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

Erik Brynjolfsson* Bharat Chandar† Ruyu Chen‡§¶

November 13, 2025

Abstract

Using high-frequency administrative data from ADP, we document six facts characterizing labor market shifts following the widespread adoption of generative AI. Early-career workers (ages 22-25) in AI-exposed occupations experienced 16% relative employment declines, controlling for firm-level shocks, while employment for experienced workers remained stable. Adjustments occur primarily via employment rather than compensation, with employment changes concentrated in occupations where AI automates rather than augments labor. Results are robust to excluding technology firms and occupations that are remotable. These six facts provide early large-scale evidence consistent with generative AI disproportionately impacting entry-level workers in the American labor market.

*Stanford University and NBER; erikb@stanford.edu

†Stanford University; chandarb@stanford.edu

‡Stanford University; ruyuchen@stanford.edu

§We thank to David Autor, Sarah Bana, Eric Bergman, Nick Bloom, Cody Cook, Chris Forman, Joshua Gans, Basil Halperin, Christina Langer, Fei-Fei Li, Frank Li, Omeed Maghzian, Jiaxin Pei, Daniel Rock, Brad Ross, Phil Trammell, Andrew Wang, and participants at the Stanford Digital Economy Lab workshop 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.

¶**Latest version: https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/**

1 Introduction

The proliferation of generative artificial intelligence (AI) has sparked a global debate about its potential impact on the labor market. This discourse spans utopian predictions of enhanced productivity, dystopian fears of widespread job displacement, and skeptical views that AI will have minimal effects on employment or productivity. Historically, technologies have affected different tasks, occupations, and industries in different ways, replacing work in some, augmenting others, and transforming still others. These heterogeneous effects suggest that there may be "canaries in the coal mine" which are harbingers of more widespread effects of AI.

Recent discussion coincides with rapid improvements in AI capabilities and adoption. From 2023 to 2024, AI systems improved from solving 4.4% of coding problems to 71.7% on SWE-Bench, a widely used benchmark for software engineering (Maslej et al., 2025). AI has also improved in areas like language understanding, subject knowledge, and reasoning. A recent paper found that current systems could match or outperform up to 47 percent of industry professionals on a predefined benchmark of economically valuable tasks (Patwardhan et al., 2025). At the same time, AI systems are increasingly widely adopted. Hartley et al. (2025) find that LLM adoption at work among U.S. survey respondents above age 18 reached 46% by June/July 2025.[1]

Given the better capabilities and widespread adoption, a central concern, amplified in recent headlines, is whether AI is beginning to supplant human labor, particularly for younger, entry-level workers in highly exposed professions like software engineering and customer service.[2] For instance, in May of 2025, Dario Amodei, co-founder and CEO of Anthropic predicted that AI could wipe out roughly 50 percent of all entry-level white-collar jobs within five years (Morris, 2025).

Despite the intensity of this debate, empirical evidence has struggled to keep pace with technological advancement, leaving many fundamental questions unanswered. This paper confronts this empirical gap by leveraging a large-scale, high-frequency administrative dataset from ADP, the largest payroll software provider in the United States. Our sample consists of monthly, individual-level payroll records through September 2025, encompassing millions of workers across tens of thousands of firms. This rich panel structure allows us to track employment dynamics with a high

[1] Similarly, Bick et al. (2024) found that in late 2024, nearly 40% of the U.S. population age 18-64 reported using generative AI, with 23% using it for work weekly and 9% daily.

[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.

degree of granularity, providing a near real-time view of labor market adjustments. By linking this data to established measures of occupational AI exposure and other variables, we can quantify the realized employment changes since the widespread adoption of generative AI.

This paper systematically presents six key facts that emerge from the data, offering an assessment of how the AI revolution is reshaping the American workforce.

Our first key finding is that we uncover substantial declines in employment for early-career workers (ages 22-25) in occupations most exposed to AI, such as software developers and customer service representatives. In contrast, employment trends for more experienced workers in the same occupations, and workers of all ages in less-exposed occupations such as nursing aides, have remained stable or continued to grow.

Our second key fact is that overall employment continues to grow robustly, but employment growth for young workers has been stagnant since late 2022. In jobs less exposed to AI, young workers have experienced comparable employment growth to older workers. In contrast, workers aged 22 to 25 have experienced a 6% decline in employment from late 2022 to September 2025 in the most AI-exposed occupations, compared to a 6-9% increase for older workers. These results suggest that declining employment in AI-exposed jobs drives stagnant overall employment growth for 22- to 25-year-olds.

Our third key fact is that not all uses of AI are associated with declines in employment. In particular, entry-level employment has declined in applications of AI that *automate* work, but not those that most *augment* it. We distinguish between automation and augmentation empirically using estimates of the extent to which observed queries to Claude, the LLM, substitute or complement the tasks in that occupation. While we find employment declines for young workers in occupations where AI primarily automates work, we find employment *growth* in occupations in which AI use is most augmentative. These findings are consistent with automative uses of AI substituting for labor while augmentative uses do not.

Fourth, we find that employment declines for young, AI-exposed workers remain after conditioning on firm-time effects. One class of explanations for our patterns is that they may be driven by industry- or firm-level shocks such as interest rate changes that correlate with sorting patterns by age and measured AI exposure. We test for a class of such confounders by controlling for firm-time effects in an event study regression, absorbing aggregate firm shocks that impact all workers at a firm regardless of AI exposure. For workers aged 22-25, we find a 15 log-point decline in relative employment for the most AI-exposed quintiles compared to the least exposed quintile, a large and statistically significant effect. Estimates for other age groups are much smaller in magnitude and not statistically significant. These findings imply that the employment trends we observe are not driven by differential shocks to firms that employ a disproportionate share of AI-exposed young workers.

Fifth, the labor market adjustments are visible in employment more than in compensation. In contrast to our findings for employment, we find little difference in annual salary trends by age or exposure quintile, suggesting possible wage stickiness. If so, AI may have larger effects on employment than on wages, at least initially, or even that AI may boost wages for as many workers as it hurts.

Sixth, the above facts are largely consistent across sample constructions designed to address various alternative explanations for the core findings. Our results are not driven solely by computer occupations or by occupations susceptible to remote work and outsourcing. We also find that the AI exposure taxonomy did not meaningfully predict employment outcomes for young workers further back in time, before the widespread use of LLMs, including during the unemployment spike driven by the COVID-19 pandemic. The patterns we observe in the data appear most acutely starting in late 2022 and early 2023, around the time of rapid proliferation of generative AI tools.[3] They hold for both occupations with a high share of college graduates and ones with a low college share, suggesting deteriorating education outcomes during COVID-19 do not drive our results. For non-college workers, we find evidence that experience may serve as less of a buffer to labor market disruption, as low college share occupations exhibit divergent employment outcomes by AI exposure up to age 40.

While we explore a variety of alternative explanations, we caution that the facts we document may in part be influenced by factors other than generative AI. Taken as a whole, our results are consistent with the hypothesis that generative AI has begun to affect entry-level employment. We intend to continue to track the data on an ongoing basis to assess whether these trends change in the future.

[3] In particular, OpenAI introduced ChatGPT in November 2022. It was reported to have over 100 million active users by January 2023 and 1.7 billion website visitors in October 2023 (DeVon, 2023).

Why might AI adversely affect exposed entry-level workers more than other age groups? One possibility is that AI disproportionately substitutes for workers using codified knowledge, including both the "book-learning" that forms the core of formal education and the insights in digital company data that can be codified by AI. AI may be less capable of replacing tacit knowledge, the idiosyncratic tips and tricks that accumulate with experience but which are never digitized. [4] As young workers supply relatively more codified knowledge than tacit knowledge, they may face greater task replacement from AI in exposed occupations, leading to greater employment reallocation (Acemoglu and Autor, 2011). Older workers with accumulated tacit knowledge may face less task replacement. These benefits of uncodified knowledge may accrue less to non-college workers in occupations with low returns to experience. In other words, AI may be automating the *codifiable, checkable* tasks that historically justified entry-level headcount, while complementing the judgment-, client-, and process-intensive tasks performed by experienced workers.

Other explanations may contribute. AI may raise the leverage of experienced staff, increasing their effective span of control (Ide, 2025). Reduced hiring may also be the lowest-friction adjustment margin, compounded by inefficient incentives to train entry-level workers who may move firms (Becker, 1994; Garicano, 2025; Garicano and Rayo, 2025). Consequently firms may primarily shrink junior inflows rather than displace incumbents.

This combination of task substitution at the apprentice margin, complementarity at the expert margin, and quantity (not price) adjustment under wage ladders can explain why early-career employment falls while employment of older workers continues to grow. An important direction for research is to further model and test these predictions.

2 Related Literature

A growing body of research has sought to measure the employment effects of AI.[5] This includes influential papers that established methodologies for estimating which occupations and tasks were susceptible to automation (Frey and Osborne, 2017; Brynjolfsson and Mitchell, 2017; Brynjolfsson et al., 2018; Felten et al., 2018, 2019; Webb, 2019; Felten et al., 2021). More recently, work such

[4]Ironically, one of the practical skills more likely to be learned on the job than in university computer science classes may be how to use AI for software development.

[5]There has also been extensive media discusssion of this topic including Thompson (2025); Raman (2025); Roose (2025); The Economist (2025); Frick (2025). See Online Appendix B for further references and discussion.

as Eloundou et al. (2024); Felten et al. (2023); Gmyrek et al. (2023); Handa et al. (2025), and Tomlinson et al. (2025) adapted this approach for Generative AI, forming the basis for exposure metrics used in this analysis. While these studies identify potential disruption, our paper connects these exposure measures to actual employment changes.

Our work broadens insights from studies that find significant effects in more specific settings, such as on online freelance platforms (Hui et al., 2023; Demirci et al., 2025) or within individual firms (Brynjolfsson et al., 2025; Dillon et al., 2025).[6] We measure labor market changes across occupations spanning the US economy.

In this sense our work complements a small but growing list of papers using economy-wide data to measure AI's impact. Recent findings have been varied. Jiang et al. (2025) find AI exposure correlates with longer work hours in the U.S.[7] Hampole et al. (2025) use job postings and LinkedIn profiles from Revelio labs from 2011 to 2023 to find limited employment impacts overall, with growing labor demand at firms offsetting relative declines in demand for exposed occupations. Chandar (2025b) uses data from the CPS to compare employment changes in more and less AI-exposed professions, finding little differential trend overall but noting difficulty with measuring changes for young workers given limited effective sample size.[8] Johnston and Makridis (2025) find employment increases in state-industry pairs more exposed to AI using Quarterly Census of Employment and Wages (QCEW) data.[9] These prior papers use data lacking either sufficient granularity or immediacy to reliably study employment changes by AI exposure and age (O'Brien, 2025).[10] In contrast, this paper uses large-scale, close to real-time data to take a step towards resolving the ongoing debate on the employment effects of AI on young workers.

After we first publicly shared our results, work by Hosseini and Lichtinger (2025) and Klein Teeselink (2025) similarly found declines in AI-exposed entry-level employment using LinkedIn data from the

[6]See also Noy and Zhang (2023); Peng et al. (2023); Dell'Acqua et al. (2023).

[7]See also Acemoglu et al. (2022); Bonney et al. (2024); Bick et al. (2024); Hartley et al. (2025); Frank et al. (2025); Chen et al. (2025).

[8]See Dominski and Lee (2025); Gimbel et al. (2025); Eckhardt and Goldschlag (2025), which also use CPS data.

[9]Johnston and Makridis (2025) measure state-industry exposure by taking an average of Eloundou et al. (2024)'s occupational exposure weighted by state-industry employment. Industry-level labor market changes may be distinct from the occupation-level changes studied in this paper if firms make capital investments or become more productive in ways that increases overall labor demand (Hampole et al., 2025).

[10]As a comparison, the CPS surveyed between 44,000 and 51,000 employed individuals in total across all age groups in each month since 2021. Between 10,000 and 12,000 of these observations were in the outgoing rotation group and included earnings records. The data in our main analysis sample includes between 250,000 and 350,000 employed individuals in each month *just between the ages of 22 and 25*, all with earnings records.

US and UK. In contrast, the most recent version of Humlum and Vestergaard (2025) found minimal effects on entry-level earnings or hours worked in Denmark.[11]

3 Data Description

3.1 Payroll Data

This study uses data from ADP, the largest payroll processing firm in America. The company provides payroll services for firms employing over 25 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.[12]

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 September 2025.

ADP observes job titles for about 70% of workers in its system. We exclude workers without a recorded job title. There are over 7,000 standardized job titles, including "Search engineer 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.

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

While the ADP data include millions of workers in each month, the distribution of firms using

[11] An important direction for future work is to reconcile differences among these recent studies, assessing the extent they arise from differences in labor market institutions, measurement of AI adoption, statistical methodology, or other factors.

[12] 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 Center for Disease Control and Prevention.

ADP services does not exactly match the distribution of firms across the broader US economy. Further details on differences in firm composition can be found in Cajner et al. (2018) and ADP Research (2025).[13]

3.2 Occupational AI Exposure

We use two 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.

The second approach we take uses generative AI usage data from the Anthropic Economic Index (Handa et al., 2025). 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.[14]

We use a 2010 SOC code to 2018 SOC code crosswalk from the BLS to merge the exposure measures to the payroll data. Table A1 shows example occupations for each AI exposure measure.

3.3 Other Data

To compare employment changes for teleworkable versus non-teleworkable occupations, we use data from Dingel and Neiman (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) data as a comparison for our main findings.

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

[14] 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 A2 reproduces Table 1 from Handa et al. (2025) and shows more details about the automation and augmentation measures.

4 Results

4.1 Fact 1: Employment for young workers has declined in AI-exposed occupations

Consider software engineers and customer service agents, two occupations frequently considered to be highly exposed to generative AI tools. Media attention has raised the specter of widespread employment disruption for young software engineers (Thompson, 2025; Raman, 2025; Allen, 2025; Horowitch, 2025).

Figure 1 shows employment changes by age group for these occupations, normalized to 1 in October 2022. While the types of work and workers in each of these occupations differs in many ways, both occupations present a similar pattern: employment for the youngest workers declines considerably after 2022, while employment for other age groups continues to grow. By September 2025, employment for software developers aged 22-25 declined nearly 20% compared to its peak in late 2022. Figure A1 shows that a similar pattern holds for computer occupations and service clerks more generally.

Figure A2 shows four other professions as case studies, spanning varying levels of AI exposure according to Eloundou et al. (2024). Marketing and sales managers, in the fourth quintile of AI exposure, show a decline in employment for young workers much like the case of software and customer service, albeit with smaller magnitudes. Front-line production and operations supervisors, in quintile 3, show an increase in employment for young workers, though the growth in employment is smaller than the increase for workers over age 35.

In contrast, the trends for occupations that Eloundou et al. (2024) rated as less exposed do not fit the pattern of the more exposed occupations. Stock clerks and order fillers, in quintile 2, show no obvious difference by age. Strikingly, the series for health aides, comprising nursing aides, psychiatric aides, and home health aides, show a quite different trend from software or customer service: employment for young workers has been growing *faster* than for older workers.

Figure 2 shows these patterns hold more generally across professions. The top left plot shows a divergence in employment outcomes for more and less exposed occupations for workers aged 22-25, with more exposed occupations experiencing declining employment. For older age groups, we find much less marked differences in employment growth across AI exposure quintiles.[15]

4.2 Fact 2: Though overall employment continues to grow, employment growth for young workers has been stagnant

Overall employment in the ADP data remains robust, coinciding with a low national unemployment rate in the post-pandemic period. However, Figure A4 suggests some leveling off in employment growth for young workers relative to other age groups, consistent with recent discussion of a worsening job market for entry-level workers (Chen, 2025; Federal Reserve Bank of New York, 2025).

Figure A5 offers insight into how these trends relate to AI exposure. For each age group, employment growth from late 2022 to September 2025 was 5-13% for the lowest three AI exposure quintiles, with no clear ordering in employment growth by age. In contrast, for the highest two exposure quintiles employment for 22-25 year olds declined by 6% between late 2022 and September 2025, while employment for workers aged 35-49 grew by over 8%. These results show that declining employment in AI-exposed jobs is driving tepid overall emplyoment growth for workers between the ages of 22 and 25.

While these findings suggest divergent employment outcomes by AI exposure for young workers, we caution the trends observed in these first two facts could be driven by other changes in the US economy. Our subsequent facts evaluate the robustness of the results to alternative analyses.

