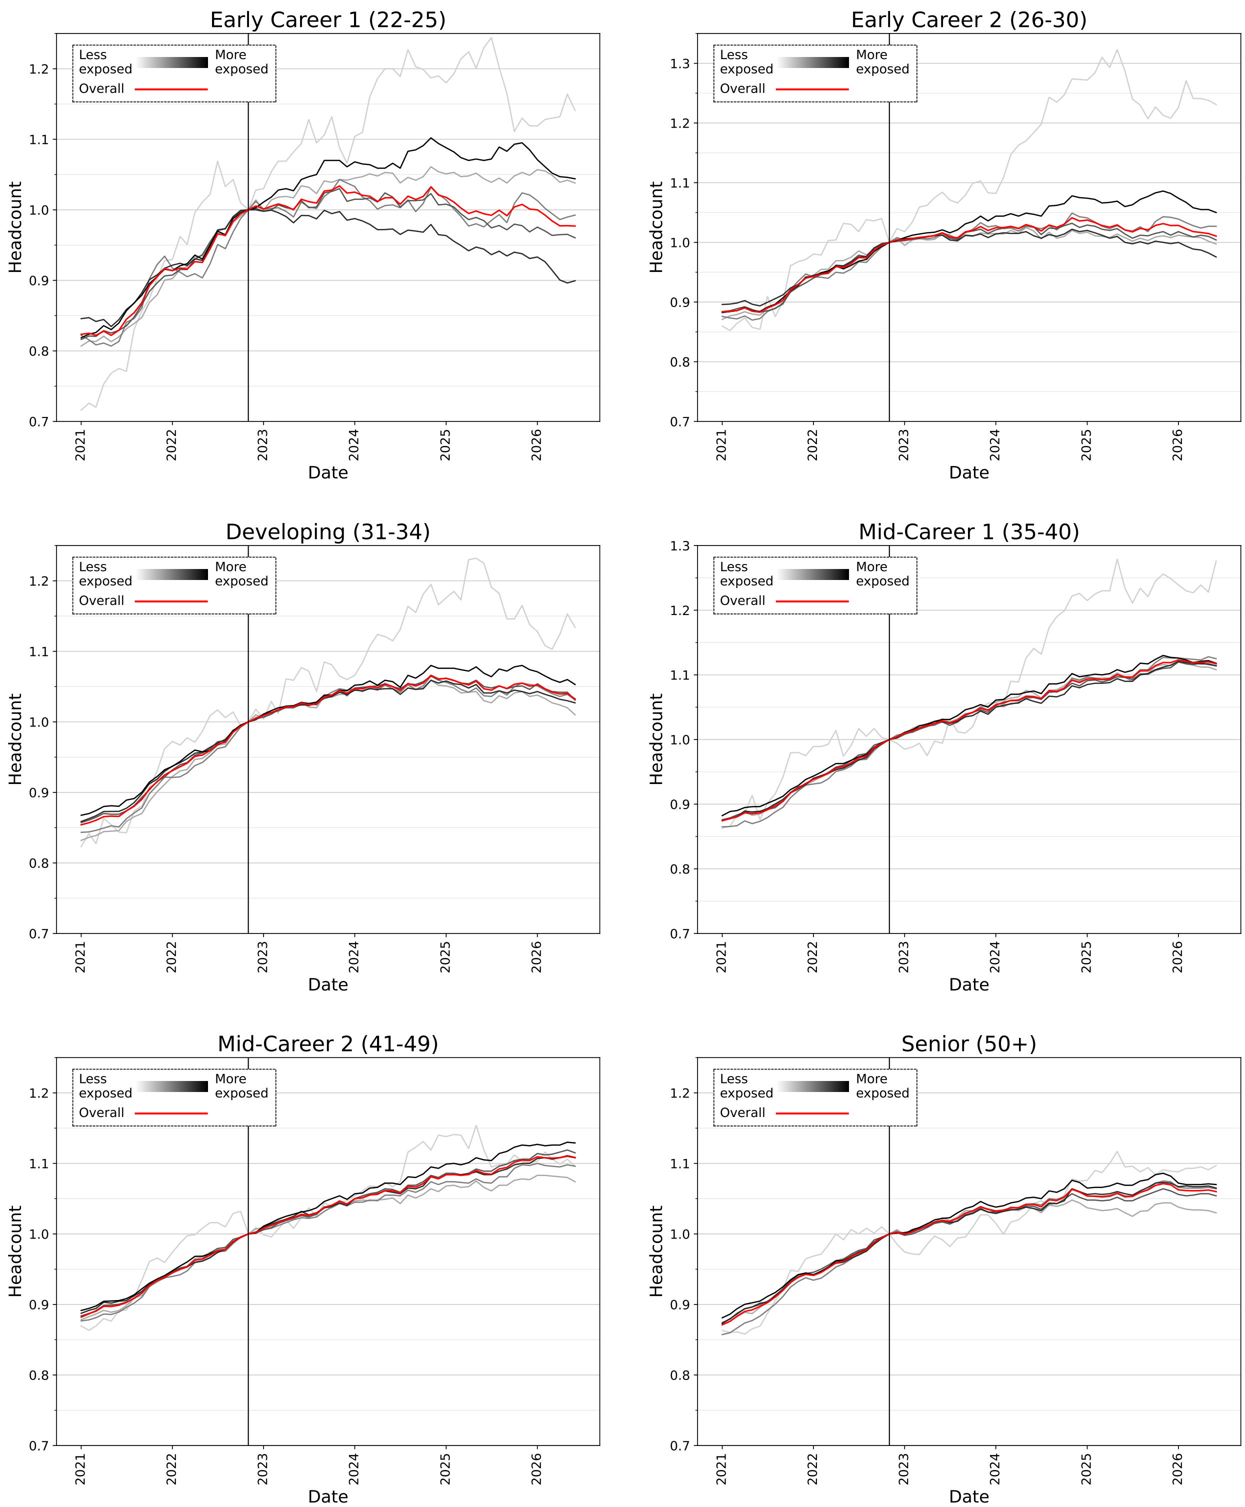

Figure F.5: Employment changes by age and complementarity ("augmentative") quintile using Claude usage data from Handa et al. (2025). Excluding occupations that have no Claude usage or are in the lowest quintile of overall Claude usage conditional on some usage.

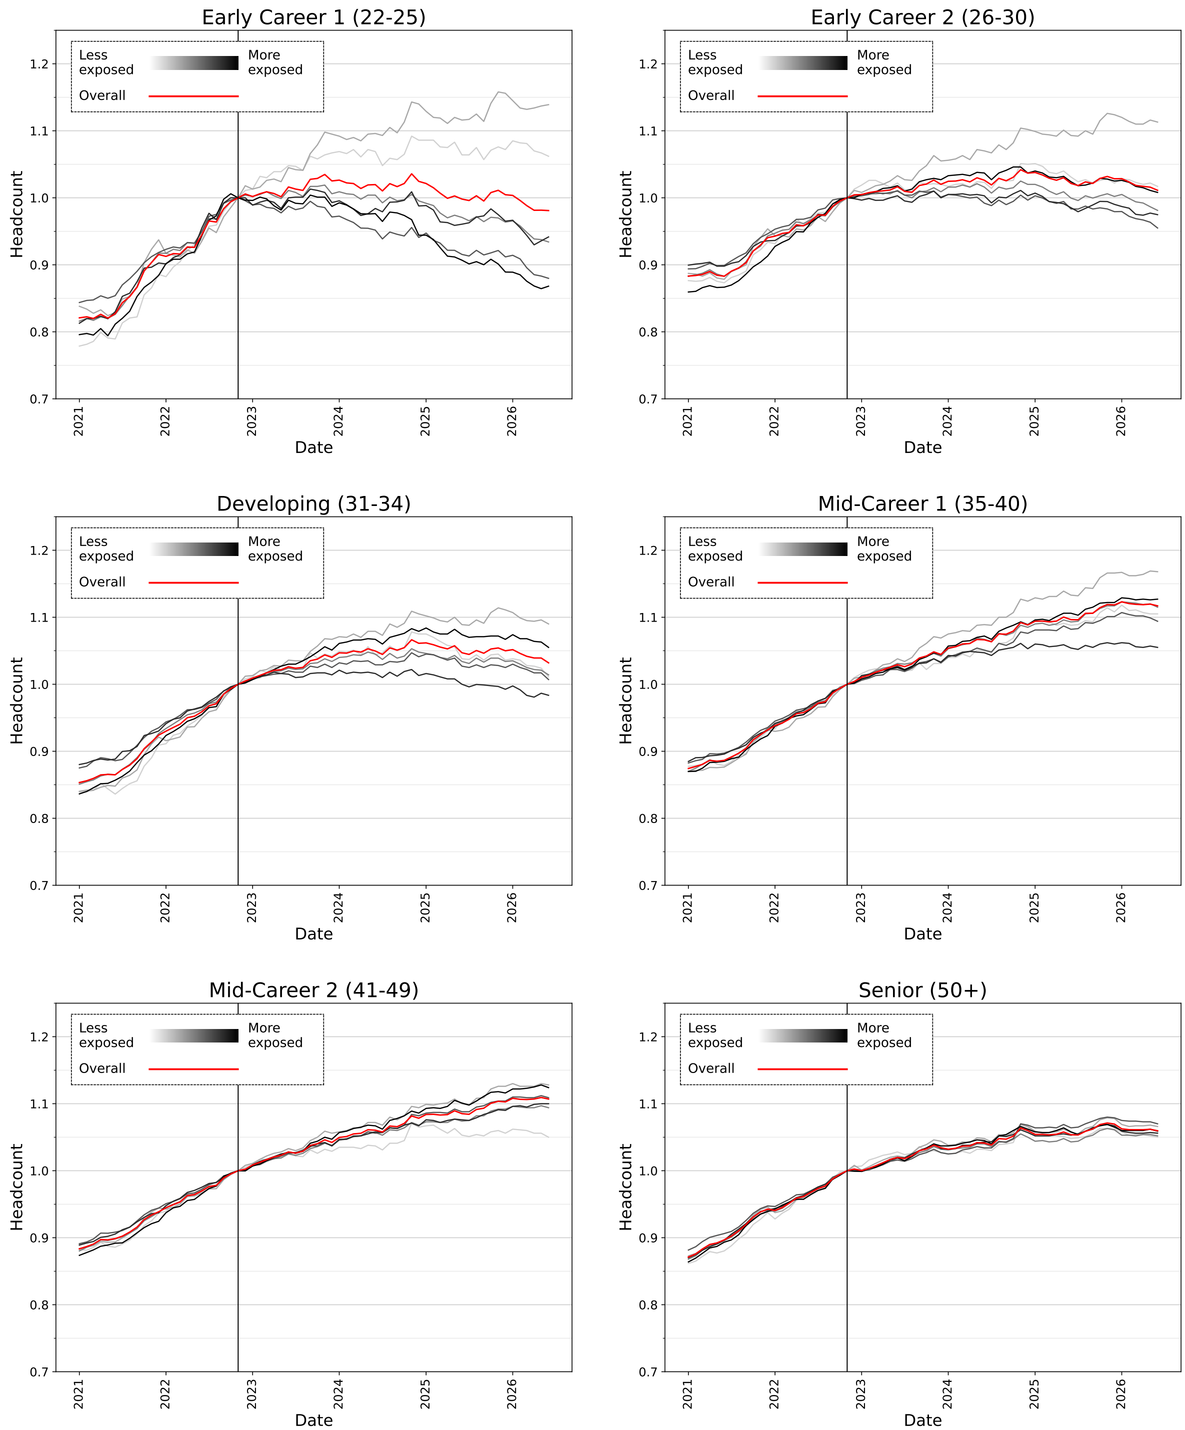

Figure F.6: Employment changes by overall Claude usage, using newer vintages of the Anthropic Economic Index.