4.3 Fact 3: Entry-level employment has declined in applications of AI that *automate* work, with muted changes for *augmentation*

AI exposure 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). 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

[15] Figure A3 shows that declining employment for young workers spans a range of occupations. Close to 70% of occupations in the first exposure quintile see rising early career employment between October 2022 and September 2025, compared to less than half of occupations in the fifth quintile.

of whether the usage of AI in an occupation is primarily a substitute or complement for labor. Table A1 shows example occupations that are in the highest and lowest exposure category for each measure.

Panel A of Figure 3 shows employment changes by overall prevalence of related Claude queries for 22-25 year-olds.[16] The patterns match the findings using the Eloundou et al. (2024) measures 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 augmentation shares have *not* experienced a similar pattern. Employment changes for young workers are not ordered by augmentation exposure, as the fifth quintile has among the fastest employment growth. The findings are consistent with automative uses of AI substituting for labor while augmentative uses do not.[17]

4.4 Fact 4: Employment declines for young, AI-exposed workers remain after conditioning on firm-time shocks

While our results so far are consistent with the hypothesis that generative AI is causing a decline in entry-level employment, there are plausible alternative explanations. One class of explanations is that our patterns are 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 a class of such confounders by controlling for a rich set of fixed effects. For each age group, we estimate the Poisson regression

$$\log(E[y_{f,q,t}]) = \sum_{q'\neq 1} \sum_{j\neq -1} \gamma_{q',j} 1\{t=j\} 1\{q'=q\} + \alpha_{f,q} + \beta_{f,t} + \epsilon_{f,q,t} \tag{4.1}$$

[16]See Figures A6 through A8 for other age groups.

[17]Occupations in the first two augmentation quintiles have very low Claude usage (0.01% and 0.09% of conversations for the average occupation, respectively), with a high share of conversations classified as neither automative nor augmentative. In contrast, occupations in the third through fifth quintiles average 0.47%, 0.39%, and 0.33% of Claude conversations. For the automation measure, overall Claude usage increases on average with the automation share, with the lowest exposure group averaging 0.05% of conversations and the highest group averaging 0.73%. Figures A9 and A10 show that automation and augmentation results are similar when dropping occupations with low overall Claude usage.

f indexes firms, q indexes Eloundou et al. (2024) exposure quintiles, and t indexes months, with $t = -1$ corresponding to October 2022. The outcome variable $y_{f,q,t}$ is employment in f, q, t. Equation 4.1 is a Poisson event study regression controlling for firm-quintile effects, $\alpha_{f,q}$, and firm-time effects, $\beta_{f,t}$. The firm-time effects absorb aggregate firm shocks that impact each exposure quintile equally. The firm-quintile effects adjust for baseline differences in hiring across quintiles within the firm. The coefficients of interest, $\gamma_{q,t}$, measure differential changes in employment growth across quintiles after accounting for firm-time effects and firm-quintile effects.[18]

We run this regression separately for each age group. For each regression, we restrict to firms that hire at least 10 workers within the age group in every period of the sample. Further, $\sum_t y_{f,q,t}$ must equal at least 100 for each q, meaning that the firm must at least employ on average about 2 workers from each exposure quintile across months in the sample.[19] Standard errors are clustered by firm.

Results are in Figure 4, which plots the $\gamma_{q,t}$ coefficients for each age group. For workers aged 22-25, estimates for higher quintiles are large and statistically significant, with a 15 log point decline in relative employment comparable in magnitude to the estimates in the raw data in Figure 2. Estimates for other age groups are generally much smaller in magnitude and not statistically significant. These findings imply that the employment trends we observe are not driven by differential shocks to firms that employ a disproportionate share of AI-exposed young workers.

One alternative confounder that would not be controlled for with firm-time effects is that even conditional on the firm workers with high AI exposure were excessively hired after the COVID-19 pandemic, leading to a subsequent contraction in their hiring. To assess such alternatives we consider various other robustness checks in Section 4.6, such as removing computer occupations and conditioning on whether the occupation is amenable to work from home.

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

[19] Results are not sensitive to these restrictions, though there must be at least one non-zero value in each firm-month and each firm-quintile for observations to not get dropped in the Poisson regression.

4.5 Fact 5: Labor market adjustments are visible in employment more than compensation

In addition to employment we observe workers' annual base compensation. We use this information to test for labor market adjustment along the compensation margin.[20] Salary data are deflated to 2017 dollars using the PCE index.[21]

Results for various occupations are in Figure A11. The findings indicate less divergence in compensation compared to employment across more and less exposed occupations. Figure 5 shows results by age and Eloundou et al. (2024)-based exposure quintile. We find little difference in compensation trends by age or exposure quintile.

Prior work by Autor and Thompson (2025) notes that technology that replaces inexpert tasks may reduce occupational employment but increase occupational wages; technology that replaces expert tasks may do the opposite. The sign of the wage effect depends on the overall share of tasks displaced as well as the whether these tasks are expert or inexpert. The limited changes we find for wages suggest that these effects may be offsetting, at least in the short run. Alternatively, the results could be explained by wage stickiness in the short run, consistent with recent evidence from Davis and Krolikowski (2025).

4.6 Fact 6: Findings are largely consistent under alternative sample constructions

We test the robustness of these results to alternative sample constructions and robustness checks.

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.[22] Figure A12 shows employment changes by age and exposure quintile after excluding computer occupations, corresponding to 2010 SOC codes that start with 15-1. Figure A13 shows results when excluding firms in information technology or computer systems design (NAICS codes

[20] Total compensation may additionally include bonuses, overtime pay, commissions, equity, tips, and other items. These may have a greater impact on overall compensation in certain professions and age groups than others.

[21] In contrast to the series for employment, results for compensation end in August 2025, the most recently available month for the PCE index.

[22] Under the Tax Cuts and Jobs Act, amendments to Internal Revenue Code §174 enacted in 2022 also disallowed companies from immediately deducting R&D expenditures, including software development costs.

51 and 5415). Results are quite similar, consistent with the above case studies showing employment changes across various occupations. Results with firm-time fixed effects in Figure 4 further show that our findings are robust to firm- or industry-level shocks that impact general hiring trends. These results indicate that our findings are not specific to technology roles.

Remote Work Figures A14 and A15 show results for occupations amenable to remote work (telework) and those that are not, according to Dingel and Neiman (2020).[23] 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.[24]

Longer Sample Figure A16 shows results when extending the balanced sample of firms to 2018. This reduces the sample size and makes the data somewhat noisier. Nonetheless, the trends remain largely ordered by exposure in the post-GPT era, whereas this is not the case before 2022. A concern is that for the Eloundou et al. (2024) measures the most exposed quintile had slower employment growth starting around 2020. This is not the case for the Anthropic exposure measures, shown in Figures A17, A18, and A19. For these measures the most exposed groups have comparable employment growth throughout the period before generative AI, with divergent trends afterwards.

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).[25] In Figure A20 we show trends for occupations in which greater than 70% of workers have a college degree according to the 2017 American Community Survey (ACS).[26] In Figure A21 we show trends for occupations in which fewer than 30% of workers

[23] 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 Figure A14 we pool together the two lowest AI exposure quintiles into one group. In Figure A15 we pool together the two highest AI exposure quintiles.

[24] Non-teleworkable occupations with high AI exposure include bank tellers, travel agents, and tax preparers.

[25] 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.

[26] 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 Figure A20.

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. 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 40 showing separation in employment trends by AI exposure. These findings suggest that deteriorating education outcomes cannot fully explain our main results. They also suggest that for non-college workers, experience may serve as less of a buffer to labor market disruption than for college workers.

Occupational Interest Rate Exposure While estimates with firm-time effects control for overall hiring trends within firms, another possibility is that even within firms occupations with high AI exposure have high interest rate exposure, leading to contraction in hiring. We use occupational interest rate exposure data from Zens et al. (2020) to test for the correlation between interest rate exposure and AI exposure. Figure A22 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. Figures A23 and A24 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.

Other Robustness Checks Figures A25 and A26 show results separately for men and women. The results are similar, suggesting that diverging prospects for men and women are not driving our findings. Figure A27 shows that results are similar when we do not take a balanced sample of firms. Figure A28 shows similar results when including part-time and temporary workers.[27]

Comparison to CPS Data A useful benchmark for our findings is to compare them to estimates from the monthly Current Population Survey (CPS). In Online Appendix C we perform similar analyses in the CPS, finding high levels of noise consistent with small sample sizes in fine age-

[27]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 explaning our findings.

occupation cells. Future work should study employment trends in other large-scale data sources such as the American Community Survey (ACS) or data from Revelio Labs as in Hosseini and Lichtinger (2025) and Klein Teeselink (2025).

5 Conclusion

We document six facts about the recent labor market effects of artificial intelligence.

- We find substantial employment declines for early-career workers in occupations most exposed to AI, such as software development and customer support.
- While economy-wide employment continues to grow, employment growth for young workers has been stagnant.
- Entry-level employment has declined in applications of AI that *automate* work, with muted effects for those that *augment* it.
- After conditioning on firm-time effects, young workers experienced a 16% relative employment decline in the most exposed occupations.
- Labor market adjustments are more visible in employment than in compensation.
- These patterns hold across various alternative analyses.

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 significantly.

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).

Tracking employment trends on an ongoing basis will help determine if the adjustment to AI follows a similar pattern. We will continue monitoring these outcomes to assess whether the trends documented in the paper accelerate in the future. Future work would benefit from better firm-level AI adoption data, which would provide sharper variation for estimating plausible causal effects of AI on employment.

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. 1

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2022): "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics*, 40, S293–S340, publisher: The University of Chicago Press. 7

ADP Research (2025): "ADP National Employment Report," . 3.1

Allen, Mike, J. V. (2025): "Behind the Curtain: Top AI CEO foresees white-collar bloodbath," *Axios*. 4.1, B

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. 5

Autor, D. and N. Thompson (2025): "Expertise," Working Paper, NBER. 4.5

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. A22

Bacon, A. (2025): "Nvidia's Jensen Huang says AI could lead to job losses 'if the world runs out of ideas' | CNN Business," *CNN*. B

Becker, G. S. (1994): *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education, Third Edition*, The University of Chicago Press, backup Publisher: National Bureau of Economic Research Type: Book. 1

Bick, A., A. Blandin, and D. J. Deming (2024): "The Rapid Adoption of Generative AI," Working Paper, National Bureau of Economic Research. 1, 7

Bonney, K., C. Breaux, C. Buffington, E. Dinlersoz, L. Foster, N. Goldschlag, J. Haltiwanger, Z. Kroff, and K. Savage (2024): "The impact of AI on the workforce: Tasks versus jobs?" *Economics Letters*, 244, 111971. 7

Bowen, T. (2025): "Graduating into a Slowdown: Class of 2025 Meets a Frozen Job Market," . 29

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. 5

Brynjolfsson, E. (2022): "The Turing Trap: The Promise and Peril of Human-Like Artificial Intelligence," *Daedalus*, 151, 272–287. 4.3

Brynjolfsson, E., D. Li, and L. Raymond (2025): "Generative AI at Work*," *The Quarterly Journal of Economics*, 140, 889–942. 2

Brynjolfsson, E. and T. Mitchell (2017): "What Can Machine Learning Do? Workforce Implications," *Science*, 358, 1530–1534. 2

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. 2

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. 5

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. 3.1, 13

Chandar, B. (2025a): "Shifts in the Composition of College Workers: Implications for US Inequality and Labor Demand," Tech. rep., Social Science Research Network. 25

——— (2025b): "Tracking Employment in AI-Exposed Jobs," Tech. rep., Social Science Research Network. 2, C, 29

Chen, J. and J. Roth (2024): "Logs with Zeros? Some Problems and Solutions*," *The Quarterly Journal of Economics*, 139, 891–936. 18

Chen, J. L. a. T.-P. (2025): “Young Graduates Are Facing an Employment Crisis,” *WSJ*, section: Economy. 4.2

Chen, W. X., S. Srinivasan, and S. Zakerinia (2025): “Displacement or Complementarity? The Labor Market Impact of Generative AI,” . 7

Davis, S. J. and P. M. Krolikowski (2025): “Sticky Wages on the Layoff Margin,” *American Economic Review*, 115, 491–524. 4.5

Dell’Acqua, F., E. McFowland III, E. R. Mollick, H. Lifshitz-Assaf, K. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani (2023): “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality,” SSRN Scholarly Paper, Social Science Research Network, Rochester, NY. 6

Demirci, O., J. Hannane, and X. Zhu (2025): “Who is AI replacing? The impact of generative AI on online freelancing platforms,” *Management Science*. 2

DeVon, C. (2023): “On ChatGPT’s one-year anniversary, it has more than 1.7 billion users—here’s what it may do next,” . 3

Dillon, E., S. Jaffe, S. Peng, and A. Cambon (2025): “Early Impacts of M365 Copilot,” Tech. Rep. MSR-TR-2025-18, Microsoft. 2

Dingel, J. I. and B. Neiman (2020): “How many jobs can be done at home?” *Journal of Public Economics*, 189, 104235. 3.3, 4.6, A14, A15

Dominski, J. and Y. S. Lee (2025): “Advancing AI Capabilities and Evolving Labor Outcomes,” Tech. rep., arXiv, arXiv:2507.08244 [econ]. 8, C

Doshay, H. and A. Bantock (2025): “The SignalFire State of Tech Talent Report - 2025,” . 29

Eckhardt, S. and N. Goldschlag (2025): “AI and Jobs: The Final Word (Until the Next One),” *Economic Innovation Group*. 8, B, C

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. 2, 9, 3.2, 4.1, 4.3, 4.4, 4.5, 4.6, 2, 4, 5, A2, A3, A12, A13, A14, A15, A16, A20, A21, A22, A23, A24, A25, A26, A27, A28, A32, A1

Ettenheim, L. E. a. K. B. |. G. b. R. (2025): "AI Is Wrecking an Already Fragile Job Market for College Graduates," *WSJ*, section: Tech. 28

Federal Reserve Bank of New York (2025): "The Labor Market for Recent College Graduates," . 4.2

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. 2

——— (2023): "How will Language Modelers like ChatGPT Affect Occupations and Industries?" Tech. rep., arXiv, arXiv:2303.01157 [econ]. 2

Felten, E. W., M. Raj, and R. Seamans (2018): "A Method to Link Advances in Artificial Intelligence to Occupational Abilities," *AEA Papers and Proceedings*, 108, 54–57. 2

——— (2019): "The Occupational Impact of Artificial Intelligence: Labor, Skills, and Polarization," SSRN Scholarly Paper, Social Science Research Network, Rochester, NY. 2

Frank, M. R., Y.-Y. Ahn, and E. Moro (2025): "AI exposure predicts unemployment risk: A new approach to technology-driven job loss," *PNAS Nexus*, 4, pgaf107. 7

Frey, C. B. and M. A. Osborne (2017): "The future of employment: How susceptible are jobs to computerisation?" *Technological Forecasting and Social Change*, 114, 254–280. 2

Frick, W. (2025): "AI Is Everywhere But the Jobs Data," *Bloomberg.com*. 5, B

Garicano, L. (2025): "The AI Becker problem," . 1

Garicano, L. and L. Rayo (2025): "DP20634 Training in the Age of AI: A Theory of Apprenticeship Viability," CEPR Discussion Paper 20634, Paris & London. 1

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," . 8

Gmyrek, P., J. Berg, and D. Bescond (2023): "Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality," SSRN Scholarly Paper, Social Science Research Network, Rochester, NY. 2

Hampole, M., D. Papanikolaou, L. D. Schmidt, and B. Seegmiller (2025): "Artificial Intelligence and the Labor Market," Working Paper, National Bureau of Economic Research. 2, 9

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]. 2, 3.2, 14, 4.3, 3, A6, A7, A8, A9, A10, A17, A18, A19, A1, A2

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. 1, 7

Hoover, A. (2025): "The AI coding apocalypse," *Business Insider*. 28

Horowitch, R. (2025): "The Computer-Science Bubble Is Bursting," *The Atlantic*, section: Economy. 4.1, 28

Hosseini, S. M. and G. Lichtinger (2025): "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data," . 2, 4.6

Hui, X., O. Reshef, and L. Zhou (2023): "The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market," SSRN Scholarly Paper, Rochester, NY. 2

Humlum, A. and E. Vestergaard (2025): "Large Language Models, Small Labor Market Effects," Working Paper, National Bureau of Economic Research. 2