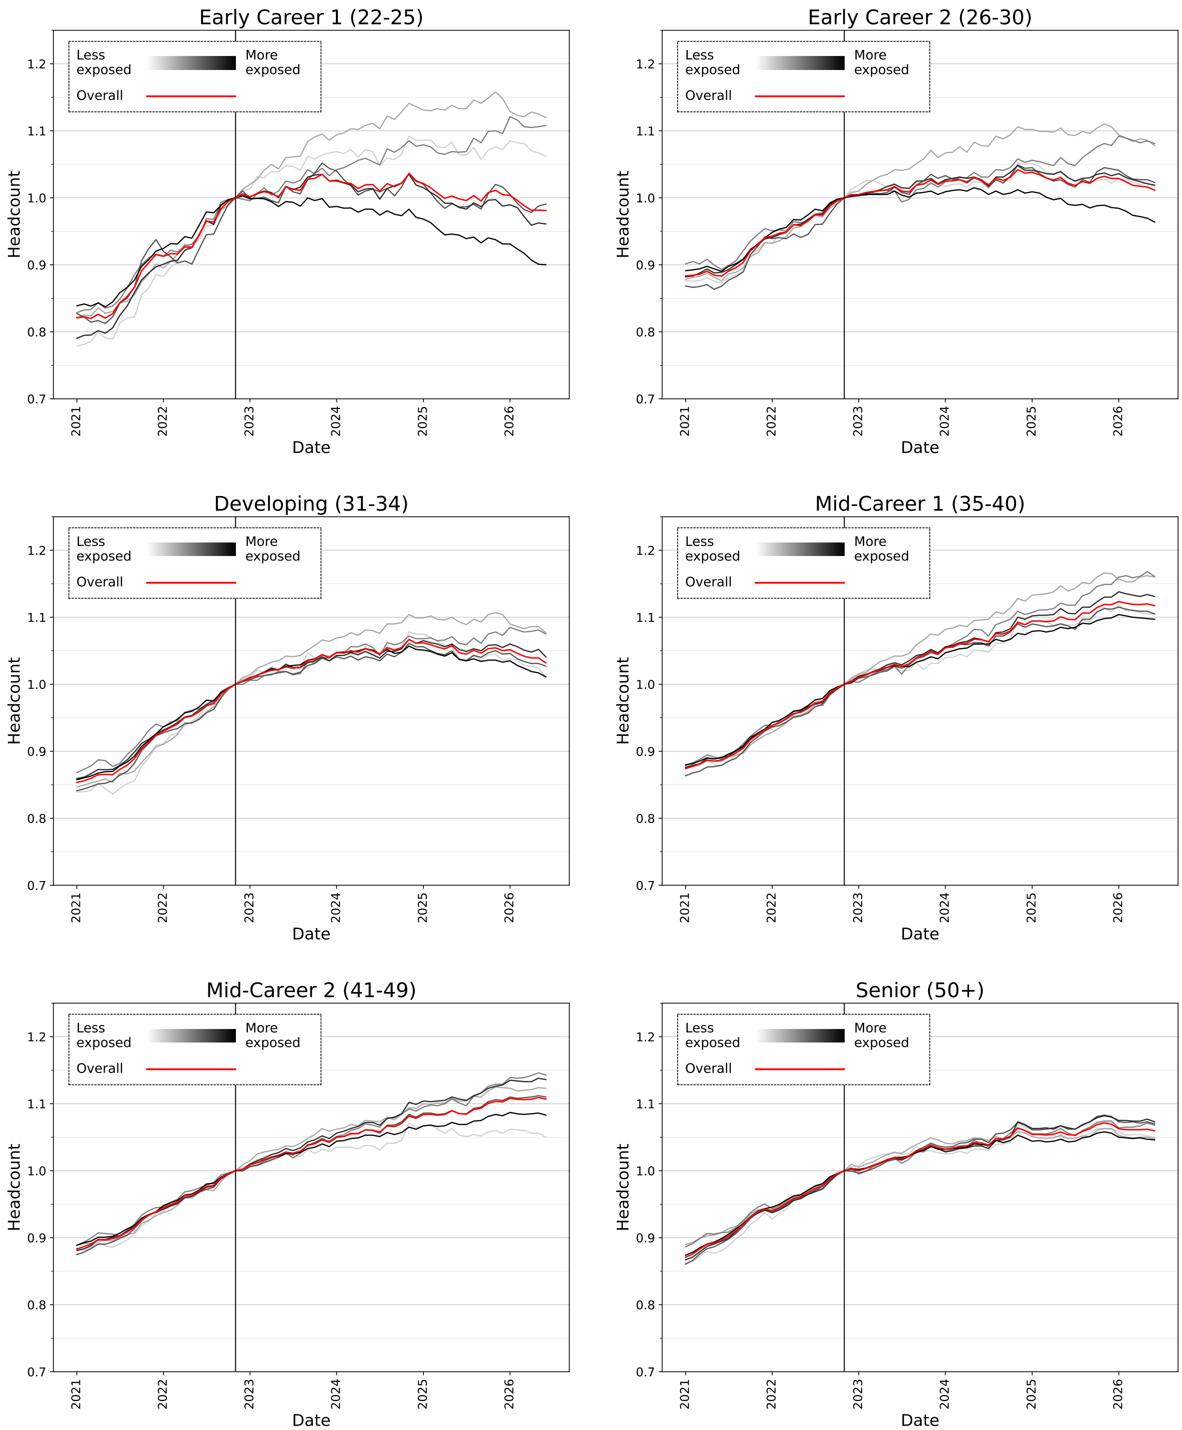

Figure F.7: Employment changes by automation exposure, using newer vintages of the Anthropic Economic Index.

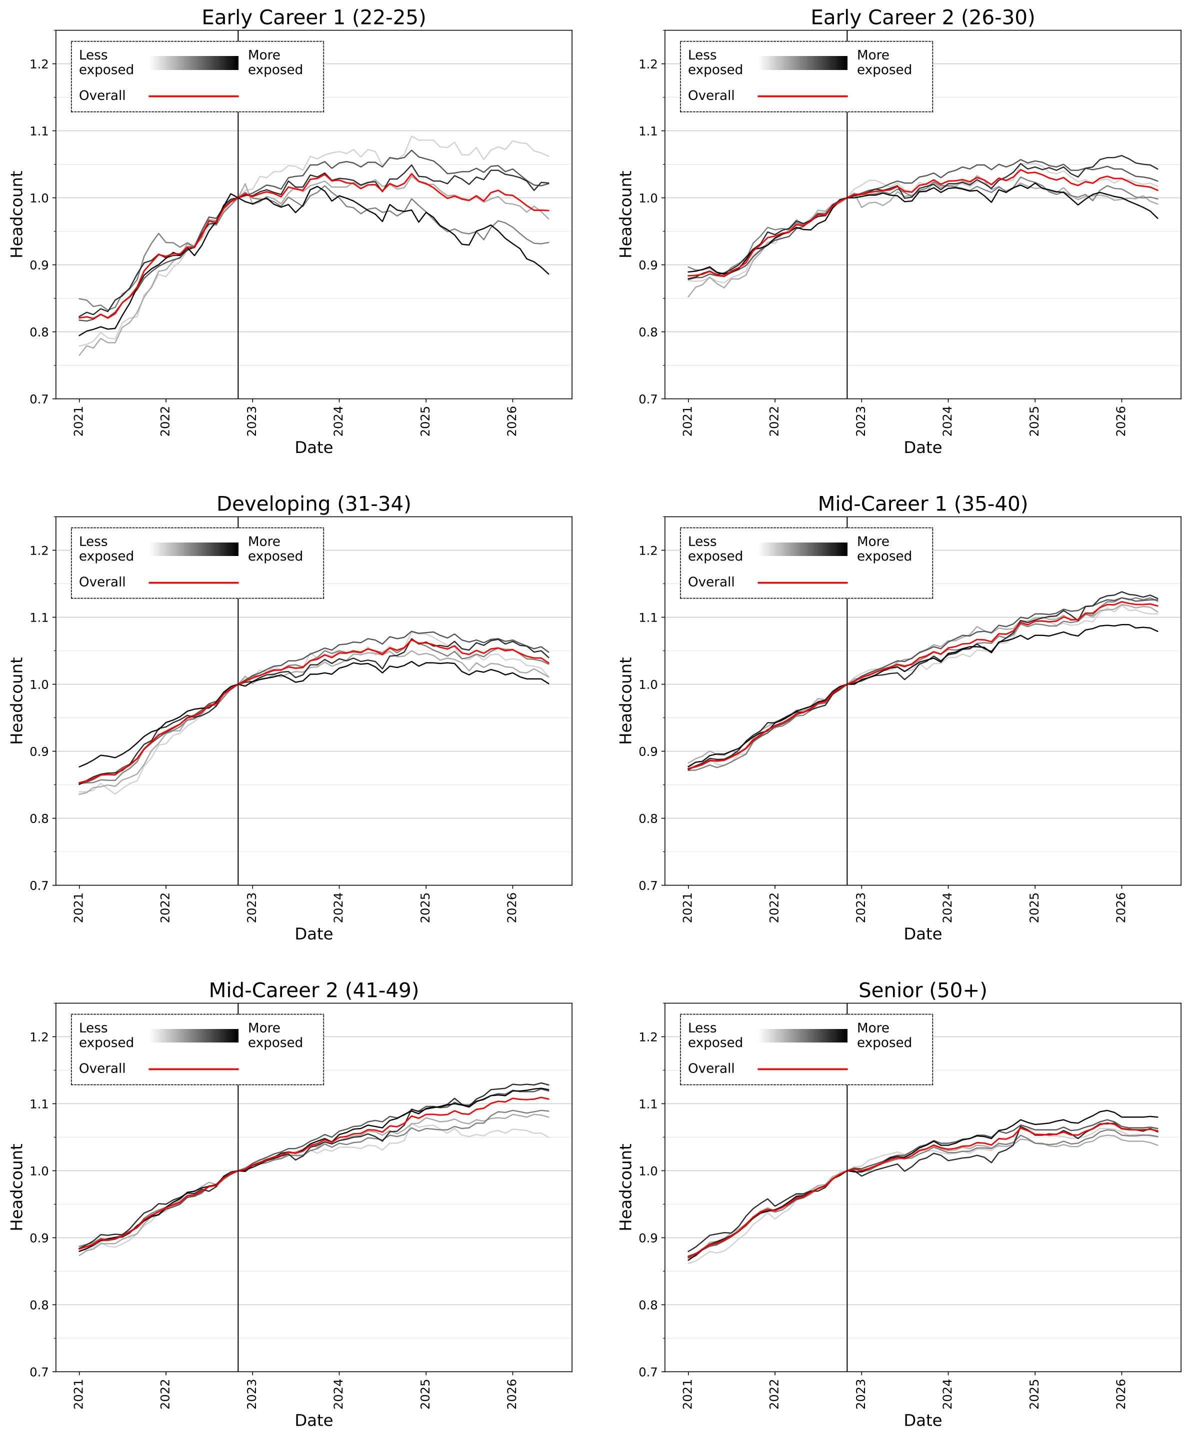

Figure F.8: Employment changes by complementarity ("augmentative") exposure, using newer vintages of the Anthropic Economic Index.