Ide, E. (2025): "Automation, AI, and the Intergenerational Transmission of Knowledge," *arXiv preprint arXiv:2507.16078*. 1

Jamali, L. (2025): "Microsoft to cut up to 9,000 jobs as it invests in AI," *BBC*. B

Jiang, W., J. Park, R. Xiao, and S. Zhang (2025): "AI and the Extended Workday: Productivity, Contracting Efficiency, and Distribution of Rents," SSRN Scholarly Paper, Social Science Research Network, Rochester, NY. 2

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. 2, 9

Klein Teeselink, B. (2025): "Generative AI and Labor Market Outcomes: Evidence from the United Kingdom," . 2, 4.6

Kochhar, R. (2023): "Which U.S. Workers Are More Exposed to AI on Their Jobs?" *Pew Research Center*. 27

Kuhfeld, M. and K. Lewis (2025): "5 years after COVID-19 hit: Test data converge on math gains, stalled reading recovery," *Brookings*. 4.6

Lenny Rachitsky (2025): "State of the product job market in 2025," *Lenny's Newsletter*, accessed: 2025-05-30. 29

Lim, S., D. Strauss, J. Burn-Murdoch, and C. Murray (2025): "Is AI killing graduate jobs?" *Financial Times*. B, C, 29

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*. 1

Milmo, D. and L. Almeida (2025): "'Workforce crisis': key takeaways for graduates battling AI in the jobs market," *The Guardian*. 28

Morris, C. (2025): "Anthropic CEO warns AI could eliminate half of all entry-level white-collar jobs," *Fortune*. 1

Noy, S. and W. Zhang (2023): "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, 381, 187–192, publisher: American Association for the Advancement of Science. 6

O'Brien, C. (2025): "A viral chart on recent graduate unemployment is misleading," *Agglomerations.* 2, C

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.* 1

Peck, E. (2025): "AI is keeping recent college grads out of work," *Axios.* 28

Peng, S., E. Kalliamvakou, P. Cihon, and M. Demirer (2023): "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot," Tech. rep., arXiv, arXiv:2302.06590 [cs]. 6

Raman, A. (2025): "Opinion | I'm a LinkedIn Executive. I See the Bottom Rung of the Career Ladder Breaking." *The New York Times.* 5, 4.1

Roose, K. (2025): "For Some Recent Graduates, the A.I. Job Apocalypse May Already Be Here," *The New York Times.* 5

Sherman, N. (2025): "Amazon boss says AI will replace jobs at tech giant," *BBC News.* B

Simon, L. K. (2025): "Is AI responsible for the rise in entry-level unemployment?" *Revelio Labs.* 29

Smith, N. (2025): "Stop pretending you know what AI does to the economy," . B

The Economist (2025): "Why AI hasn't taken your job," *The Economist.* 5, B

Thompson, D. (2025): "Something Alarming Is Happening to the Job Market," *The Atlantic*, section: Economy. 5, 4.1, B

Tomlinson, K., S. Jaffe, W. Wang, S. Counts, and S. Suri (2025): "Working with AI: Measuring the Occupational Implications of Generative AI," Tech. rep., arXiv, arXiv:2507.07935 [cs]. 2

Webb, M. (2019): "The Impact of Artificial Intelligence on the Labor Market," SSRN Scholarly Paper, Social Science Research Network, Rochester, NY. 2

Wu, T. (2025): "Opinion | A 'White-Collar Blood Bath' Doesn't Have to Be Our Fate," *The New York Times*. 28

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. 4.6, A22, A23, A24

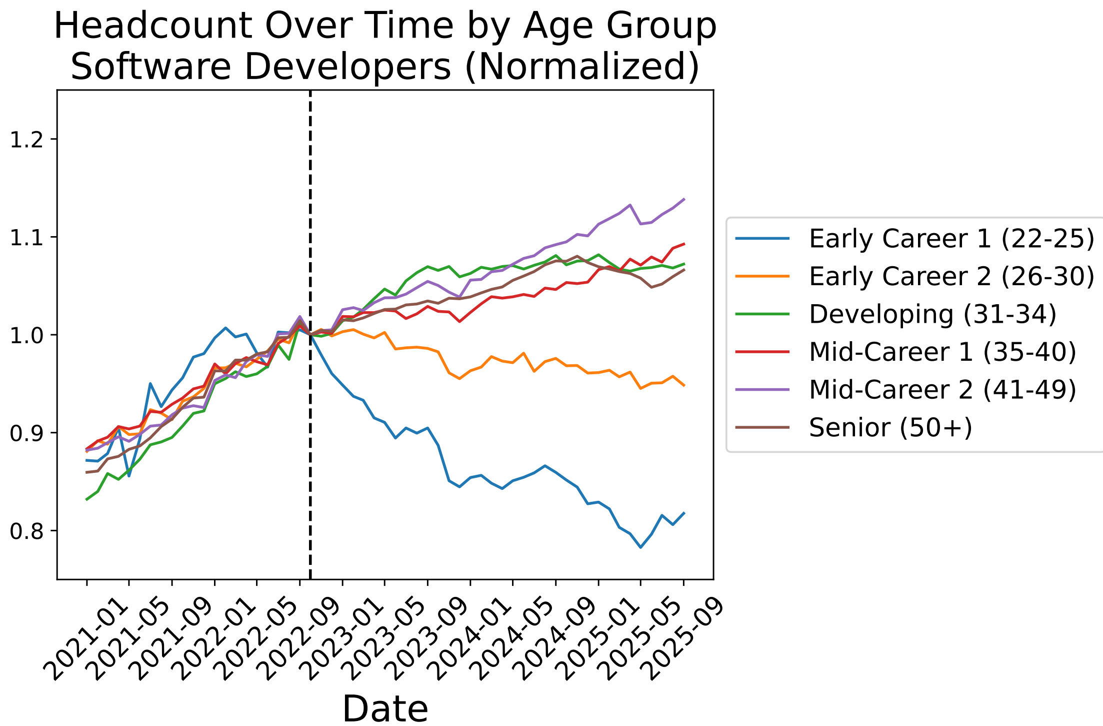

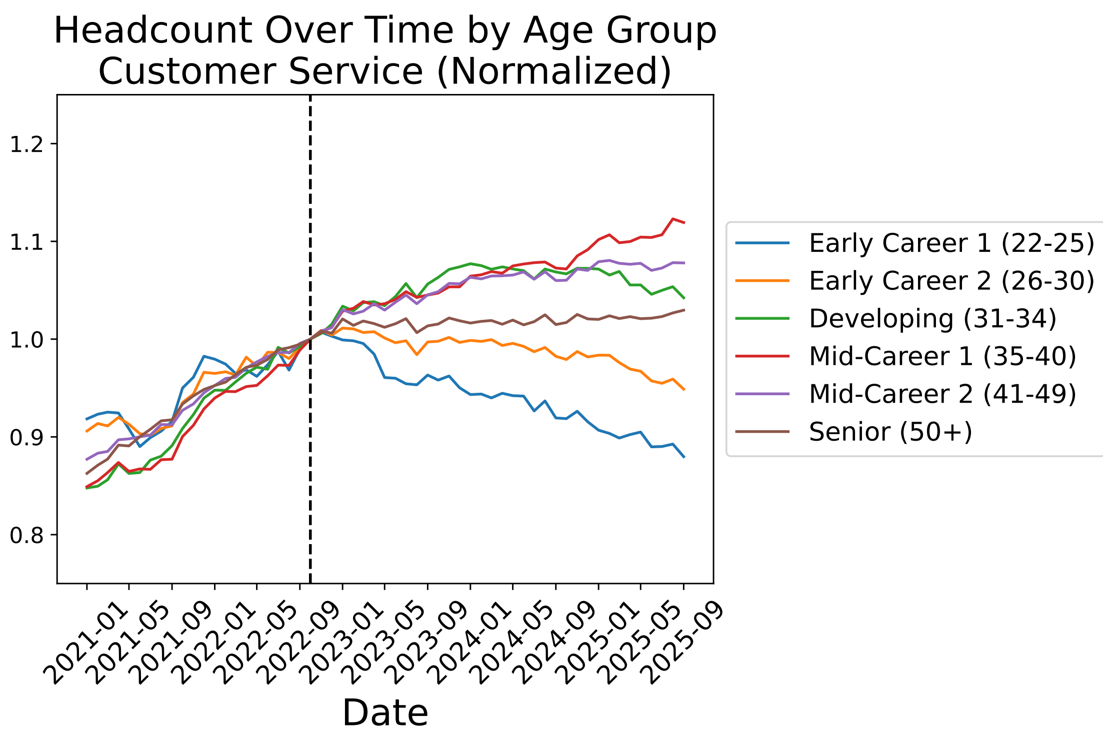

Figure 1: Employment changes for software developers and customer service agents by age, normalized to 1 in October 2022.

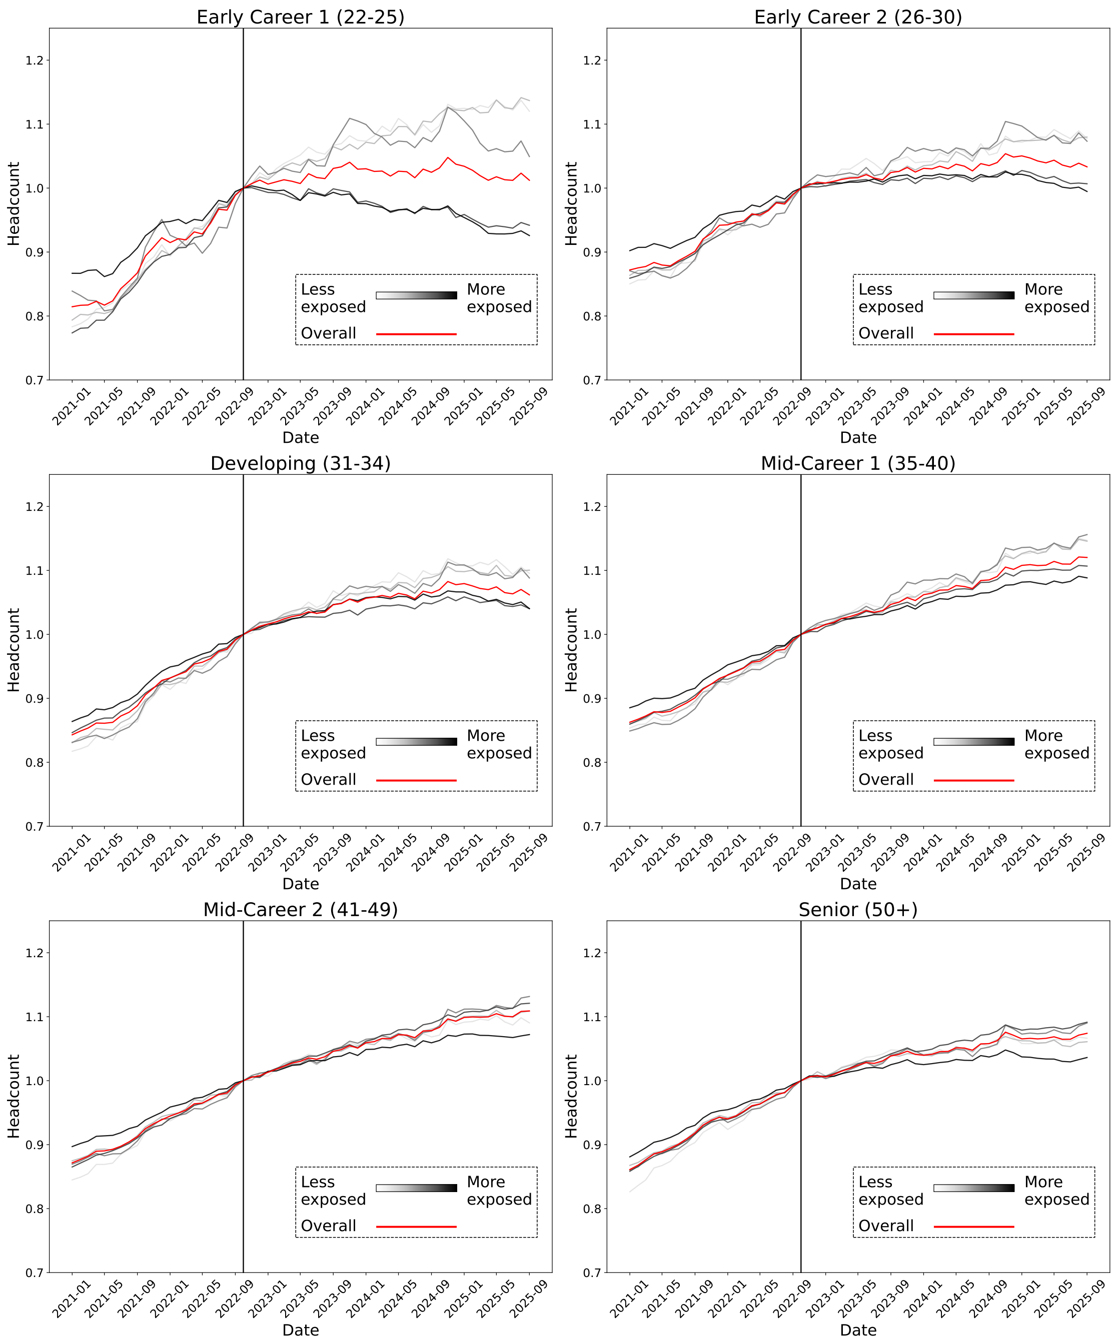

Figure 2: GPT-4 β. Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Exposure quintiles are defined based on the GPT-4 β measures. Darker lines are more exposed quintiles. The red line shows the overall trend pooling across quintiles.

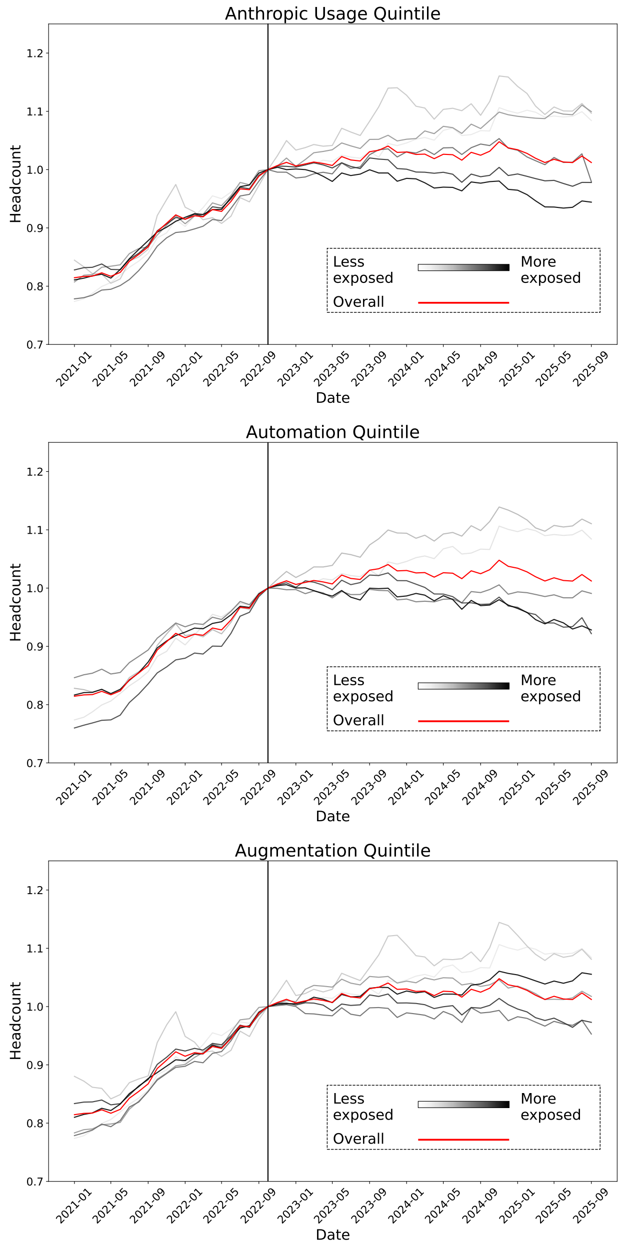

Figure 3: *Automation and Augmentation.* Employment changes by age and exposure quintile using measures from Handa et al. (2025). Panel A: *Overall usage.* Exposure quintiles are defined based on the share of queries to Claude that relate to tasks associated with an occupation. Occupations whose associated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

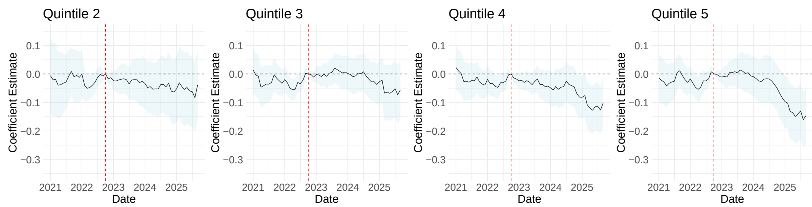

(a) Age 22-25

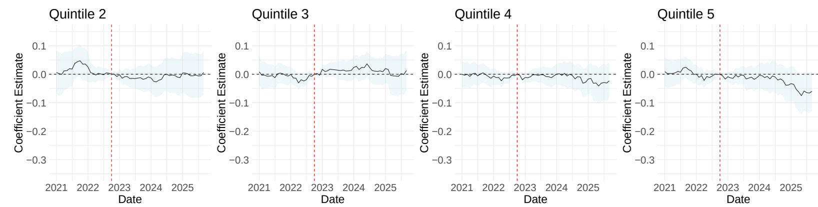

(b) Age 26-30

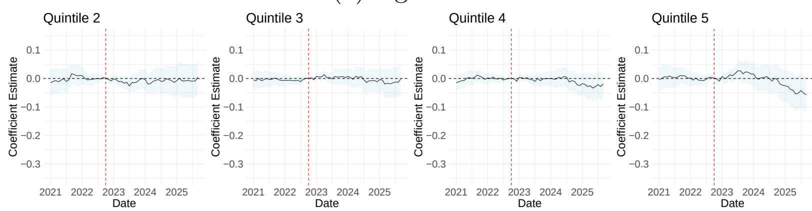

(c) Age 31-34

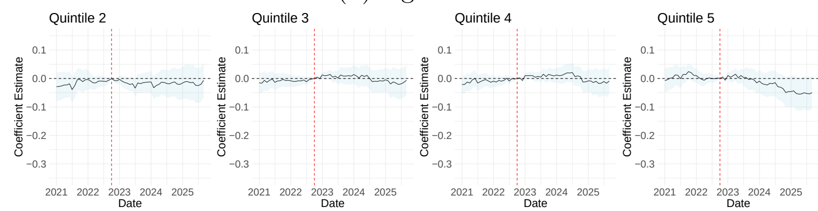

(d) Age 35-40

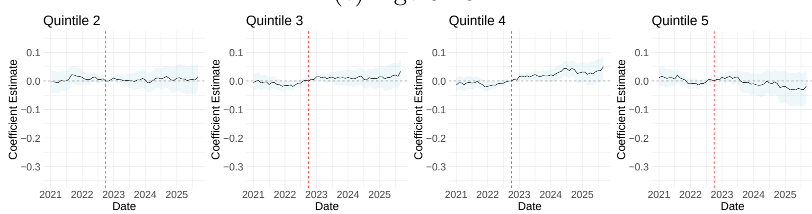

(e) Age 41-49

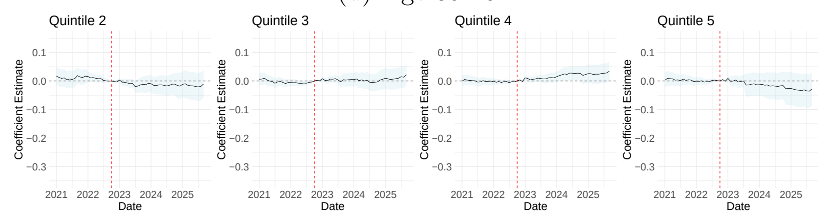

(f) Age 50+

Figure 4: Poisson regression event study estimates for employment changes by age and AI exposure. 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 following Equation 4.1. Shaded regions are 95% confidence intervals. Standard errors are clustered by firm.

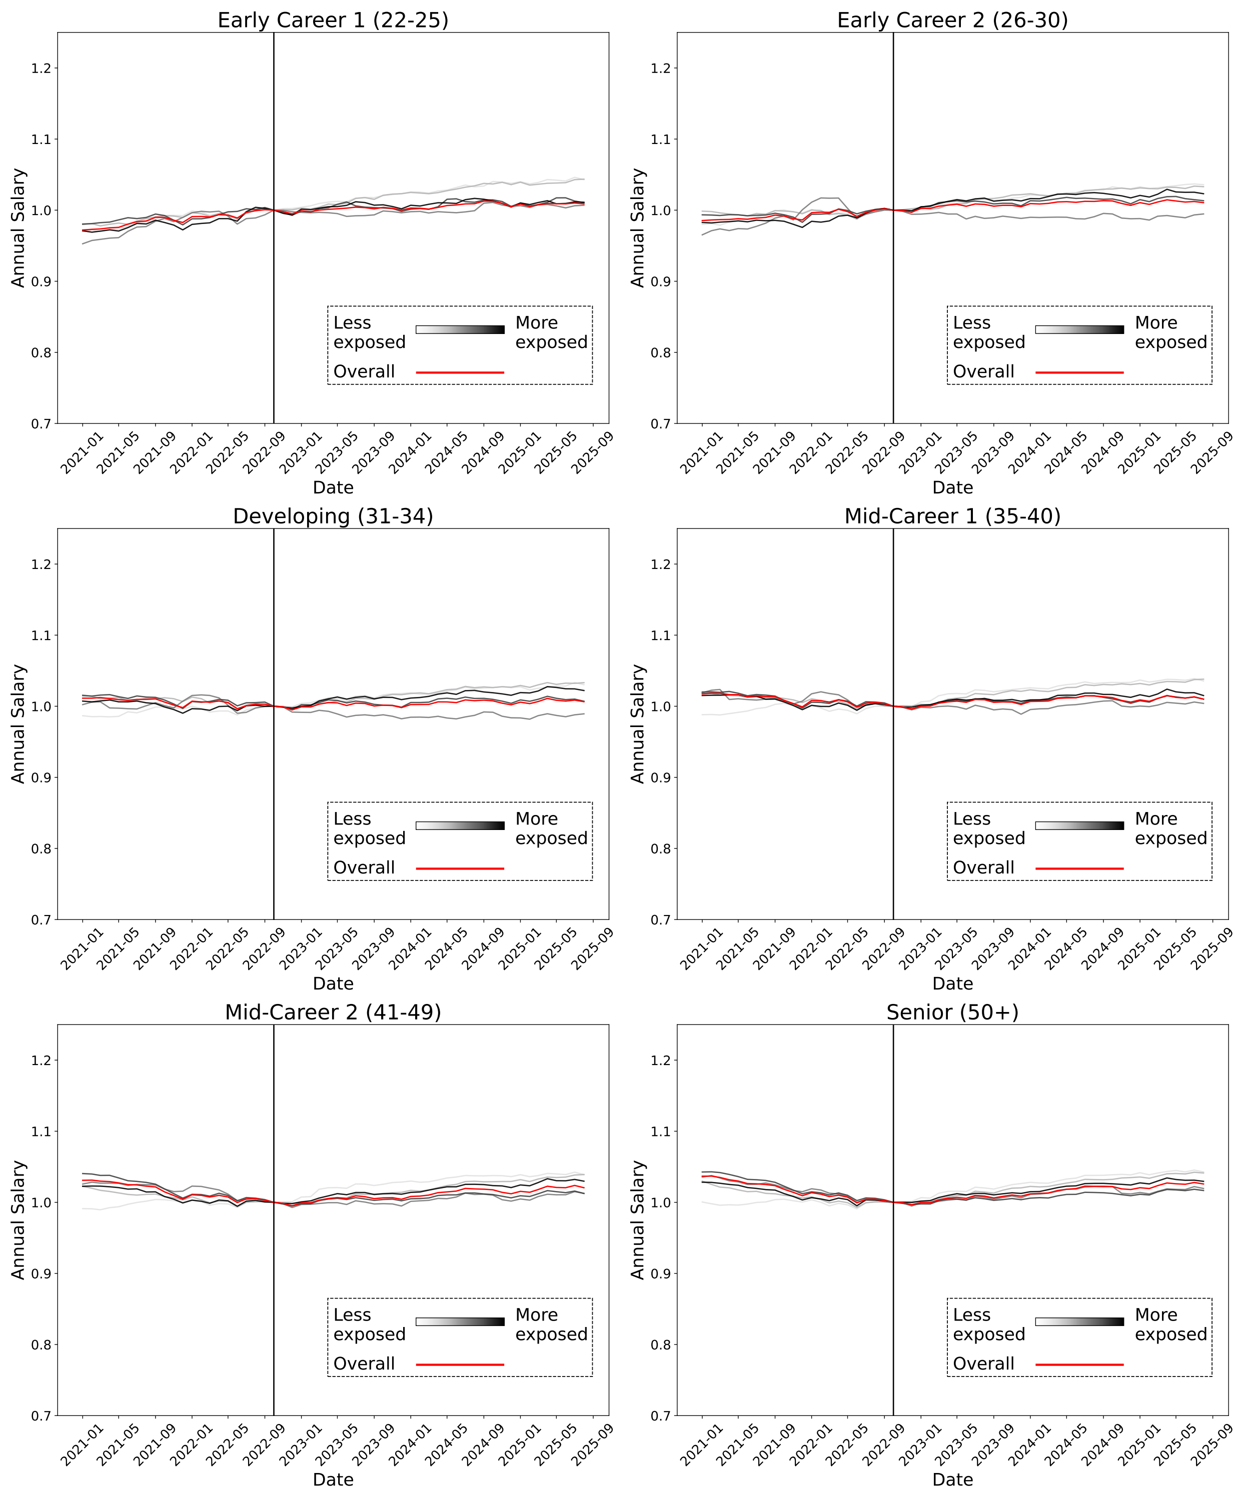

Figure 5: *Annual base compensation.* Changes in annual base compensation by age and exposure quintile. Exposure quintiles are defined based on the GPT-4 β measures from Eloundou et al. (2024). Darker lines are more exposed quintiles. The red line shows the overall trend pooling across quintiles. Annual base compensation is deflated to 2017 dollars using the PCE deflator.

Online Appendix

A Additional Figures and Tables

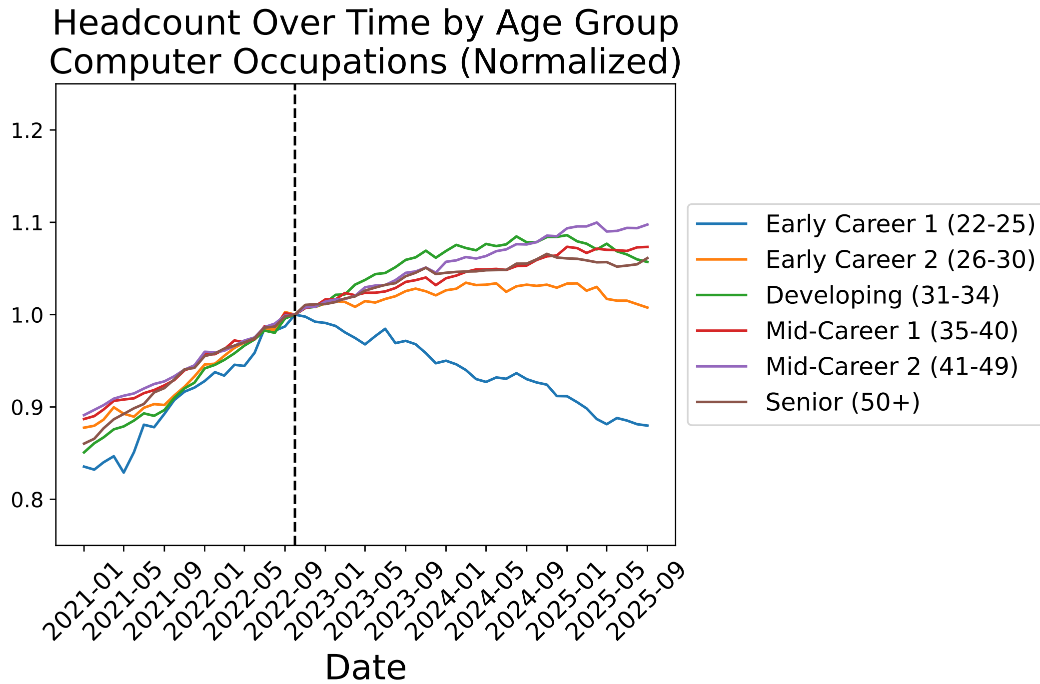

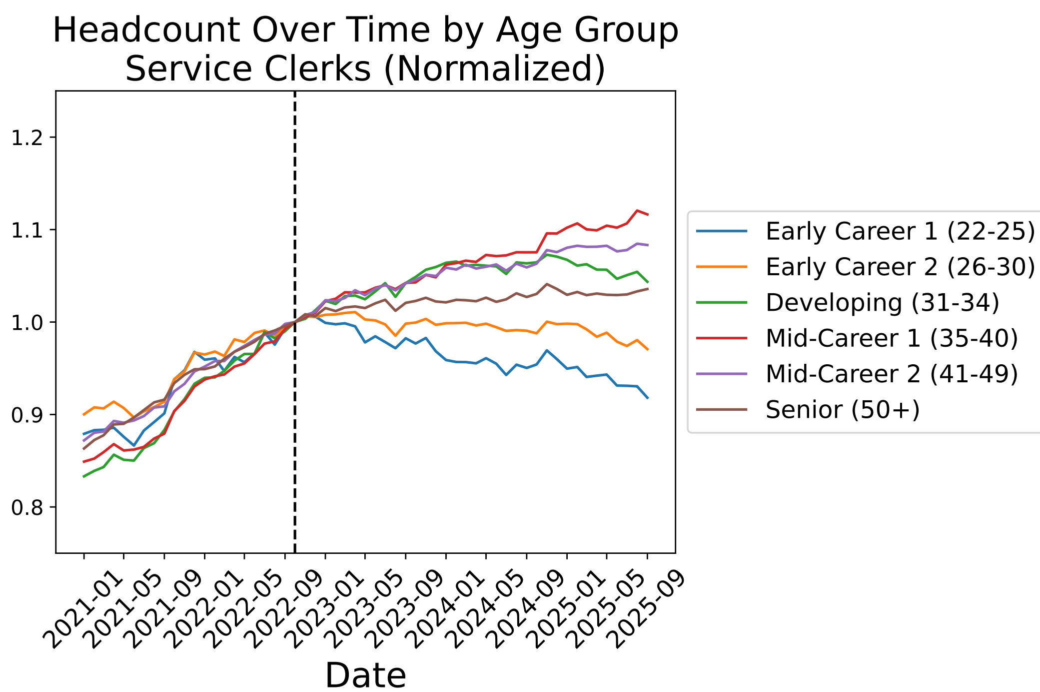

Figure A1: Employment changes for computer occupations (2010 SOC codes starting with 15-1) and service clerks (starting with 43-4), normalized to 1 in October 2022.