G Compensation

This section reports changes in real base compensation. Figure G.1 shows changes by age and occupation, and Figure G.2 by age and exposure quintile. Figures G.3 and G.4 separate the two margins along which pay can adjust, reporting base pay for job-stayers and for newly hired workers. All compensation measures are deflated to 2017 dollars using the PCE deflator.

Figure G.1: Changes in annual base compensation by age and occupation. Annual base compensation is deflated to 2017 dollars using the PCE deflator.

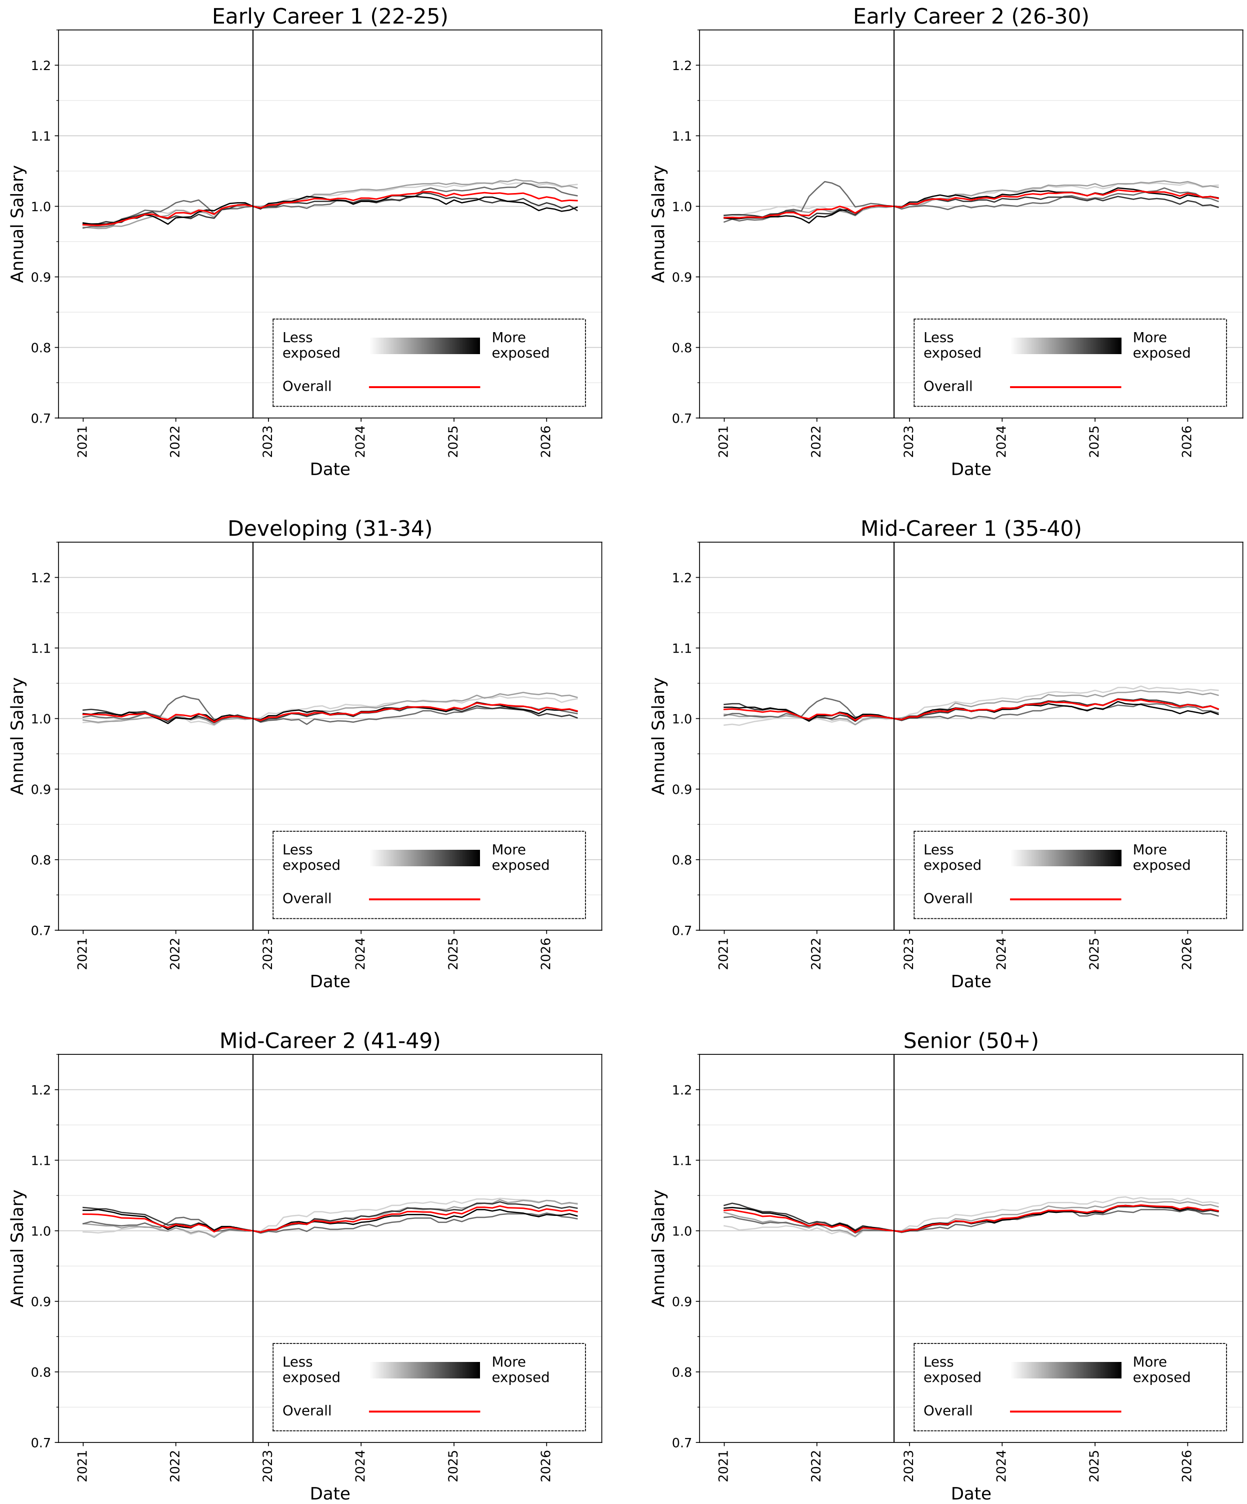

Figure G.2: *Annual base compensation.* Changes in annual base compensation by age and exposure quintile. Exposure quintiles are defined based on the GPT-4 β measures from Eloundou et al. (2024). Darker lines are more exposed quintiles. The red line shows the overall trend pooling across quintiles. Annual base compensation is deflated to 2017 dollars using the PCE deflator.

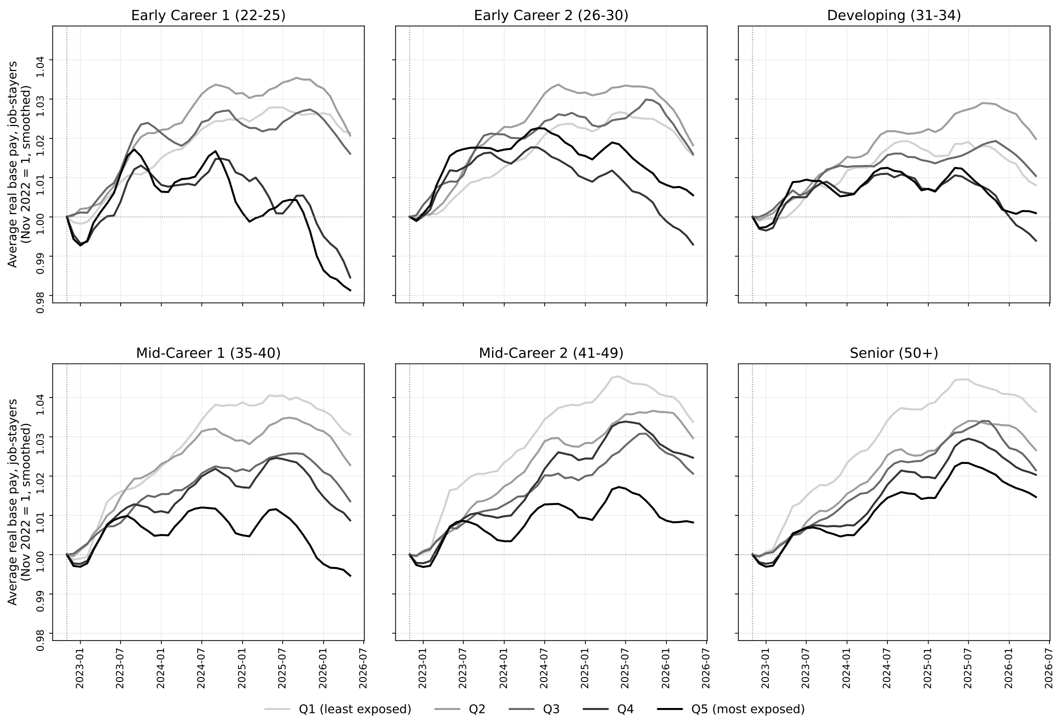

Figure G.3: *Job-stayer base pay.* Average real annual base pay for job-stayers (workers who remain employed at the same firm), by age group and GPT-4 β exposure quintile, shown as a three-month trailing moving average indexed to November 2022. Darker lines are more exposed quintiles; the vertical line marks November 2022. Base pay is deflated to 2017 dollars using the PCE deflator.