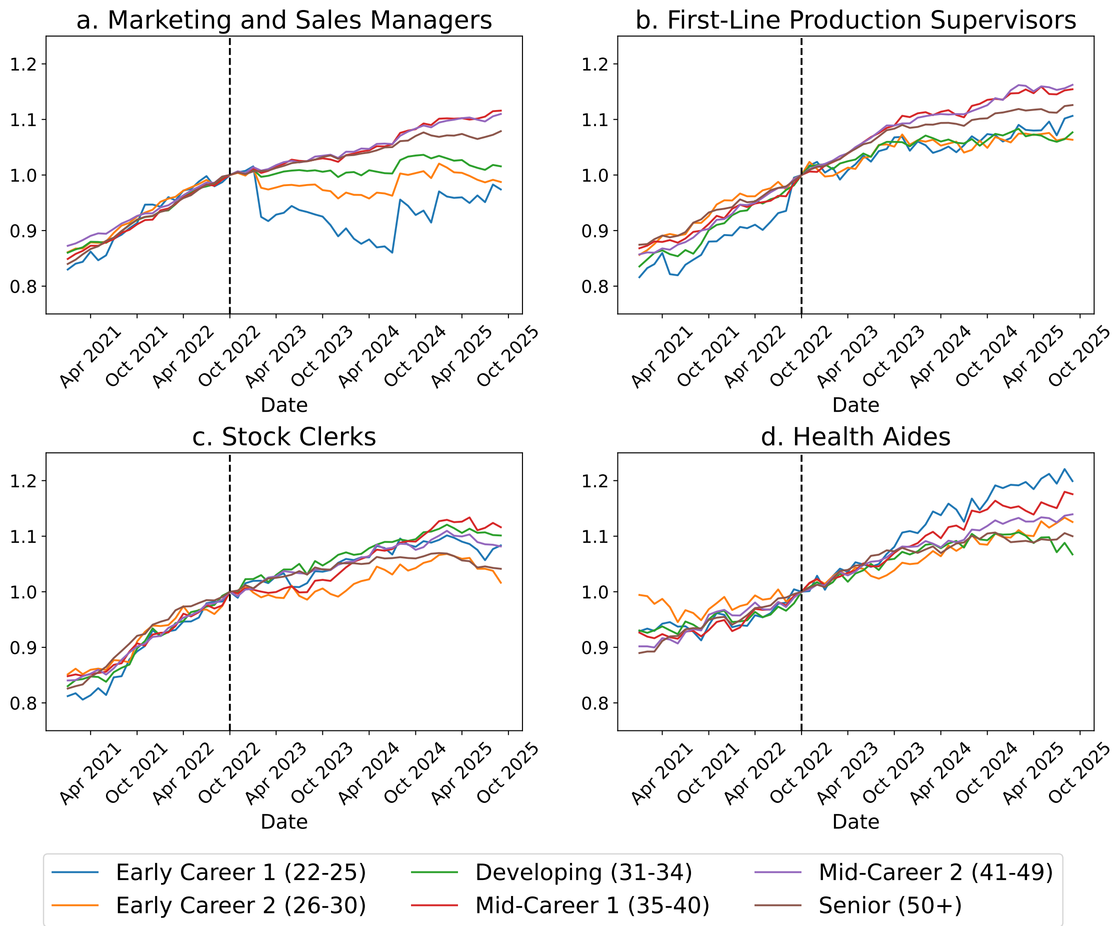

Figure A2: Employment changes for marketing and sales managers (exposure quintile 4), first-line production supervisors (exposure quintile 3), stock clerks and order fillers (exposure quintile 2), and health aides (exposure quintile 1), normalized to 1 in October 2022. Exposure quintiles are defined based on the Eloundou et al. (2024) GPT-4 β measure.

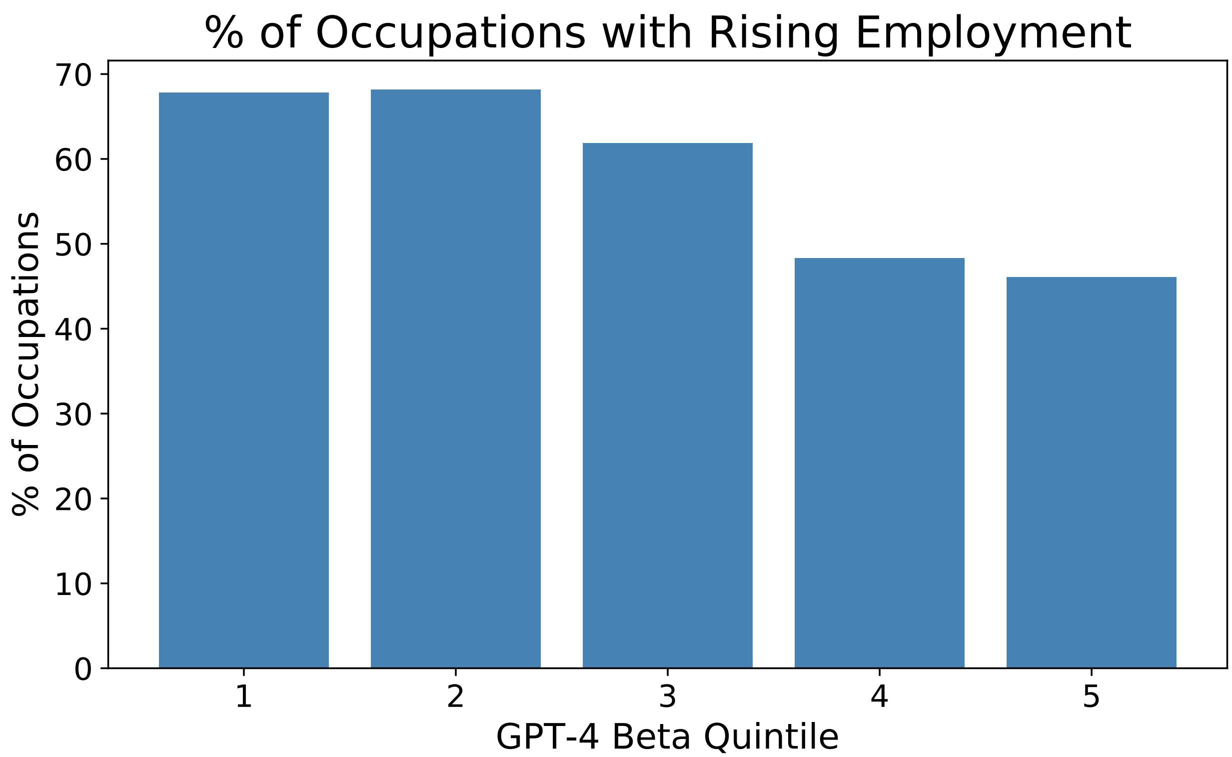

Figure A3: Percent of occupations with an increase in employment for 22-25 year old workers between October 2022 and September 2025. Exposure quintiles are defined based on the Eloundou et al. (2024) GPT-4 β measure.

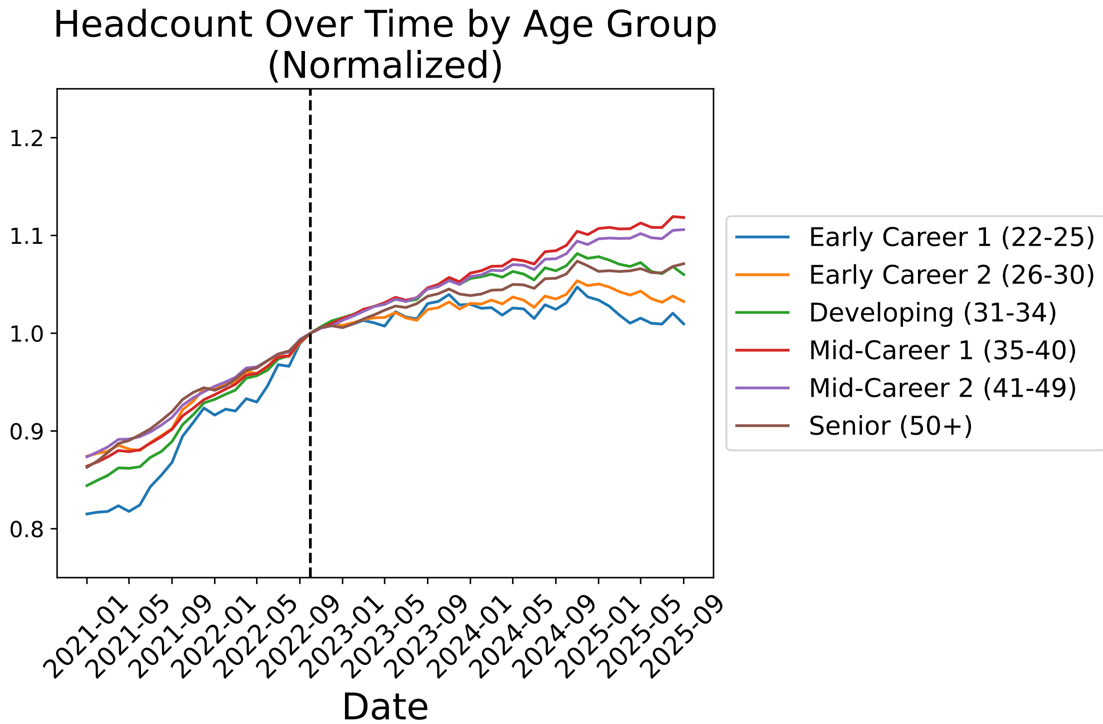

Figure A4: Employment changes by age. Including all occupations.

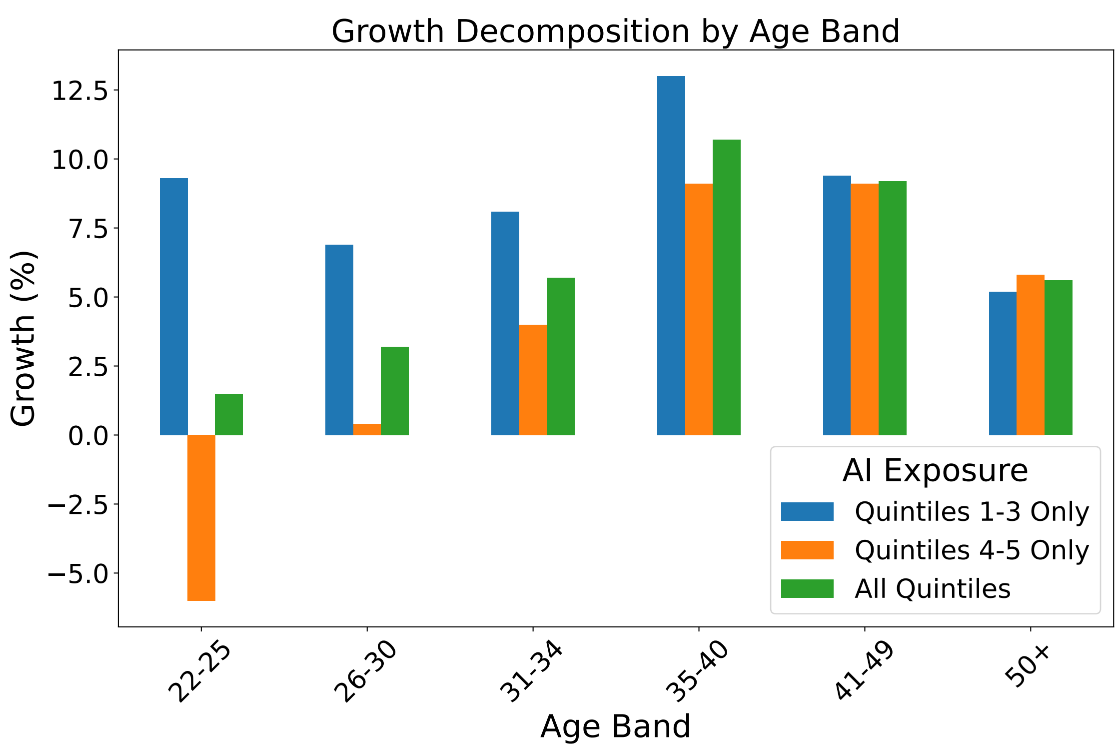

Figure A5: Growth in employment between October 2022 and September 2025 by age and GPT-4 β-based AI exposure group.

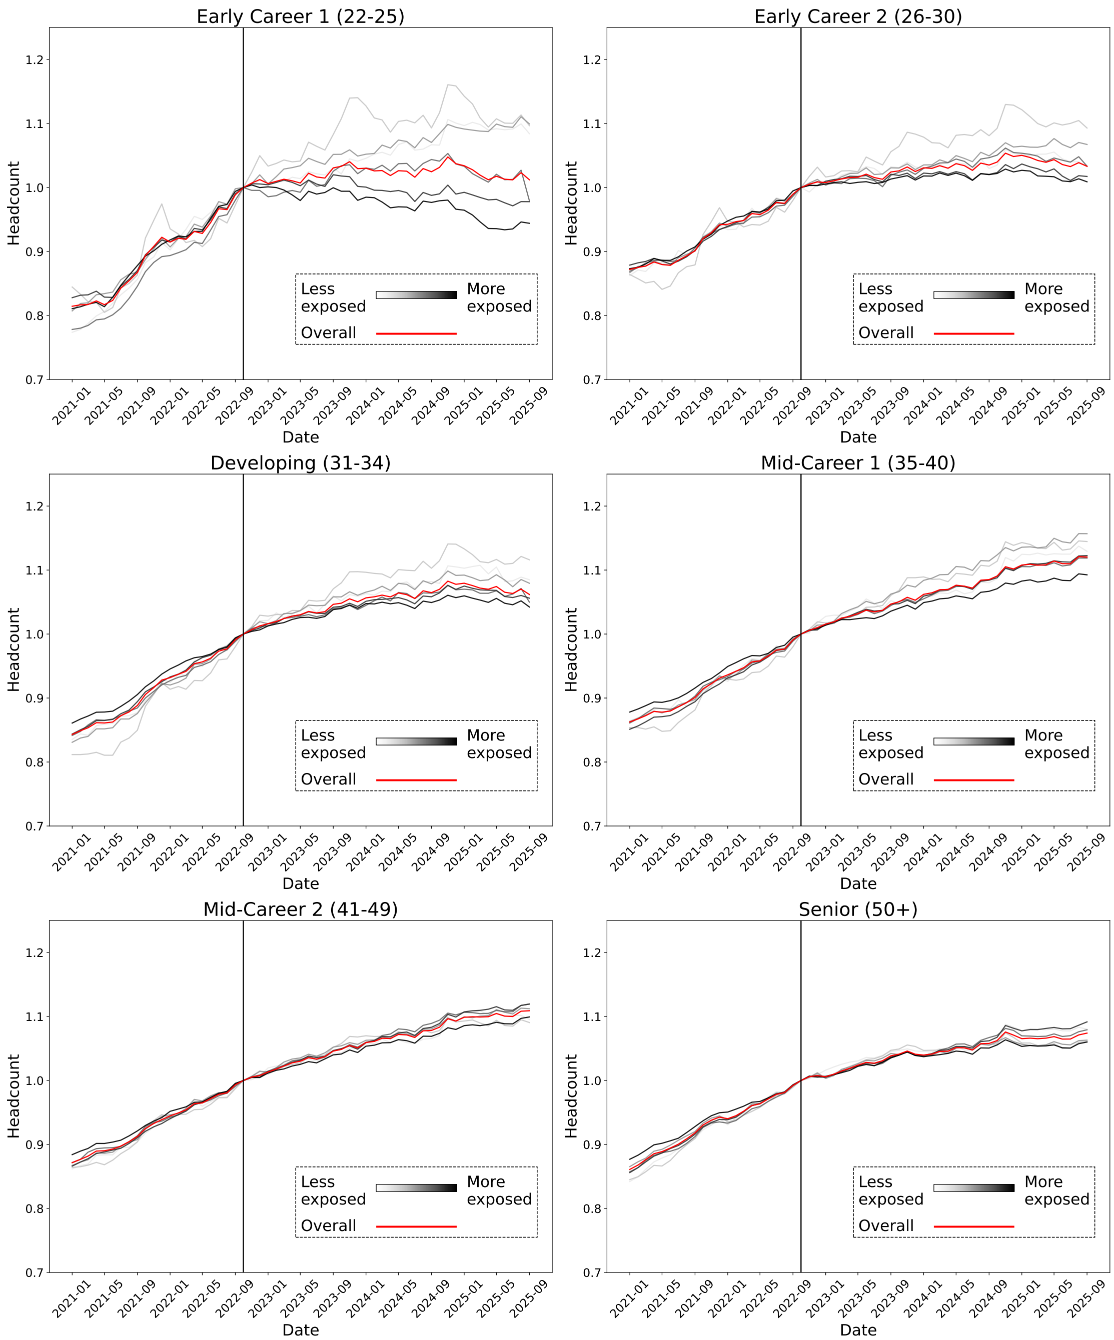

Figure A6: *Overall Claude usage.* Employment changes by age and exposure quintile using Claude usage data from Handa et al. (2025). Exposure quintiles are defined based on the share of queries to Claude that relate to tasks associated with an occupation. Darker lines are more exposed quintiles. The 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.