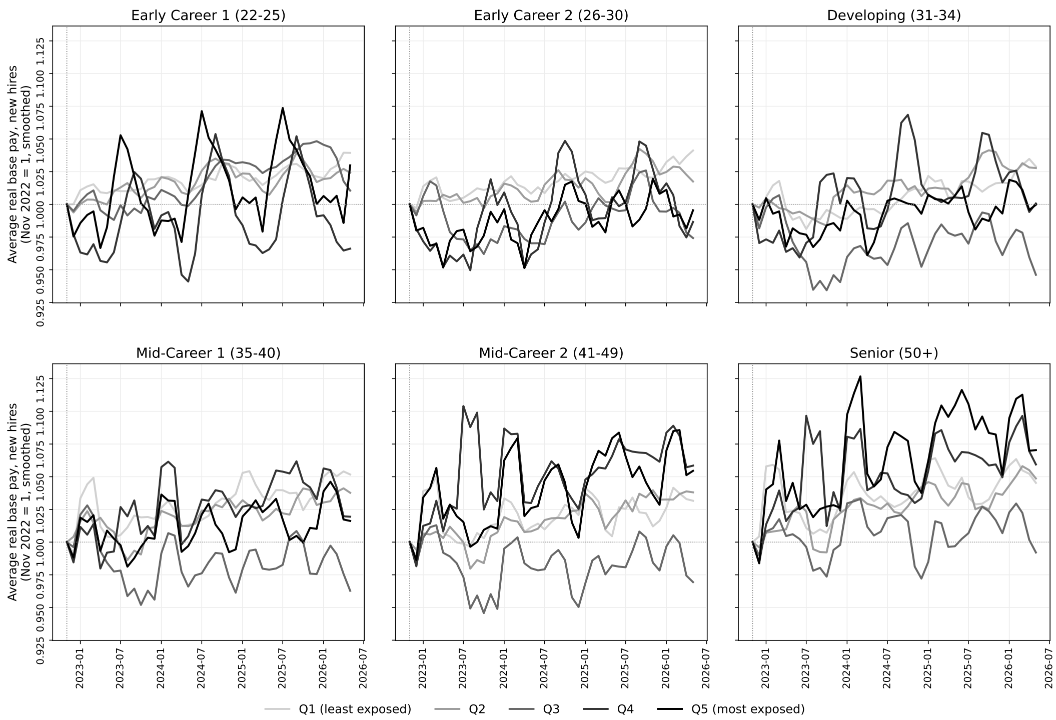

Figure G.4: *New-hire base pay.* Average real annual base pay of newly hired workers, by age group and GPT-4 β exposure quintile, shown as a three-month trailing moving average indexed to November 2022. Darker lines are more exposed quintiles; the vertical line marks November 2022. Base pay is deflated to 2017 dollars using the PCE deflator.

H Comparison to CPS and ACS Data

The CPS surveys about 60,000 households nationwide each month to collect data on employment and other labor force characteristics. These data are released a few weeks after the reference month, giving close to real-time estimates of employment statistics. A number of prior analyses have used the CPS to assess how AI is impacting entry-level work (Chandar, 2025b; Dominski and Lee, 2025; Lim et al., 2025; Eckhardt and Goldschlag, 2025). We compare some of our main findings in the ADP analysis sample to estimates from the CPS.

Figures H.1 through H.3 show employment changes by age for software developers, customer service representatives, and home health aides by age using data from the CPS. Though there are millions of workers employed in these professions across the US, the estimates are highly volatile, with common fluctuations of 20% or greater in estimated employment month-to-month. Figure H.4 shows estimated employment changes by age and exposure quintile using the CPS, also suggesting a high degree of volatility in the estimates.

This volatility in CPS microdata reflects small sample sizes and the fact that the CPS is not stratified to target employment statistics for these demographic-occupation subgroups.[42] The sample size and sampling procedure of the CPS may therefore make it challenging to assess employment changes by age and AI exposure with a high degree of confidence over the time horizon considered in this paper (Chandar, 2025b; O'Brien, 2025).

Other large-scale data sources such as the American Community Survey (ACS) may offer a more reliable comparison to the ADP analysis sample, though the ACS is released with a significant lag compared to the data from ADP. Fig. H.5 through Fig. H.7 show changes in the ACS for three case study occupations. Figure H.8 shows ACS employment by age and exposure quintile, using the same crosswalk and quintile definitions as the ADP results. Among the youngest workers, the two most exposed quintiles show the weakest employment growth by 2024, but the trends are considerably less pronounced than in the ADP analysis sample.[43]

[42] The CPS includes between 23 and 48 young software developers aged between 22 and 25 per month over our sample period. It includes between 30 and 63 young customer service representatives, and between 17 and 41 young home health aides.

[43] There are known questions about the weighting of the 2024 ACS, which incorporates a substantially revised population control whose effect the Census Bureau notes "was not evenly distributed amongst all demographic groups" and appears "most in the weights for the non-householder population" (U.S. Census Bureau, 2025a). Young workers in less exposed occupations are disproportionately non-householders—73 percent in the least exposed quintile against 62 percent in the most exposed—and their average person weight fell by more between 2022 and 2024. Table H.1

Table H.1 compares the two sources for workers aged 22–25. Across all employed young workers, the ACS gap between the most and least exposed quintiles is −0.022, with a confidence interval that includes zero. Restricting the ACS to civilian wage and salary workers employed full-time, which approximates ADP's employment relationship, leaves the estimate essentially unchanged at −0.019.

Table H.2 shows how these estimates vary by industry. ADP is concentrated in professional, information, and financial services, which account for 34 percent of its young employment against 22 percent in the ACS, and where the exposure gradient is steeply negative in both sources. It is correspondingly light in education, health, and public administration, 16 percent against 24 percent, where the ACS gradient is positive.

Table H.3 repeats this sector breakdown alongside the ratio of unweighted respondent counts. The supplementary diagnostic is motivated by Census guidance cautioning against comparing population totals across years, due in part to updates to the population controls (U.S. Census Bureau, 2025b). Doing so moves the ACS gap from −0.022 to −0.047, and from −0.019 to −0.042 on the sample restricted to full-time civilian wage and salary workers. Weighting changes impact the education, health, and public administration sector in particular, with Q1 employment growth 16 percentage points higher when unweighted and the Q1 to Q5 gap shrinking by 13 percentage points.

reports a comparison based on unweighted respondent counts.

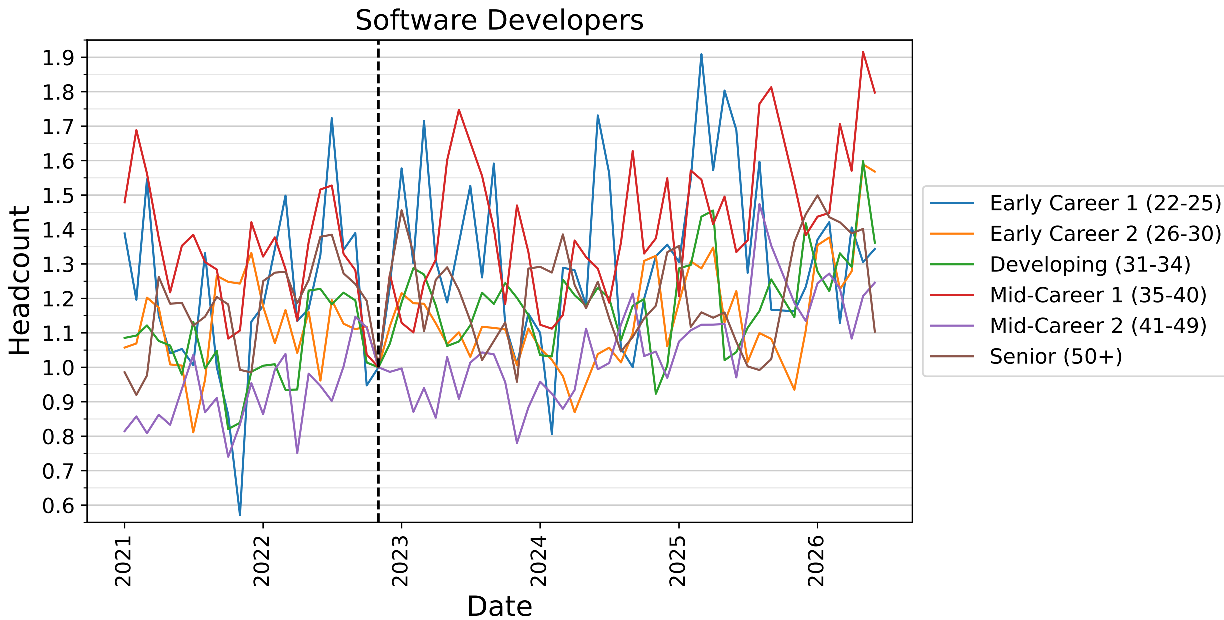

Figure H.1: Employment changes for software developers by age, normalized to 1 in November 2022. Data come from the monthly CPS through June 2026.

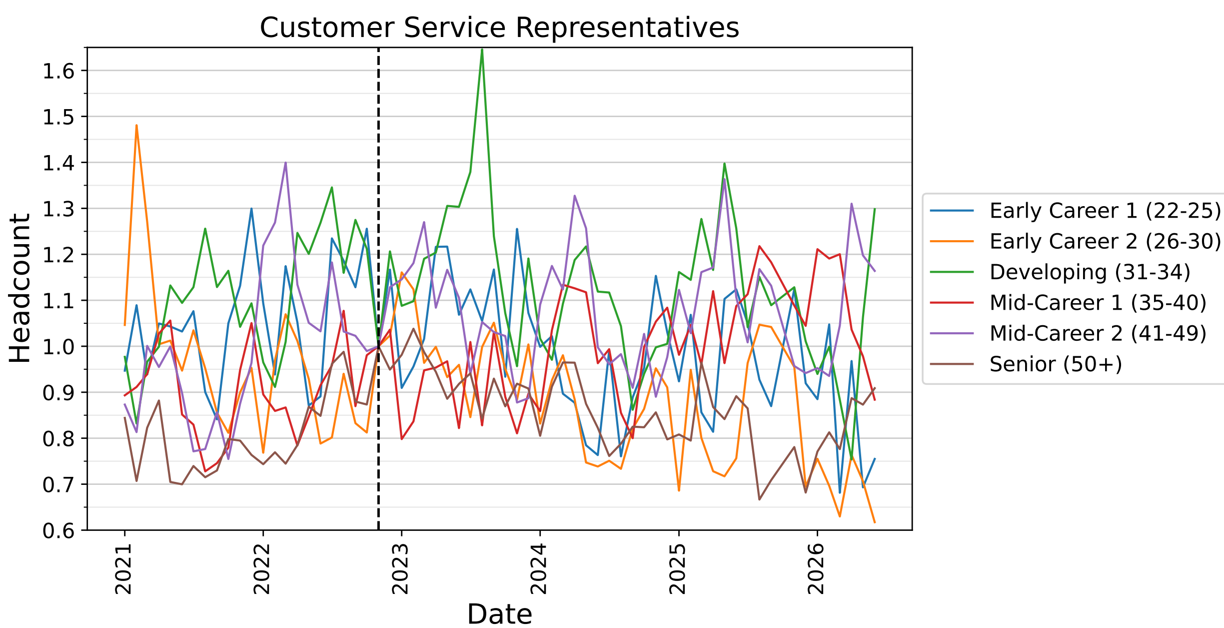

Figure H.2: Employment changes for customer service representatives by age, normalized to 1 in November 2022. Data come from the monthly CPS through June 2026.

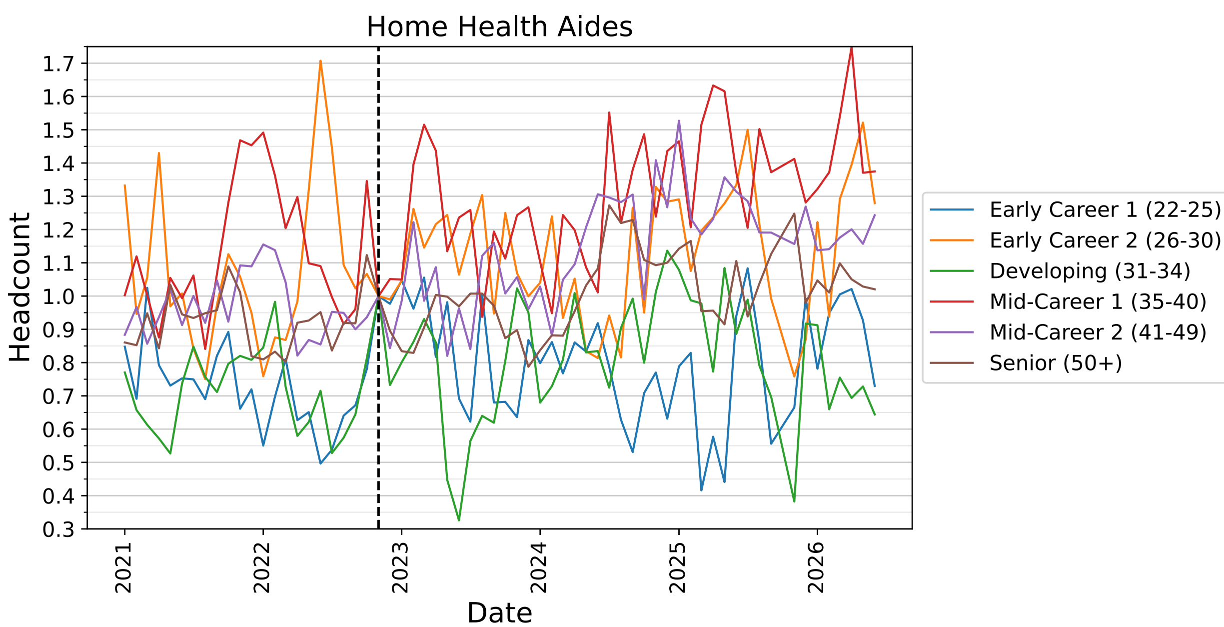

Figure H.3: Employment changes for home health aides by age, normalized to 1 in November 2022. Data come from the monthly CPS through June 2026.

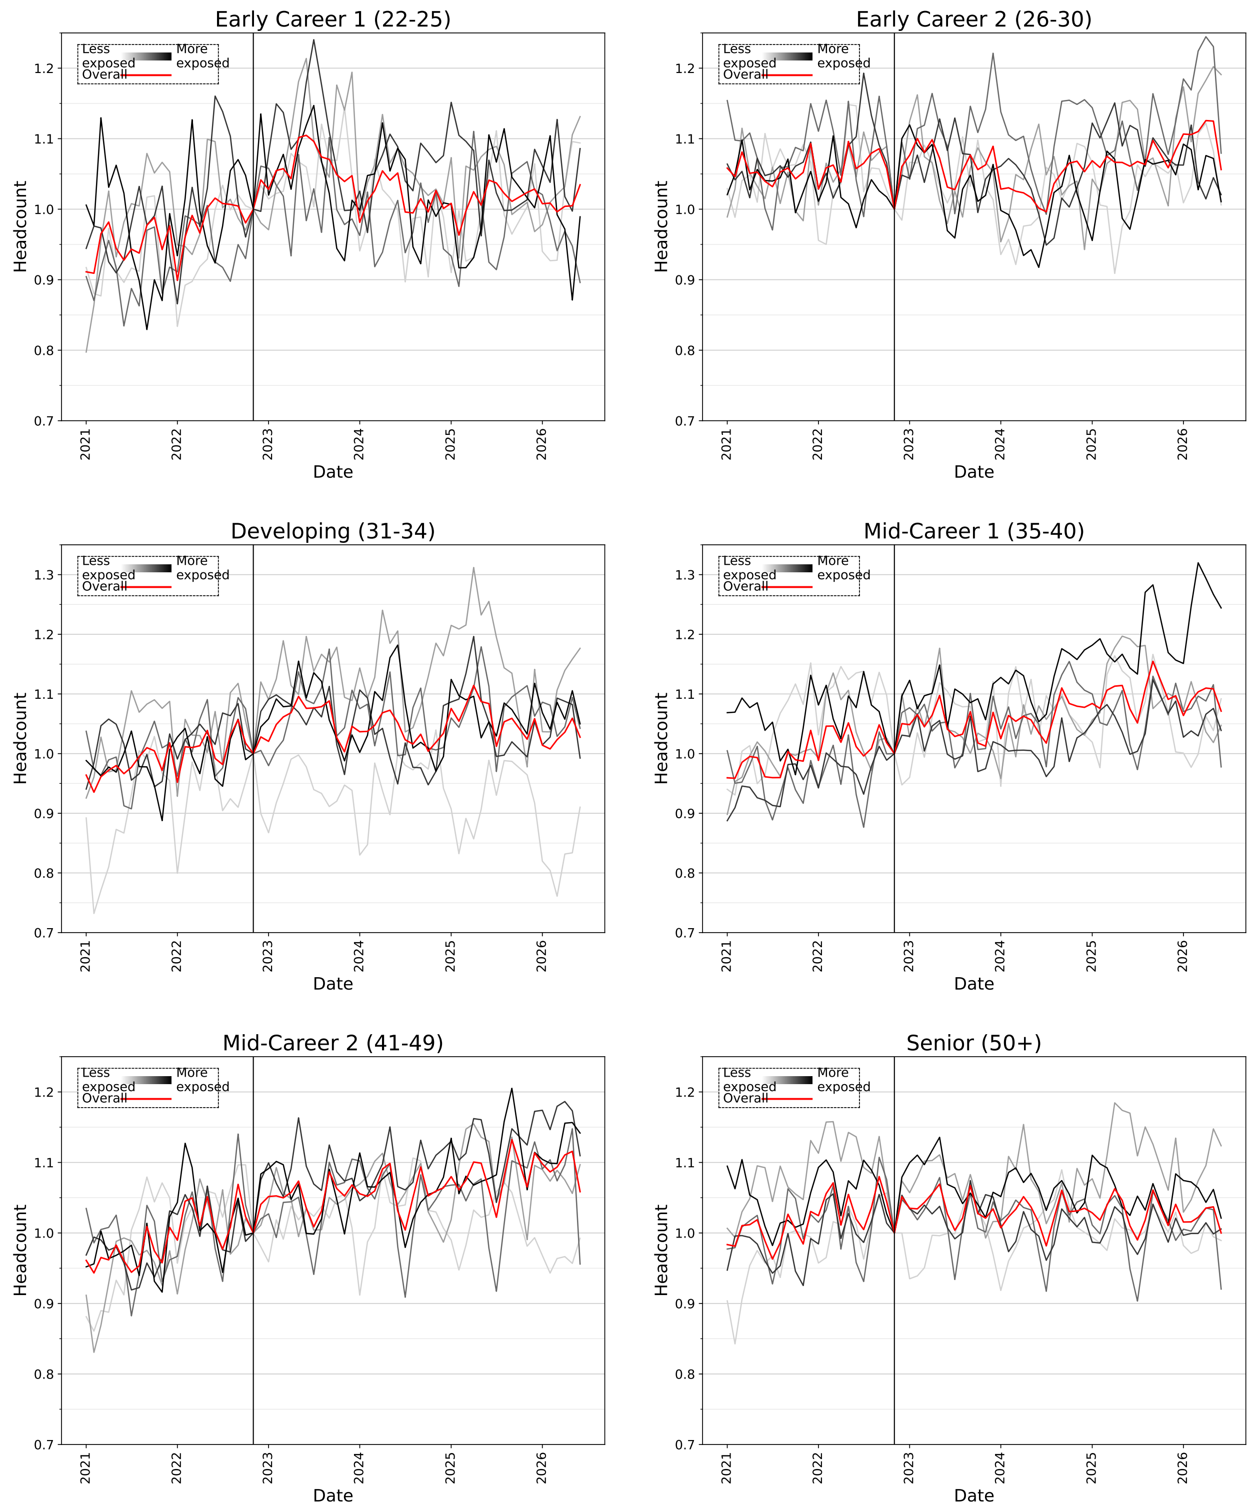

Figure H.4: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Exposure quintiles use the pipeline crosswalk from 2010/2018 Census occupation codes and are directly comparable to the ADP quintiles. Data come from the monthly CPS through June 2026.

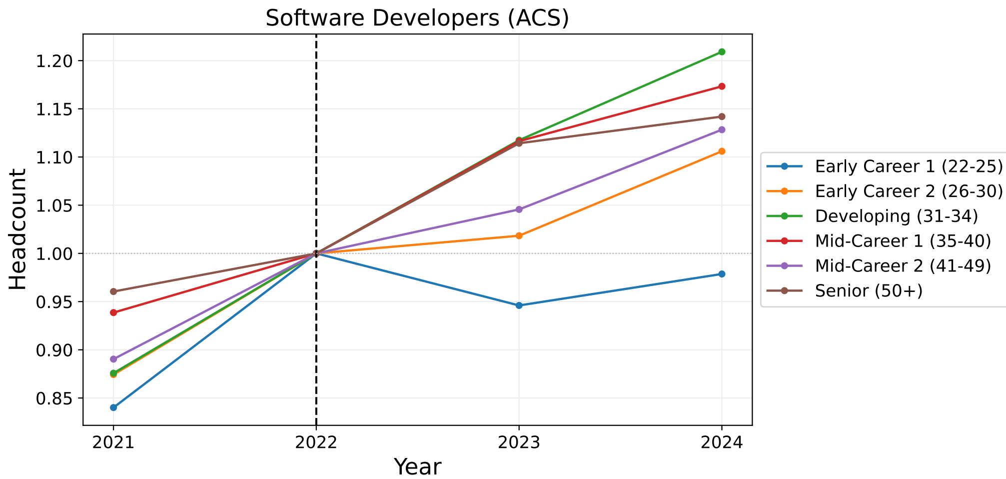

Figure H.5: Employment changes for software developers by age using annual ACS data, normalized to 1 in 2022. Compare with the monthly-CPS version in Fig. H.1.

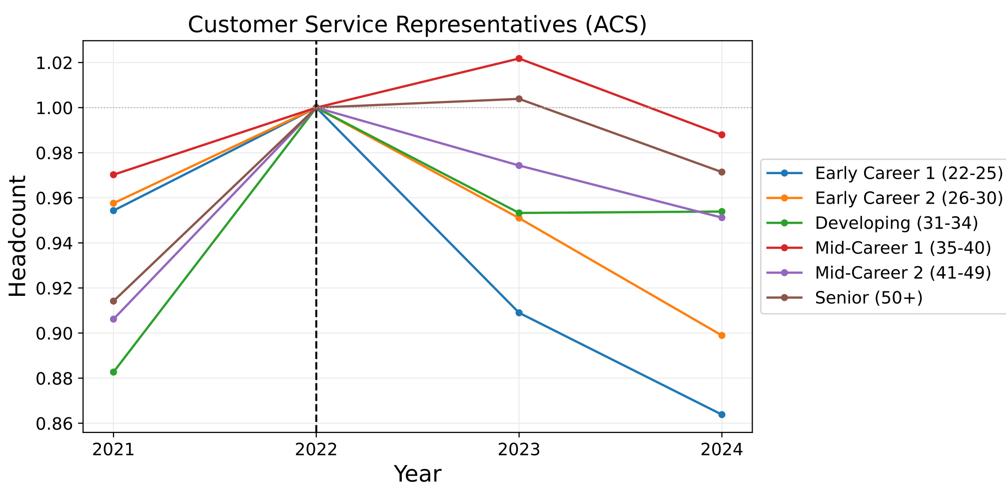

Figure H.6: Employment changes for customer service representatives by age using annual ACS data, normalized to 1 in 2022. Compare with the monthly-CPS version in Fig. H.2.