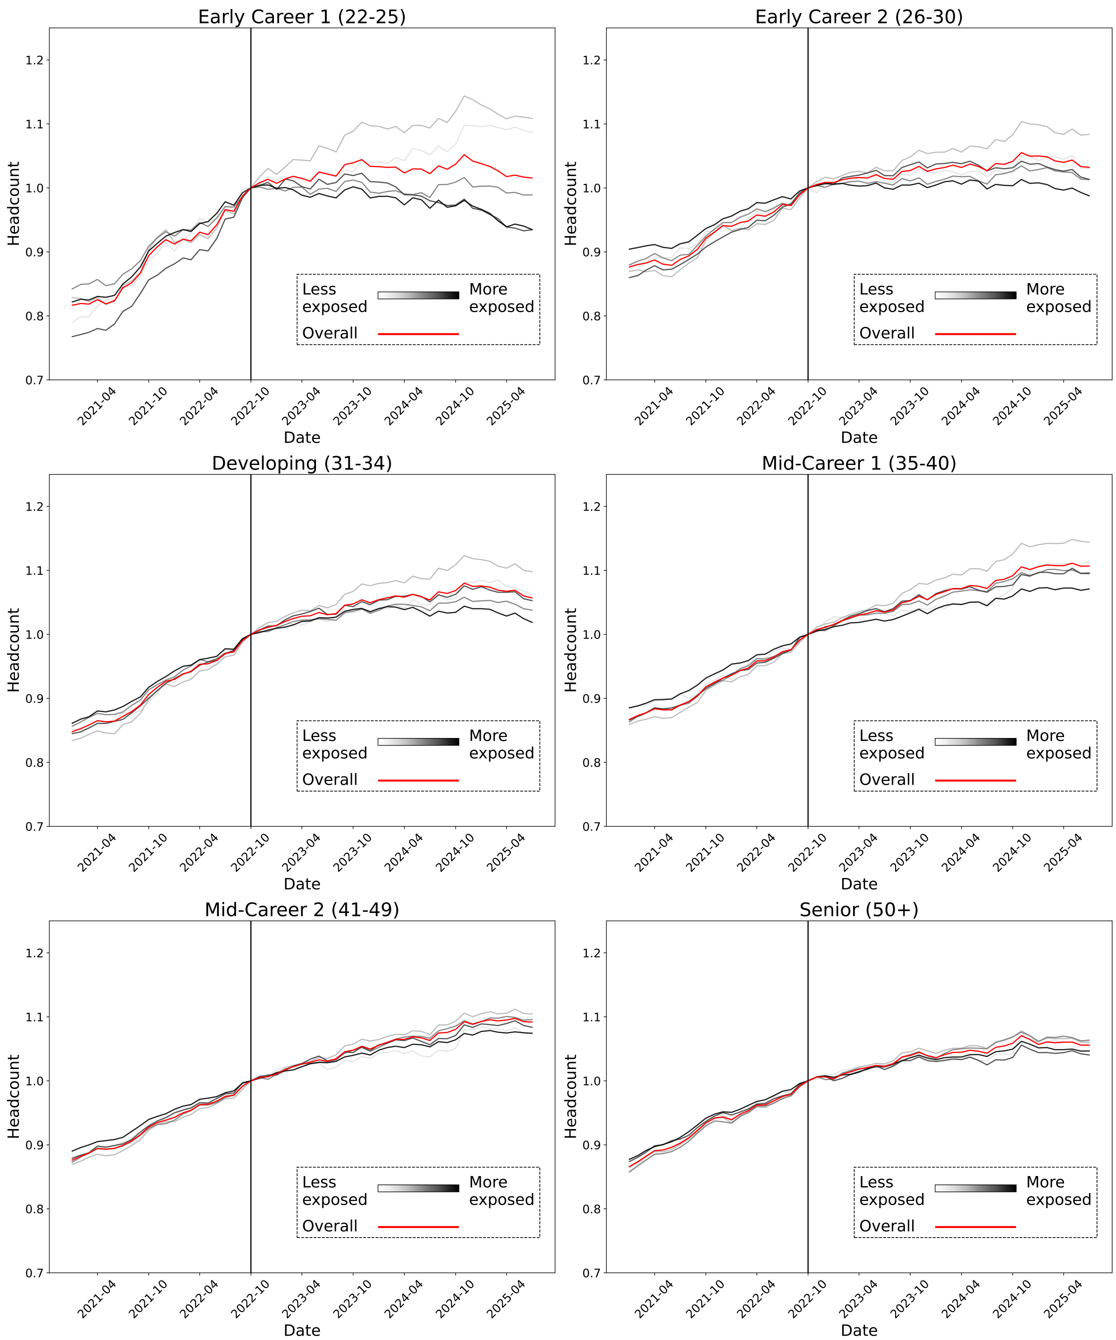

Figure A7: *Automation.* Employment changes by age and automation level using Claude usage data from Handa et al. (2025). Automation levels are defined based on the share of queries related to an occupation that are classified by Claude as automative in nature. Darker lines are more automative. The red line shows the overall trend pooling across automation levels. 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. Note that greater than 20% of occupations above the minimum query threshold have an estimated automation share of 0. All occupations in the first and second quintile are consequently grouped together in level 1. The remaining quintiles are coded as 2, 3 and 4.

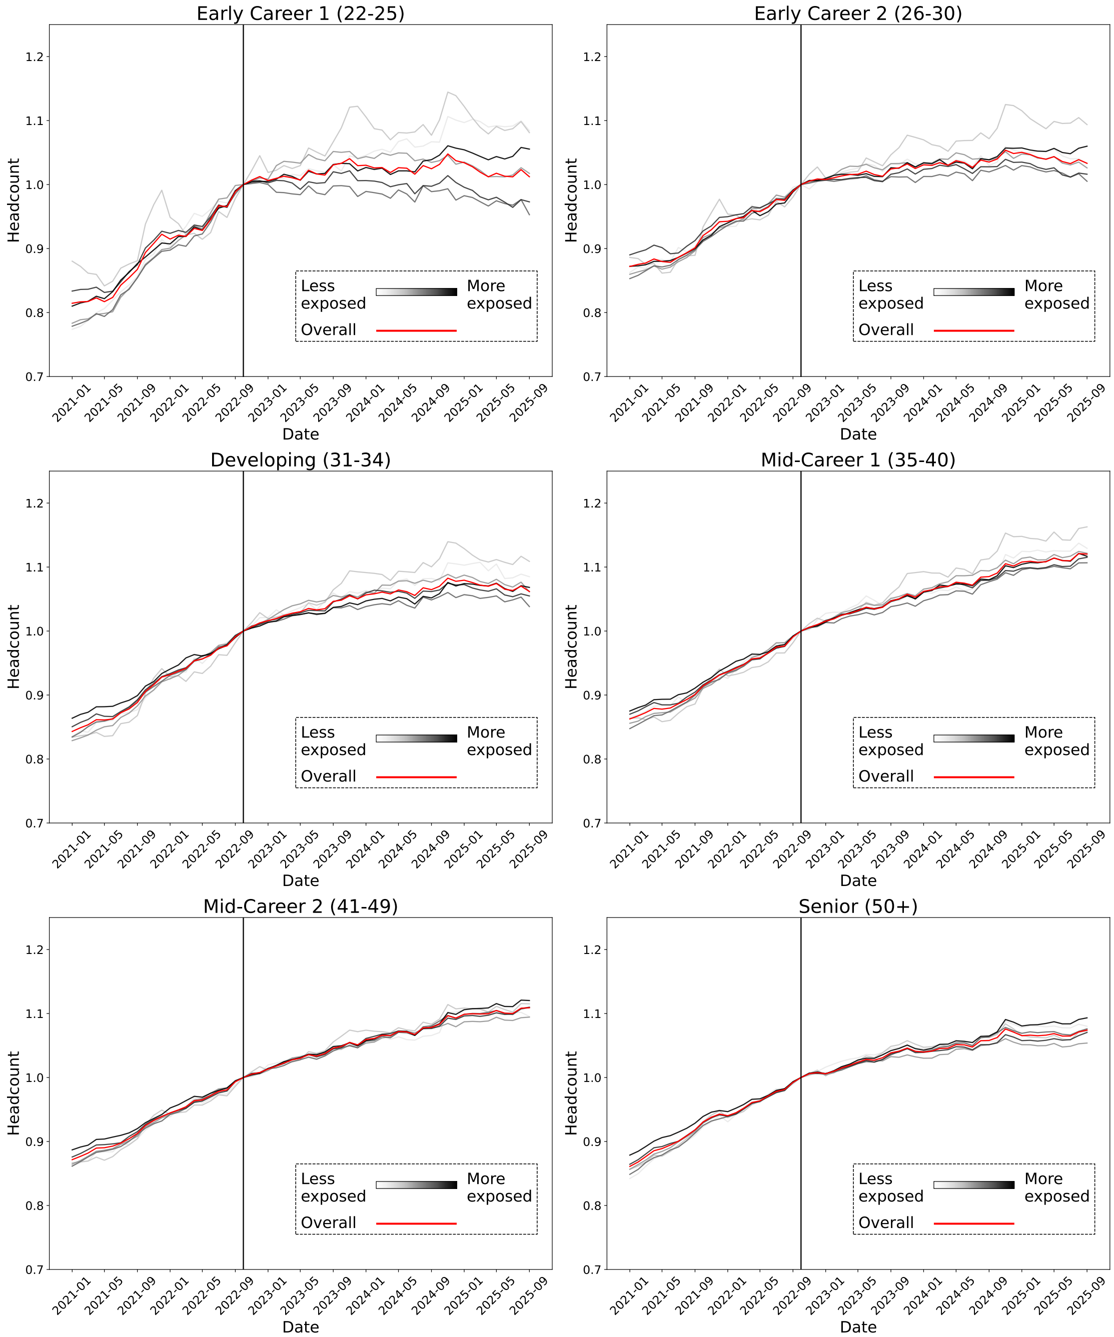

Figure A8: *Augmentation.* Employment changes by age and augmentation quintile using Claude usage data from Handa et al. (2025). Augmentation quintiles are defined based on the share of queries related to an occupation that are classified by Claude as augmentative. Darker lines are more augmentative 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.

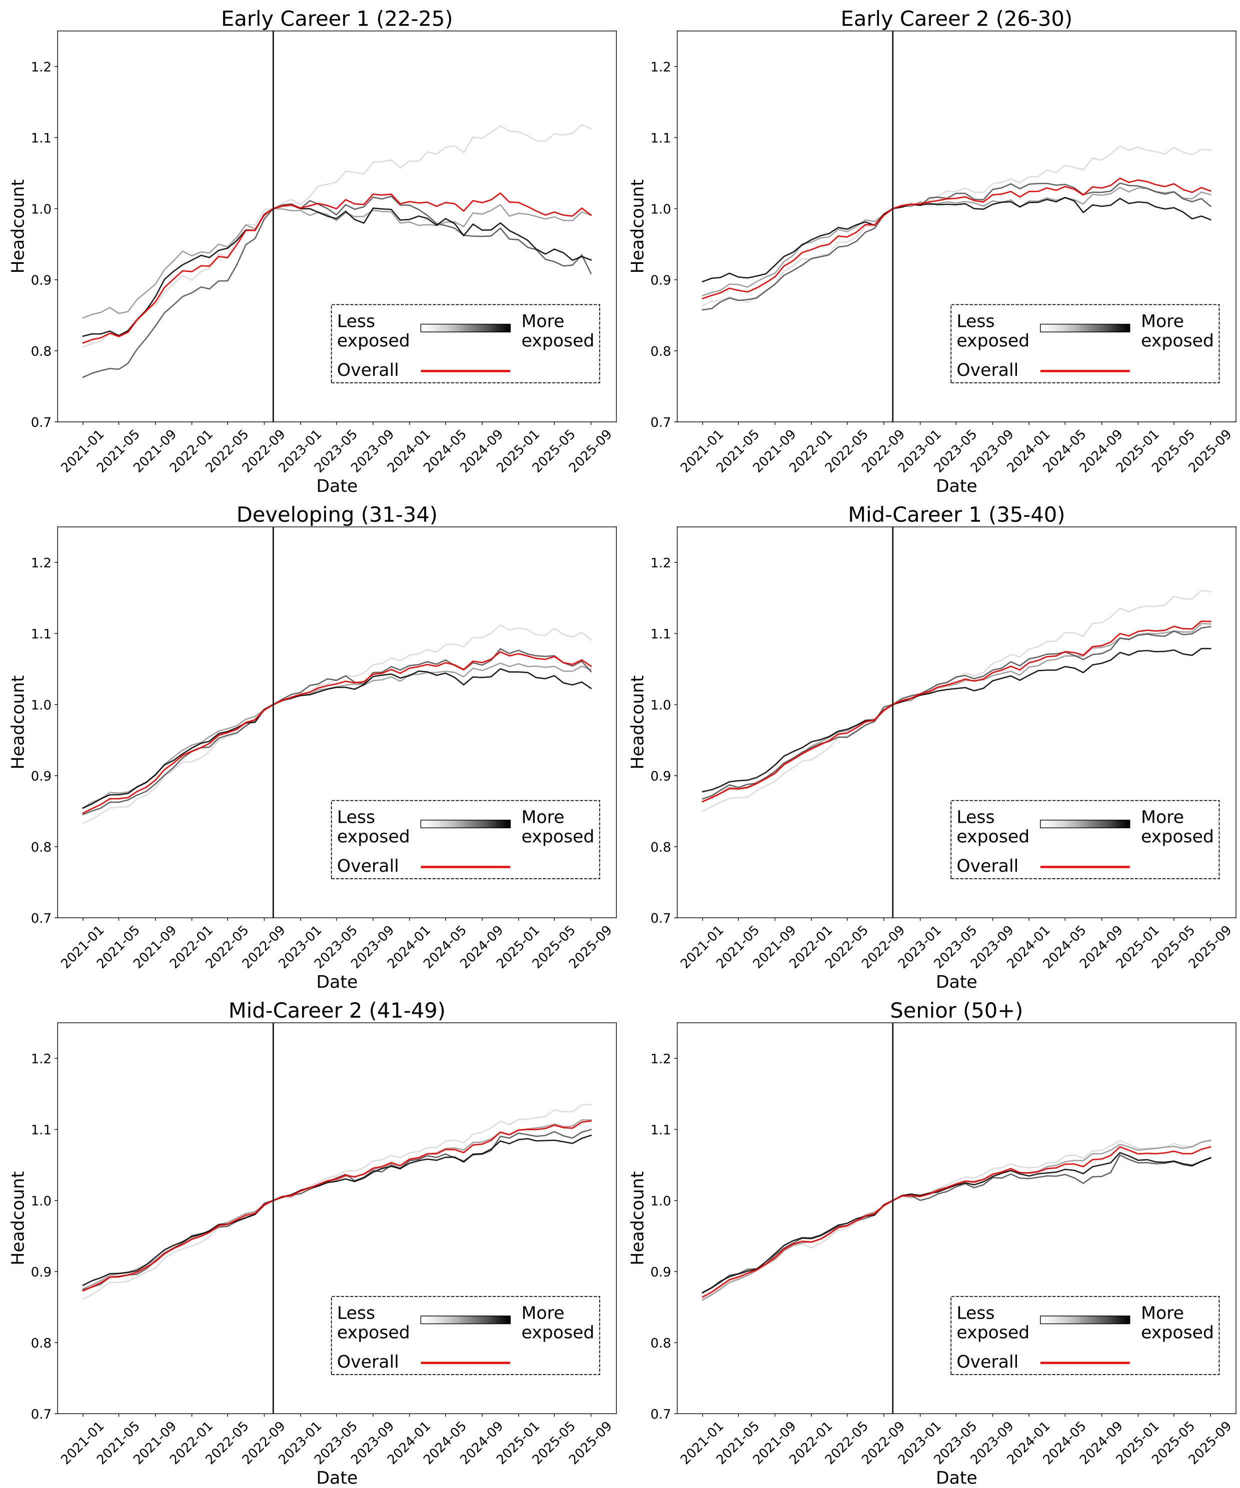

Figure A9: Employment changes by age and automation level using Claude usage data from Handa et al. (2025). Excluding occupations that have no Claude usage or are in the lowest quintile of overall Claude usage conditional on some usage.