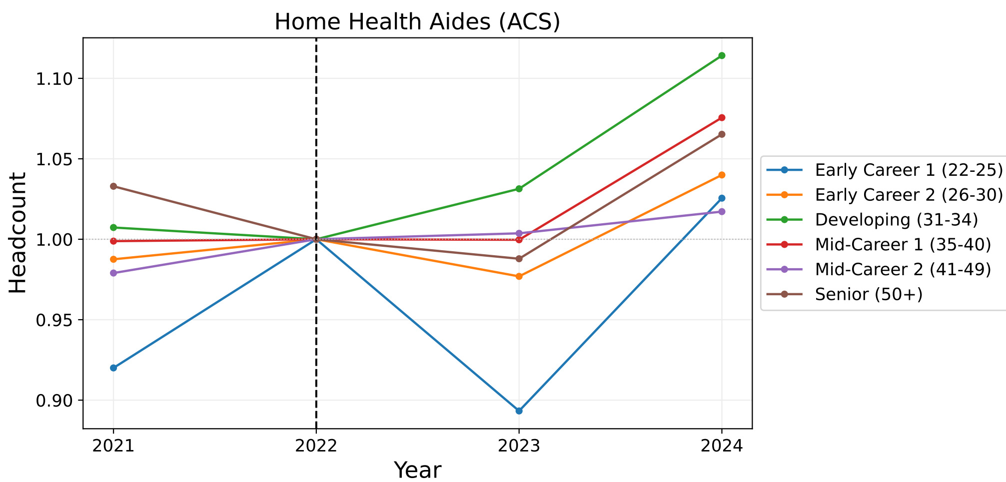

Figure H.7: Employment changes for home health aides by age using annual ACS data, normalized to 1 in 2022. Compare with the monthly-CPS version in Fig. H.3.

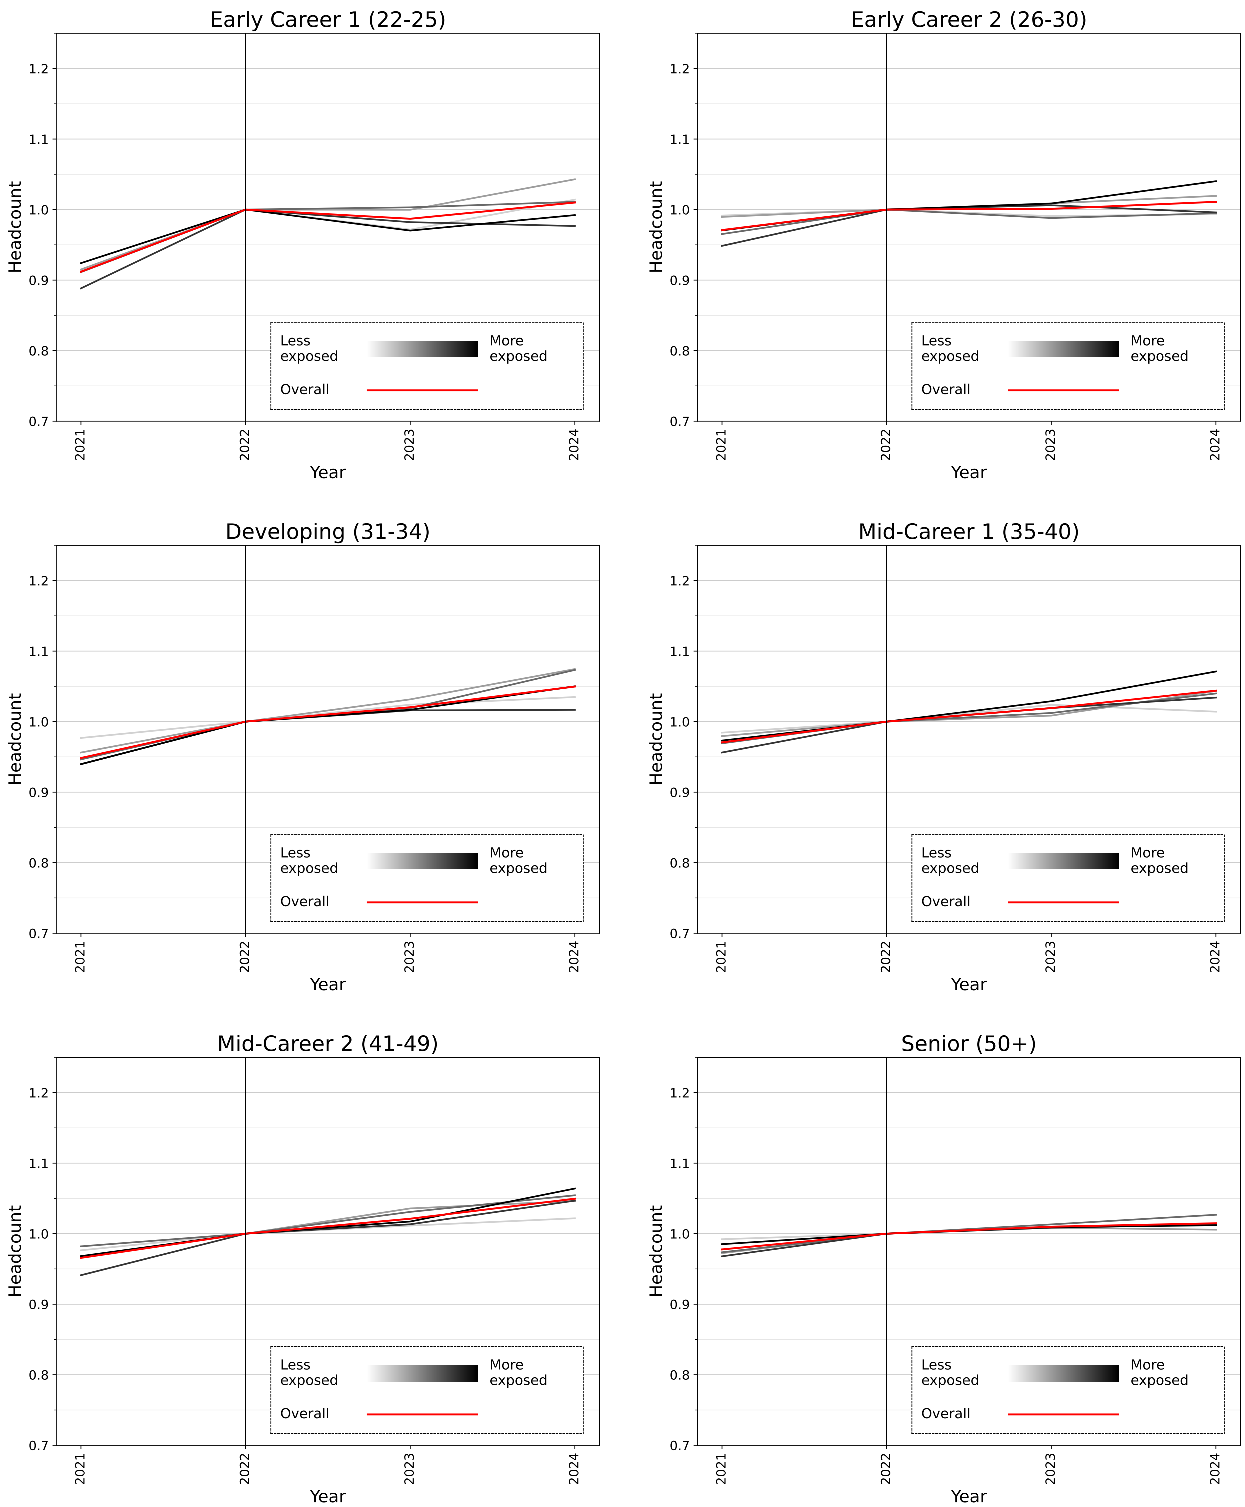

Figure H.8: Employment changes by age and exposure quintile in the ACS, normalized to 1 in 2022. Darker lines are more exposed quintiles; the red line is the overall trend pooling across quintiles. Annual ACS data through 2024, the latest available wave.

Table H.1: Comparing the ACS and ADP estimates of the young-worker AI-exposure gradient. Each row reports the difference between the fifth and first exposure quintiles in employment growth from 2022 to 2024 for workers aged 22–25, as the ACS sample is restricted toward ADP's sample conditions. Confidence intervals come from the 80 Census successive-difference replicate weights. The final ACS row reports the ratio of unweighted respondent counts, which is invariant to the Census person weights and therefore carries no interval.

	Quintile 5 − Quintile 1
ACS, all employed aged 22–25	−0.022 [-0.055, +0.011]
less armed forces	−0.024 [-0.057, +0.010]
less unpaid family workers	−0.024 [-0.058, +0.010]
less self-employed	−0.020 [-0.055, +0.015]
full-time only	−0.019 [-0.059, +0.020]
same, unweighted respondent counts	−0.042
ADP (balanced panel)	−0.132
Residual, published weights	*−0.112*
Residual, respondent counts	*−0.089*

Table H.2: The young-worker AI-exposure gradient by sector, ACS versus ADP. The first two columns give each sector's share of employment among workers aged 22–25 in 2022. The remaining columns report, within each sector, employment in 2024 relative to 2022 for the least and most exposed quintiles separately, and the difference between them. ACS 95% confidence intervals appear in brackets beneath each estimate and are computed from the 80 Census successive-difference replicate weights.

	Employment share		ACS			ADP		
Sector	ACS	ADP	Q1	Q5	Q5 − Q1	Q1	Q5	Q5 − Q1
Professional, information, and finance	0.215	0.340	1.133	0.920	−0.213	1.170	0.939	−0.231
			[0.998, 1.268]	[0.882, 0.958]	[-0.353, -0.072]			
Goods-producing	0.195	0.229	0.965	1.022	+0.057	1.113	1.075	−0.037
			[0.915, 1.015]	[0.959, 1.085]	[-0.023, 0.136]			
Trade, transportation, and utilities	0.208	0.181	0.994	0.945	−0.049	1.151	1.041	−0.109
			[0.923, 1.066]	[0.882, 1.008]	[-0.142, 0.043]			
Education, health, and public administration	0.244	0.155	0.881	1.126	+0.245	1.146	1.073	−0.074
			[0.760, 1.003]	[1.066, 1.187]	[0.104, 0.386]			
Leisure and other services	0.137	0.096	1.074	1.060	−0.015	1.125	1.045	−0.080
			[0.986, 1.163]	[0.956, 1.163]	[-0.141, 0.111]			

Table H.3: The young-worker AI-exposure gradient by sector in the ACS, under the published Census weights and as a ratio of unweighted respondent counts. Columns report employment in 2024 relative to 2022 for the least and most exposed quintiles and the difference between them. The published-weight columns replicate the ACS block of Table H.2; 95% confidence intervals appear in brackets beneath each estimate and are computed from the 80 Census successive-difference replicate weights. The respondent-count columns are invariant to the Census person weights and carry no intervals.