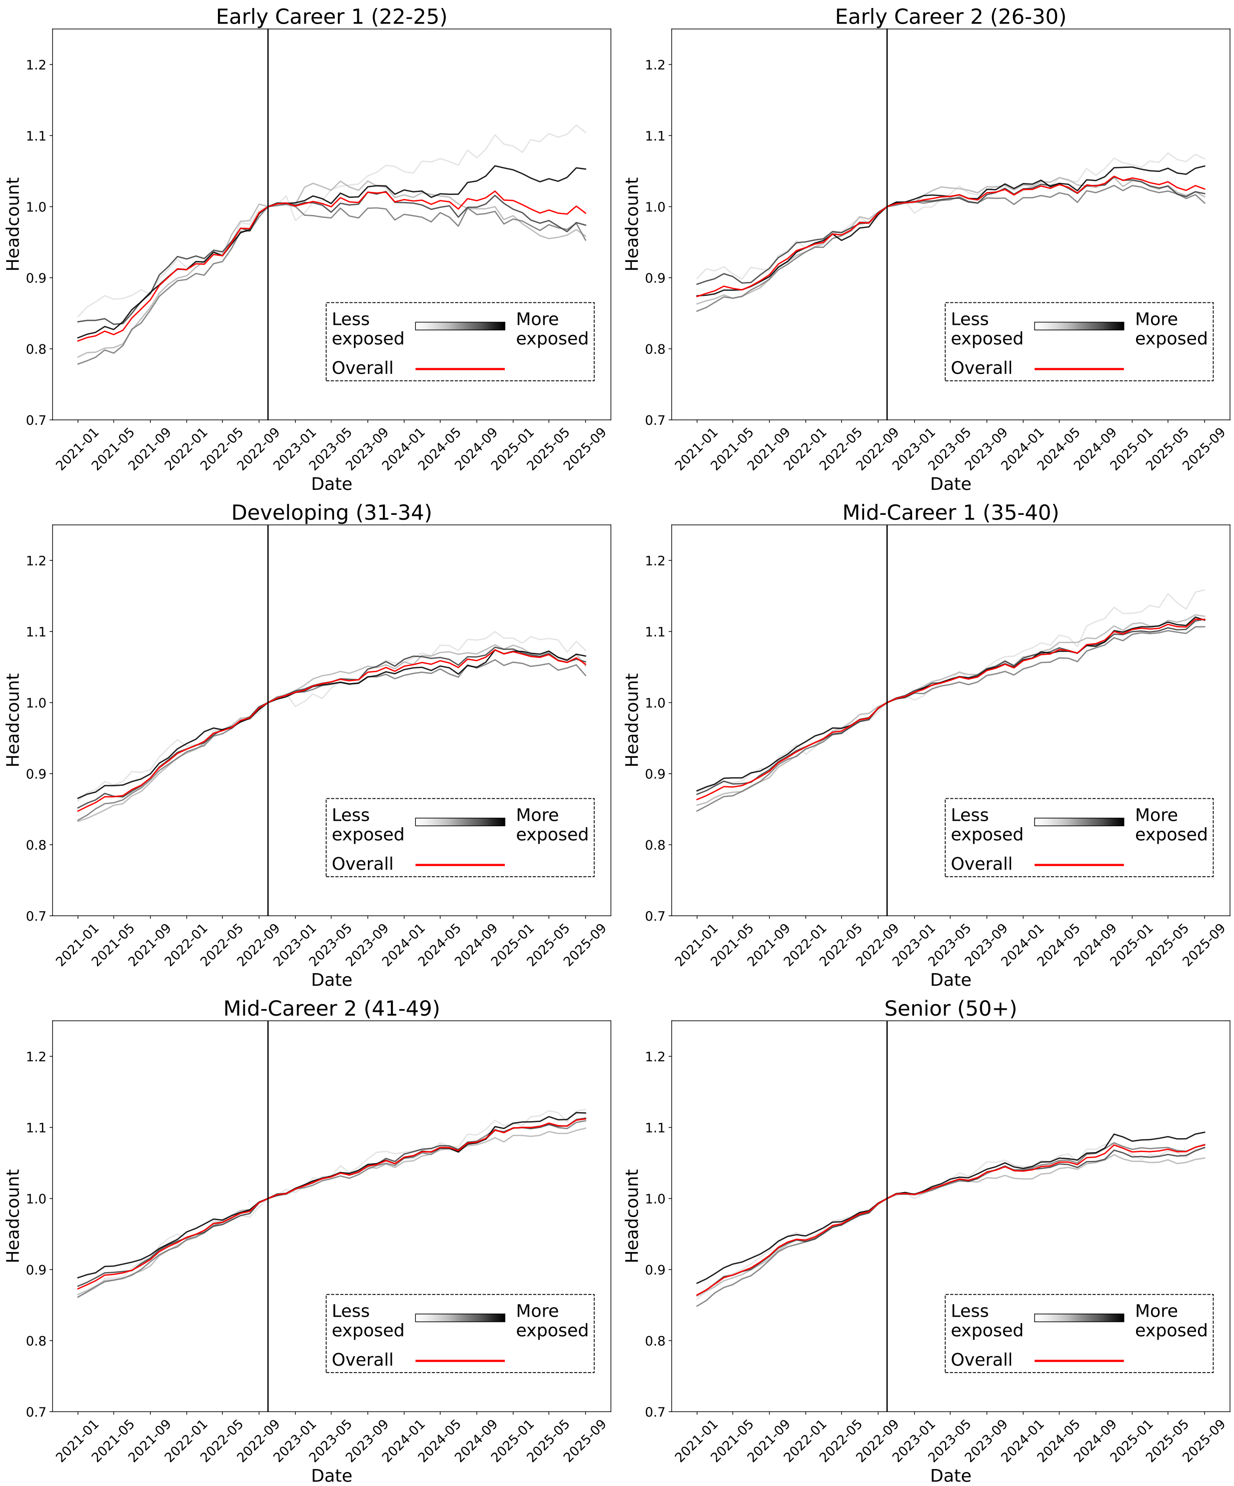

Figure A10: Employment changes by age and augmentation quintile using Claude usage data from Handa et al. (2025). Excluding occupations that have no Claude usage or are in the lowest quintile of overall Claude usage conditional on some usage.

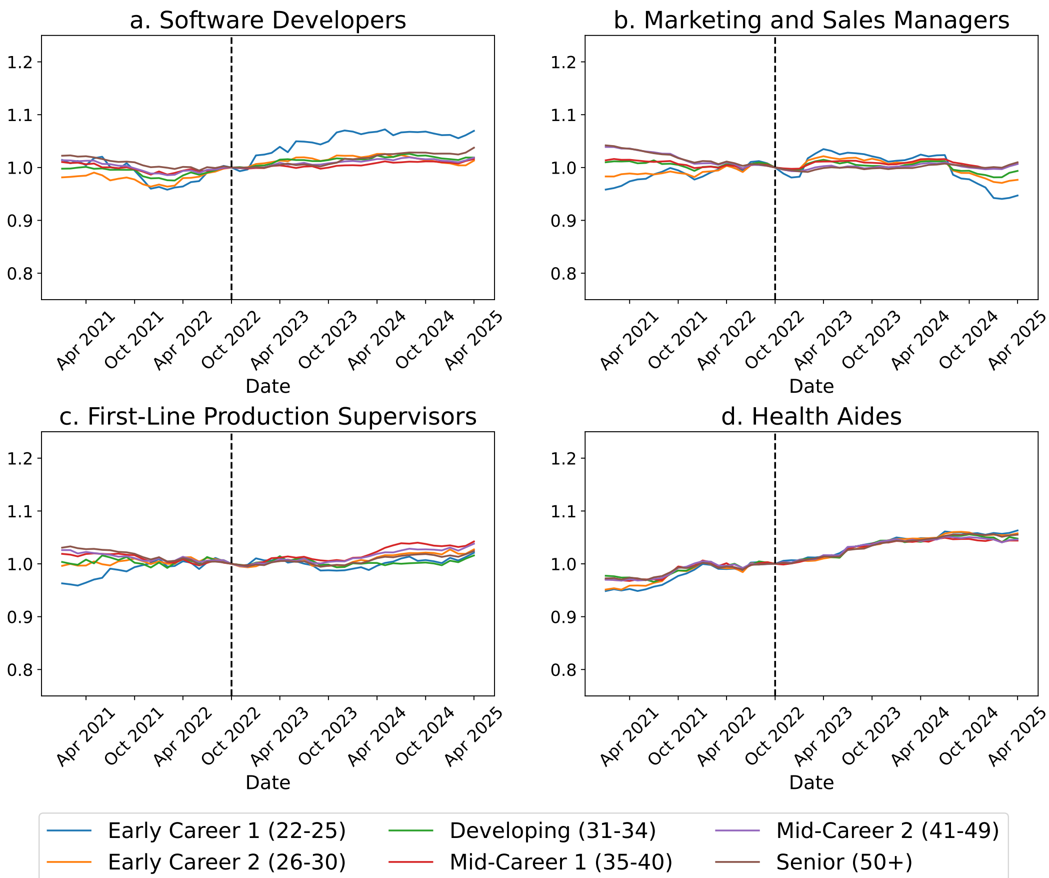

Figure A11: Changes in annual base compensation by age and occupation. Annual base compensation is deflated to 2017 dollars using the PCE deflator.

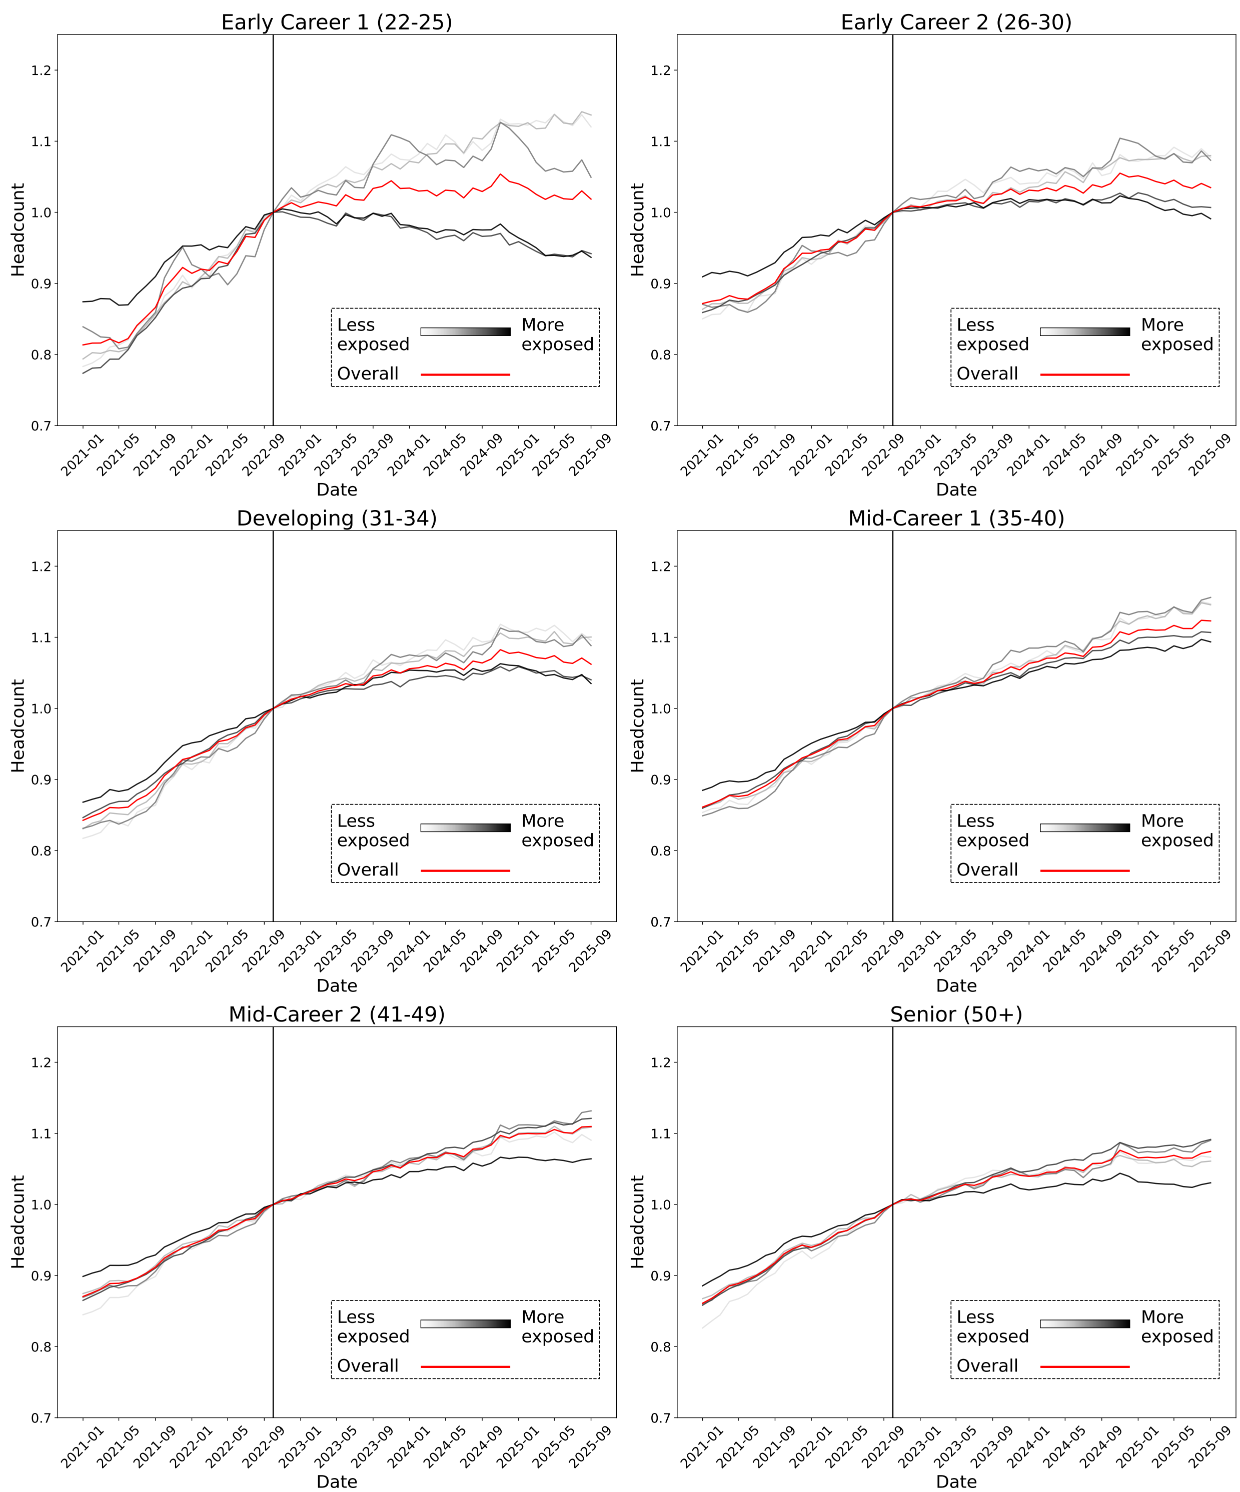

Figure A12: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Excluding computer occupations (2010 SOC codes starting with 15-1).