	Published weights			Unweighted respondent counts		
Sector	Q1	Q5	Q5 − Q1	Q1	Q5	Q5 − Q1
All sectors	1.001	0.982	−0.019	1.042	0.999	−0.042
	[0.968, 1.034]	[0.955, 1.008]	[-0.059, 0.020]			
Professional, information, and finance	1.133	0.920	−0.213	1.181	0.937	−0.244
	[0.998, 1.268]	[0.882, 0.958]	[-0.353, -0.072]			
Goods-producing	0.965	1.022	+0.057	0.987	1.043	+0.056
	[0.915, 1.015]	[0.959, 1.085]	[-0.023, 0.136]			
Trade, transportation, and utilities	0.994	0.945	−0.049	1.042	0.945	−0.097
	[0.923, 1.066]	[0.882, 1.008]	[-0.142, 0.043]			
Education, health, and public administration	0.881	1.126	+0.245	1.046	1.158	+0.112
	[0.760, 1.003]	[1.066, 1.187]	[0.104, 0.386]			
Leisure and other services	1.074	1.060	−0.015	1.108	1.066	−0.042
	[0.986, 1.163]	[0.956, 1.163]	[-0.141, 0.111]			

I Replicating the Work-from-Home Specification

Lambert and Schindler (2026) study the relationship between the entry-level share of new hires, AI exposure, and remote work. We replicate their specification in the ADP analysis sample as closely as possible.

We estimate their state-by-occupation event-study specification, their equation (8):

$$y_{it} = \sum_{k \in K} \sum_{\tau \neq \tau_0} \phi_{k\tau} \left(Z_i^k \times \mathbf{1}\{\text{year}(t) = \tau\} \right) + \alpha_i + \delta_t + \varepsilon_{it}, \tag{I.1}$$

where i is a state-by-detailed-occupation cell, y_{it} is the junior share of new hires in percentage points, and Z_i^k is the standardized exposure of occupation i to shock $k \in K = \{\text{WFH}, \text{GenAI}\}$. AI exposure is the GPT-4 β measure of Eloundou et al. (2024) and work-from-home exposure is the measure of Lambert and Schindler (2026). The specification includes state-by-detailed-occupation unit fixed effects α_i and state-by-month time fixed effects δ_t, is weighted by cell hires, and standard errors are two-way clustered by occupation and state. The coefficient $\phi_{k\tau}$ is the year-specific exposure gradient per standard deviation of Z_i^k, relative to the omitted year τ_0. When only one exposure is included, K is a singleton; the contrast between the single- and joint-treatment estimates is the object of interest.

Our replication uses 2019 as the reference year, as in Lambert and Schindler (2026). Table I.1 and Figure I.1 report the estimates.

Estimated separately, both exposures enter negatively in the post-ChatGPT years: occupations more exposed to AI and occupations more amenable to remote work both saw their junior hiring share fall. The AI gradient is the larger of the two and grows steadily, from −0.48 in 2022 to −1.06 by 2025.

When entering both jointly, the work-from-home gradient shrinks. The AI gradient grows, from −0.64 in 2022 to −1.52 by 2025. In the ADP analysis sample the joint variation therefore loads on AI exposure rather than on remote work, the opposite of what Lambert and Schindler (2026) find in the Revelio data, where the AI gradient is the one that does not survive joint entry. As the specifications are identical the different outcomes are driven by differences in the samples.

Table I.1: Estimates of equation (I.1) in the ADP data, replicating the state-by-occupation specification of Lambert and Schindler (2026). The outcome is the junior share of new hires in percentage points, annualized as in their specification. AI exposure is the GPT-4 β measure of Eloundou et al. (2024) and work-from-home exposure is the measure of Lambert and Schindler (2026); both are standardized, so coefficients are percentage points per standard deviation of exposure relative to the omitted year, 2019, the same reference year Lambert and Schindler (2026) use. 2026 is excluded because it is a partial year. Each exposure is estimated on its own in the first two panels and the two are estimated jointly in the third. Standard errors, two-way clustered by occupation and state, in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	2018	2019 (omitted)	2020	2021	2022	2023	2024	2025
AI only	-0.43^{**}	0	-0.10	-0.17	-0.48^{**}	-0.76^{***}	-0.81^{***}	-1.06^{***}
	(0.18)		(0.15)	(0.15)	(0.21)	(0.26)	(0.29)	(0.30)
WFH only	-0.15	0	0.12	0.00	-0.18	-0.31	-0.46^{*}	-0.32
	(0.17)		(0.16)	(0.18)	(0.20)	(0.25)	(0.24)	(0.31)
Joint: AI	-0.59^{***}	0	-0.34^{*}	-0.32	-0.64^{**}	-0.98^{***}	-0.85^{***}	-1.52^{***}
	(0.21)		(0.18)	(0.22)	(0.26)	(0.30)	(0.31)	(0.38)
WFH	0.20	0	0.32^{*}	0.20	0.20	0.28	0.05	0.59^{*}
	(0.20)		(0.19)	(0.26)	(0.25)	(0.27)	(0.21)	(0.30)

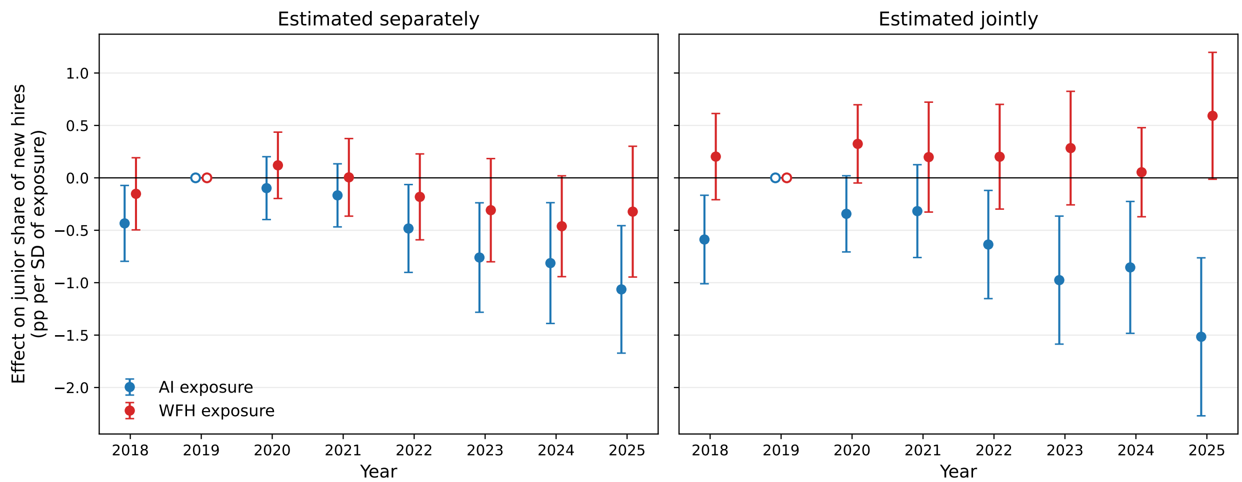

Figure I.1: Effect of AI exposure and work-from-home exposure on the junior share of new hires, by year, replicating the Lambert and Schindler (2026) specification in the ADP analysis sample. Both exposures are standardized, so estimates are percentage points per standard deviation, relative to 2019. Bars are 95% confidence intervals. In the left panel each exposure is estimated on its own; in the right panel they are estimated jointly.

J Codified and Tacit Knowledge Index

We construct new scores for whether a job requires codified and tacit knowledge. We start with the 7,514 raw job titles in ADP. Each of these job titles also includes a job description constructed by ADP.

We query GPT-5.4-mini (OpenAI) with temperature = 0 to score each job independently on two dimensions: codified knowledge (1–10) and tacit knowledge (1–10). Each job receives two separate API calls. The model receives a shared system prompt that teaches it the codified/tacit distinction using O*NET knowledge domains, work activities, and cross-functional skills as training examples, followed by a dimension-specific user prompt.

Below are the shared system prompt and the individual prompts for the codified and tacit dimensions.

System prompt (shared across both dimensions)

You are an expert at classifying job descriptions based on the type of knowledge they require.

Definitions

Codified knowledge: formal, standardized, documented knowledge that can be taught through education, textbooks, or written procedures. It is explicit and transferable through text.

Tacit knowledge: implicit, experience-based knowledge that is difficult to write down or teach in a classroom. It is acquired through practice, mentorship, hands-on experience, and repeated exposure to real situations.

Learn From the Following Classification

We have classified O*NET occupational knowledge domains, work activities, and cross-functional skills into codified vs. tacit. Study the descriptions and classifications below to understand what makes knowledge codified or tacit, then apply this understanding to score job descriptions.

Knowledge Domains

Codified

Mathematics: arithmetic, algebra, geometry, calculus, statistics and their applications. → Formalized in textbooks, taught through curricula.

Physics: physical principles, laws, fluid/material/atmospheric dynamics, mechanical/electrical/atomic structures. → Scientific laws documented in textbooks.

Chemistry: chemical composition, structure, properties, processes, interactions, production techniques. → Taught through formal science education.

Biology: plant/animal organisms, cells, functions, interdependencies, interactions with environment. → Academic science with documented knowledge base.

Engineering and Technology: applying engineering science and technology to design and production of goods and services. → Technical standards and specifications from formal education.

Computers and Electronics: circuit boards, processors, hardware, software, applications, programming. → Learned through courses, documentation, and manuals.

Economics and Accounting: economic/accounting principles, financial markets, banking, financial data analysis. → Business theory from degree programs and standards.

Law and Government: laws, legal codes, court procedures, government regulations, executive orders, agency rules. → Explicitly documented rules and regulations.

Medicine and Dentistry: diagnosing/treating injuries and diseases, drug properties and interactions, preventive health care. → Taught through medical school curricula.

English Language: structure/content of English, meaning/spelling of words, rules of composition and grammar. → Grammar and composition taught in formal education.

Tacit

Administration and Management: strategic planning, resource allocation, leadership technique, production methods, coordination of people and resources. → Organizational and leadership knowledge gained through managerial experience.

Customer and Personal Service: customer needs assessment, meeting quality standards for services, evaluation of customer satisfaction. → Interpersonal service skills acquired through practice.

Personnel and Human Resources: recruitment, selection, training, compensation, labor relations, negotiation. → People management learned through experience.

Building and Construction: materials, methods, tools for construction or repair of buildings, highways, roads. → Craft knowledge via hands-on apprenticeship.

Mechanical: machines and tools, their designs, uses, repair, maintenance. → Equipment intuition from hands-on work.

Psychology: human behavior/performance, individual differences in ability/personality/interests, learning/motivation. → Understanding people through interaction and observation.

Public Safety and Security: equipment, policies, procedures, strategies for protection of people, data, property, institutions. → Situational awareness from field experience.

Food Production: techniques/equipment for planting, growing, harvesting food products for consumption. → Craft/culinary skills via apprenticeship.

Therapy and Counseling: diagnosis, treatment, rehabilitation of physical/mental dysfunctions, career counseling. → Clinical intuition built through patient interaction.

Education and Training: curriculum/training design, teaching, instruction for individuals/groups, measurement of training effects. → Pedagogical skill developed through teaching practice.

Work Activities

Codified

Processing Information: compiling, coding, categorizing, calculating, auditing, verifying information or data. → Rule-based handling of structured information.

Documenting/Recording Information: entering, transcribing, recording, storing, maintaining information. → Standardized recording procedures.

Getting Information: observing, receiving, obtaining information from relevant sources. → Structured retrieval from documented sources.

Analyzing Data or Information: identifying underlying principles, reasons, facts by breaking down information. → Formal analytical methods.

Working with Computers: using computers and computer systems to program, write software, set up functions, enter data, process information. → Technical tool use with documented interfaces.

Evaluating Compliance with Standards: using relevant information and judgment to determine whether events or processes comply with laws, regulations, or standards. → Checking against written standards.

Tacit

Judging Qualities of Objects, Services, or People: assessing the value, importance, or quality of things or people. → Experiential quality assessment.

Making Decisions and Solving Problems: analyzing information and evaluating results to choose the best solution. → Context-dependent judgment under ambiguity.

Thinking Creatively: developing, designing, or creating new applications, ideas, relationships, systems, products. → Non-codifiable ideation and innovation.

Coaching and Developing Others: identifying developmental needs, coaching, mentoring, helping others improve knowledge or skills. → Mentoring skills built through experience.

Resolving Conflicts and Negotiating: handling complaints, settling disputes, resolving grievances, negotiating with others. → Interpersonal situation management.

Guiding, Directing, and Motivating Subordinates: providing guidance and direction, including setting performance standards and monitoring performance. → Leadership skills from experience.

Establishing and Maintaining Interpersonal Relationships: developing constructive and cooperative working relationships with others. → Relationship-building through practice.

Cross-Functional Skills

Codified

Programming: writing computer programs for various purposes. → Formal syntax and logic learned from documentation.

Systems Analysis: determining how a system should work and how changes in conditions, operations, and the environment will affect outcomes. → Systematic methods taught in formal training.

Operations Analysis: analyzing needs and product requirements to create a design. → Structured analytical frameworks.

Quality Control Analysis: conducting tests and inspections of products, services, or processes to evaluate quality or performance. → Standardized testing procedures.

Technology Design: generating or adapting equipment and technology to serve user needs. → Technical design principles from formal education.

Operations Monitoring: watching gauges, dials, or other indicators to make sure a machine is working properly. → Following documented procedures and specifications.

Tacit

Social Perceptiveness: being aware of others' reactions and understanding why they react as they do. → Reading people, learned through social experience.

Negotiation: bringing others together and trying to reconcile differences. → Interpersonal skill built through practice.

Persuasion: persuading others to change their minds or behavior. → Influence skills developed through experience.

Coordination: adjusting actions in relation to others' actions. → Situational responsiveness learned through teamwork.

Instructing: teaching others how to do something. → Pedagogical skill from practice.

Service Orientation: actively looking for ways to help people. → Empathy and responsiveness developed through interaction.

Complex Problem Solving: identifying complex problems and reviewing related information to develop and evaluate options and implement solutions. → Solving novel, ill-defined problems in real-world settings.

Judgment and Decision Making: considering the relative costs and benefits of potential actions to choose the most appropriate one. → Weighing trade-offs under uncertainty, learned through experience.

Management of Personnel Resources: motivating, developing, and directing people as they work, identifying the best people for the job. → People leadership developed through managerial practice.

User prompt: codified dimension

```
JOB TITLE: {title}
JOB DESCRIPTION: {description}
```

Read the job description above. For each task described, infer what types of knowledge it draws upon using the framework you learned.

How much does this job rely on CODIFIED knowledge — formal, documented, taught through education? Score 1–10 (1 = not at all, 10 = entirely).

Return ONLY valid JSON:

```
{"codified": ⟨int 1-10⟩, "reasoning": "⟨which tasks draw on codified knowledge⟩"}
```

User prompt: tacit dimension

```
JOB TITLE: {title}
JOB DESCRIPTION: {description}
```

Read the job description above. For each task described, infer what types of knowledge it draws upon using the framework you learned.

How much does this job rely on TACIT knowledge — experience-based, learned through practice and mentorship? Score 1–10 (1 = not at all, 10 = entirely).

Return ONLY valid JSON:

```
{"tacit": ⟨int 1-10⟩, "reasoning": "⟨which tasks draw on tacit knowledge⟩"}
```

Constructing Codified and Tacit Knowledge Quadrants To map from the codified and tacit knowledge scores for raw job titles to the quadrants used in Online Appendix Figure J.1, we first collapse the scores to the SOC code level. We do this by taking an unweighted average of the raw job title codified and tacit ratings that map to a given SOC code.

Given these SOC-level codified and tacit scores, we then take the November 2022 employment weighted mean of each score. We code each SOC occupation according to whether its score in each dimension is above or below the weighted mean. This results in a quadrant of possible classifications for each occupation: high tacit and high codified, high tacit and low codified, low tacit and high codified, or low tacit and low codified.

Results Figure J.1 plots employment over time by age group and codified/tacit quadrant, normalized to November 2022. For early-career workers, high-codified occupations grow more slowly than low-codified ones. For older workers, high-tacit occupations grow more quickly than low-tacit ones.

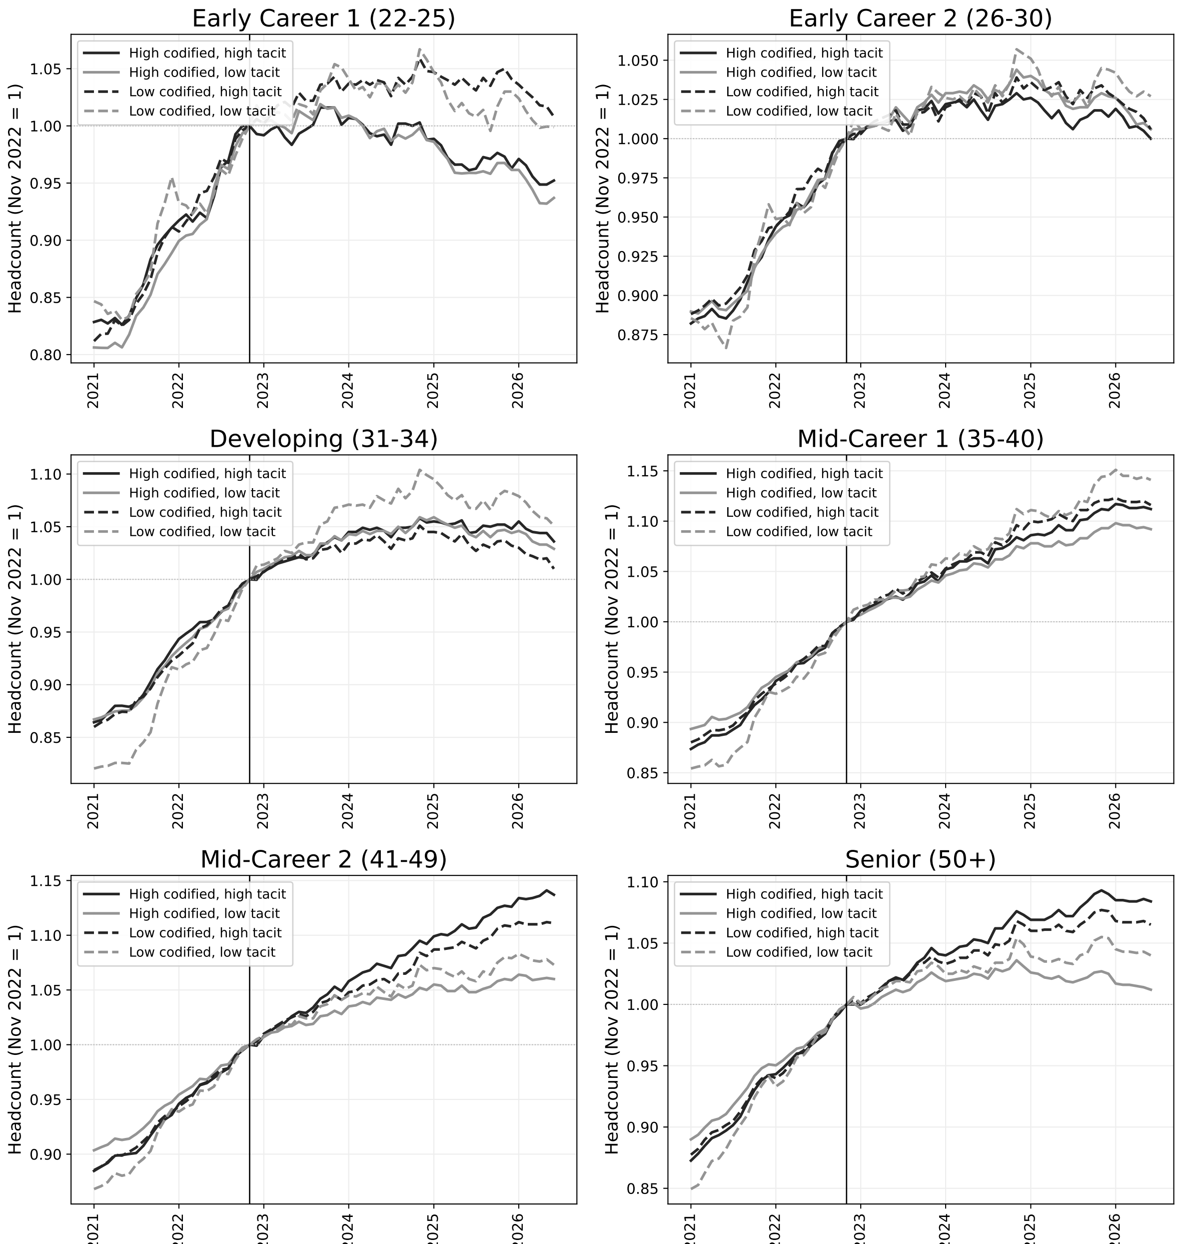

Figure J.1: Employment over time by age group and codified/tacit knowledge quadrant, normalized to 1 in November 2022. Occupations are classified as high or low on each dimension relative to the November 2022 employment-weighted mean of the SOC-level scores.

K Hiring and Separation Flows

We measure hiring and separations for each (age group, exposure quintile) cell. A hire is a new worker–firm match: a worker observed at a firm in month t but not in month $t-1$. A separation is the reverse, a worker observed at a firm in month $t-1$ but not in month t. Both are cumulated over the twelve months ending in the reported month and divided by the cell's headcount twelve months earlier, so that rates are comparable across cells of different size. Employment in a cell also changes as workers cross age-band boundaries and as workers move between occupations in different exposure quintiles; we measure those flows separately and report in Fact 4 that neither accounts for the patterns we document.[44]

L Changes from the August 2025 Version

Relative to the initially circulated version of this paper (August 2025), this draft (i) extends the data through June 2026, (ii) improves the crosswalk merging ADP occupation codes to external exposure measures, and (iii) imputes missing occupation codes from workers' own employment histories (Section A.1). The descriptive facts are qualitatively unchanged. The within-firm Poisson event-study estimates are most sensitive to the pipeline changes. Under our minimal-filter specification, the estimates for workers aged 22–25 as of June 2026 are -6.8 log points ($p = 0.03$) for the most exposed quintile and -11.3 log points ($p = 0.003$) for the fourth quintile. The August 2025 release instead applied stricter sample filters, requiring firms to hire at least 10 workers within the age group in every period and $\sum_t y_{f,q,t} \geq 100$ for each exposure quintile, and reported -11.7 log points ($p = 0.02$) for the most exposed quintile and -10.5 log points ($p = 0.02$) for the fourth quintile. Applying those same filters to the current data gives -5.3 log points ($p = 0.26$) and -12.8 log points ($p = 0.02$), respectively. Holding the specification fixed, the fourth-quintile estimate is therefore stable across vintages, while the estimate for the most exposed quintile attenuates and is no longer statistically significant.

[44] People within a firm who simultaneously switch between an age cell and an occupation cell get coded as both an age change and an occupation change, with the age change measured first.