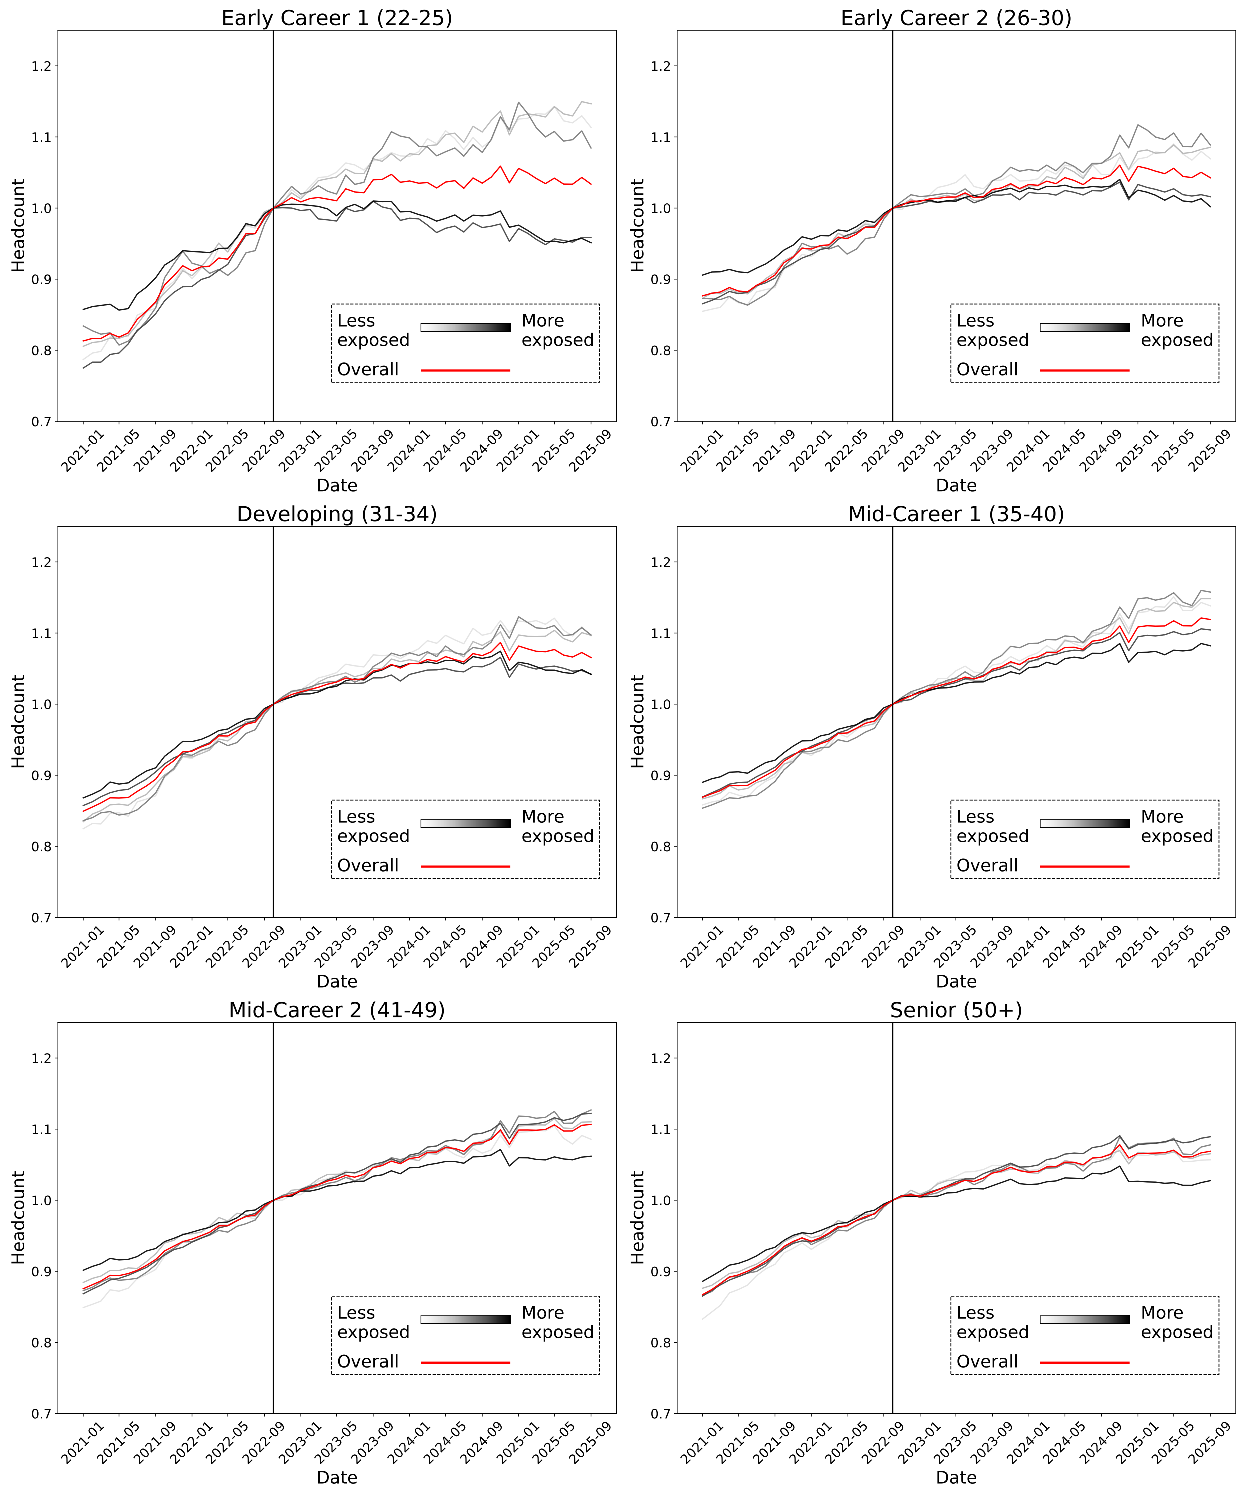

Figure A13: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Excluding firms in the information sector (NAICS code 51).

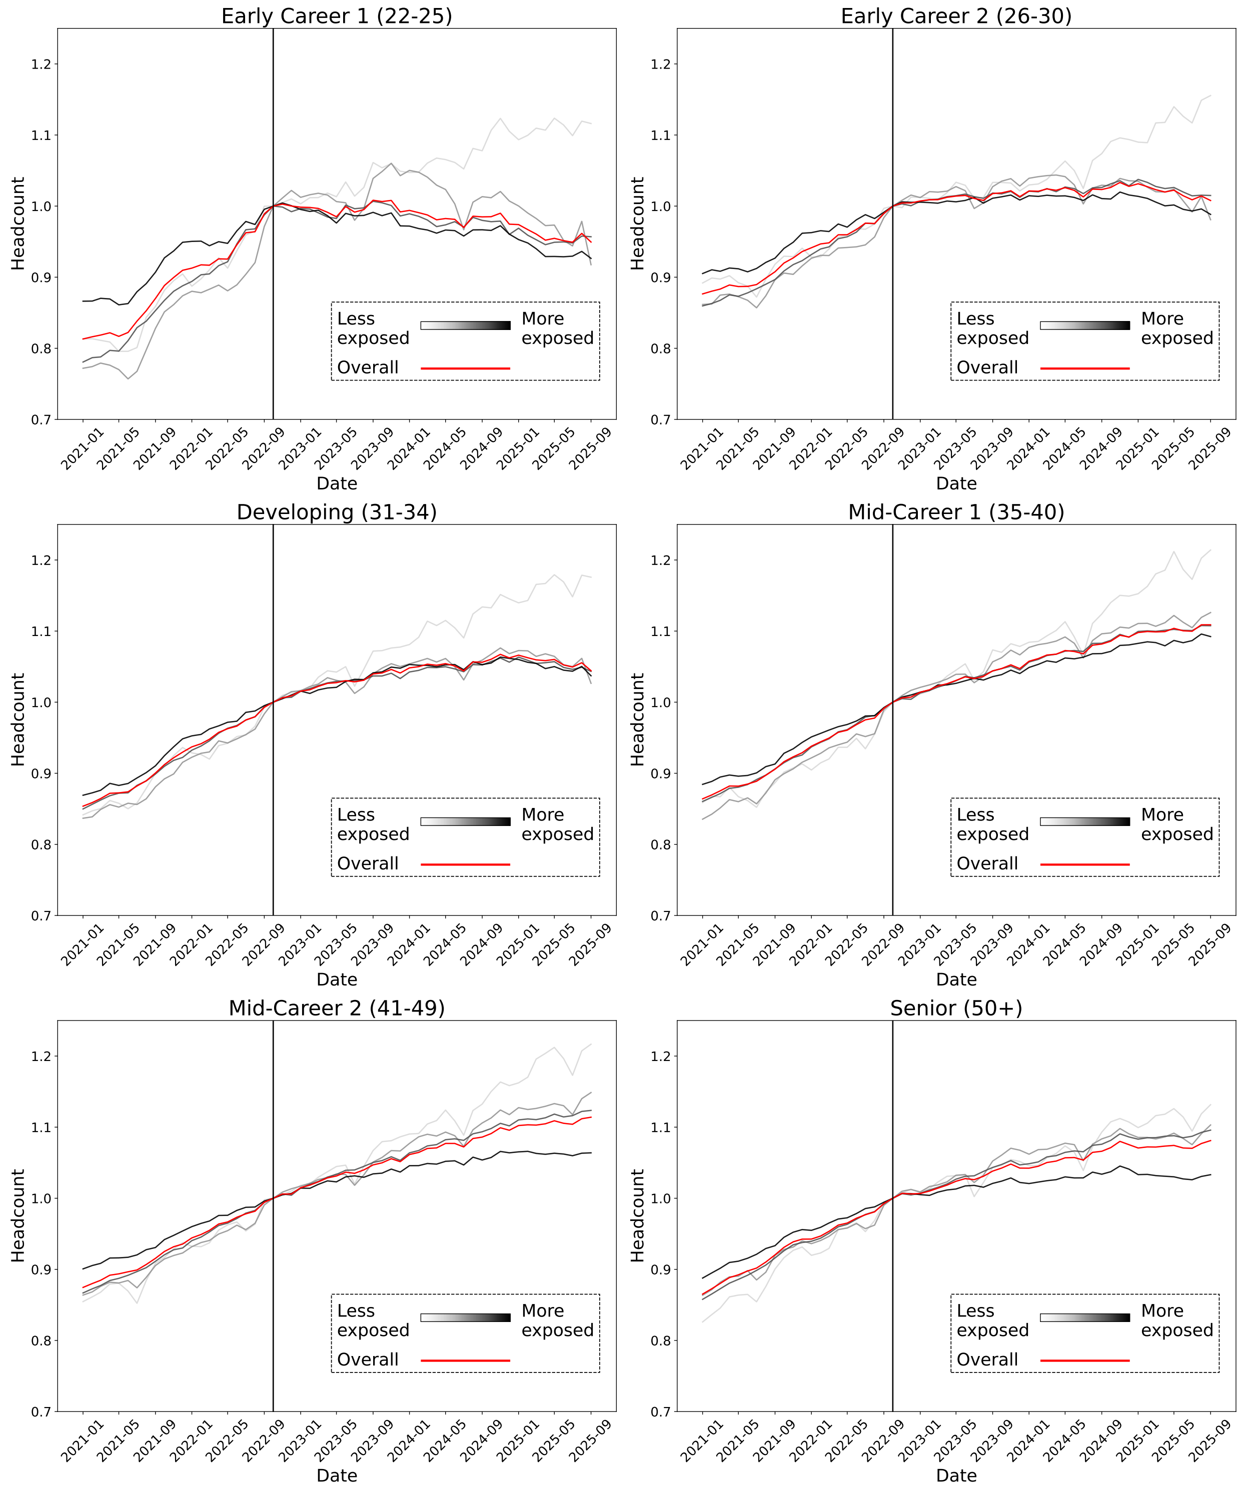

Figure A14: Employment changes by age and exposure group using measures from Eloundou et al. (2024). Including only teleworkable occupations according to Dingel and Neiman (2020). Note that very few teleworkable occupations fall in the lowest exposure quintile. All occupations in the first and second quintile are consequently grouped together in level 1. The remaining quintiles are coded as 2, 3, and 4.

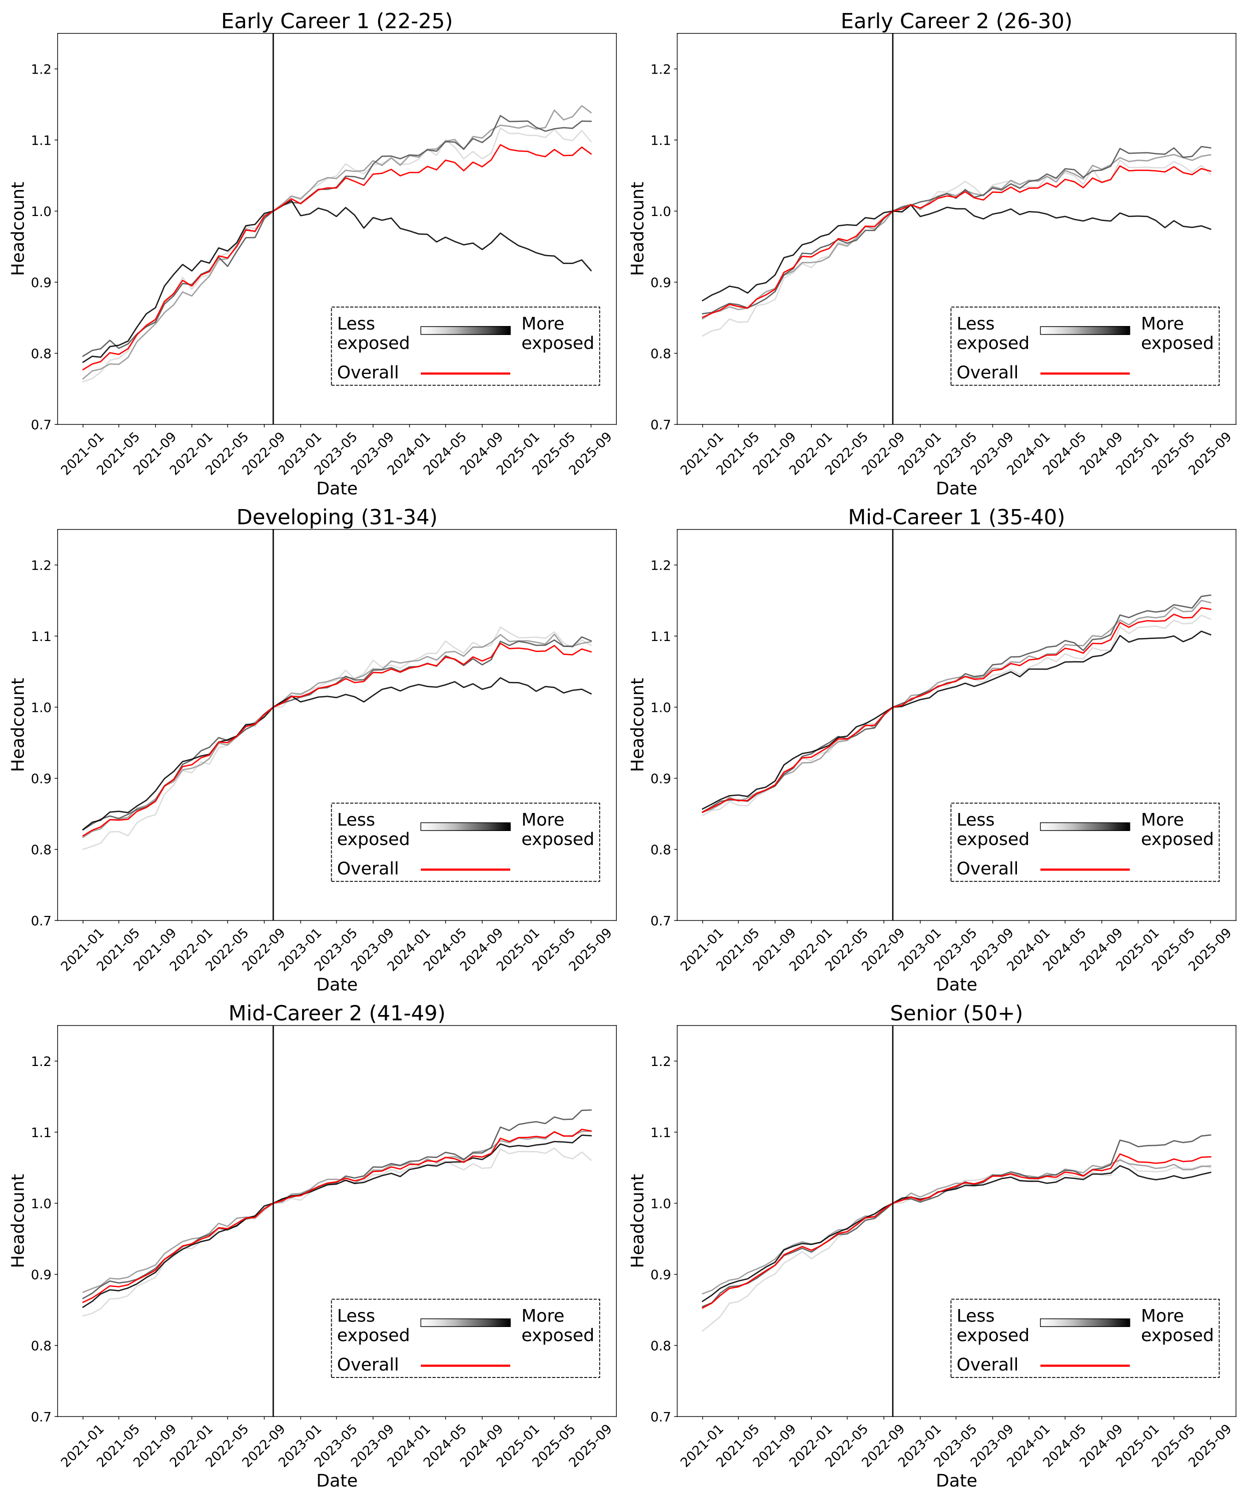

Figure A15: Employment changes by age and exposure group using measures from Eloundou et al. (2024). Including only non-teleworkable occupations according to Dingel and Neiman (2020). Note that very few non-teleworkable occupations fall in the highest exposure quintile. All occupations in the fourth and fifth quintile are consequently grouped together in level 4. The remaining quintiles are coded as 1, 2, and 3.

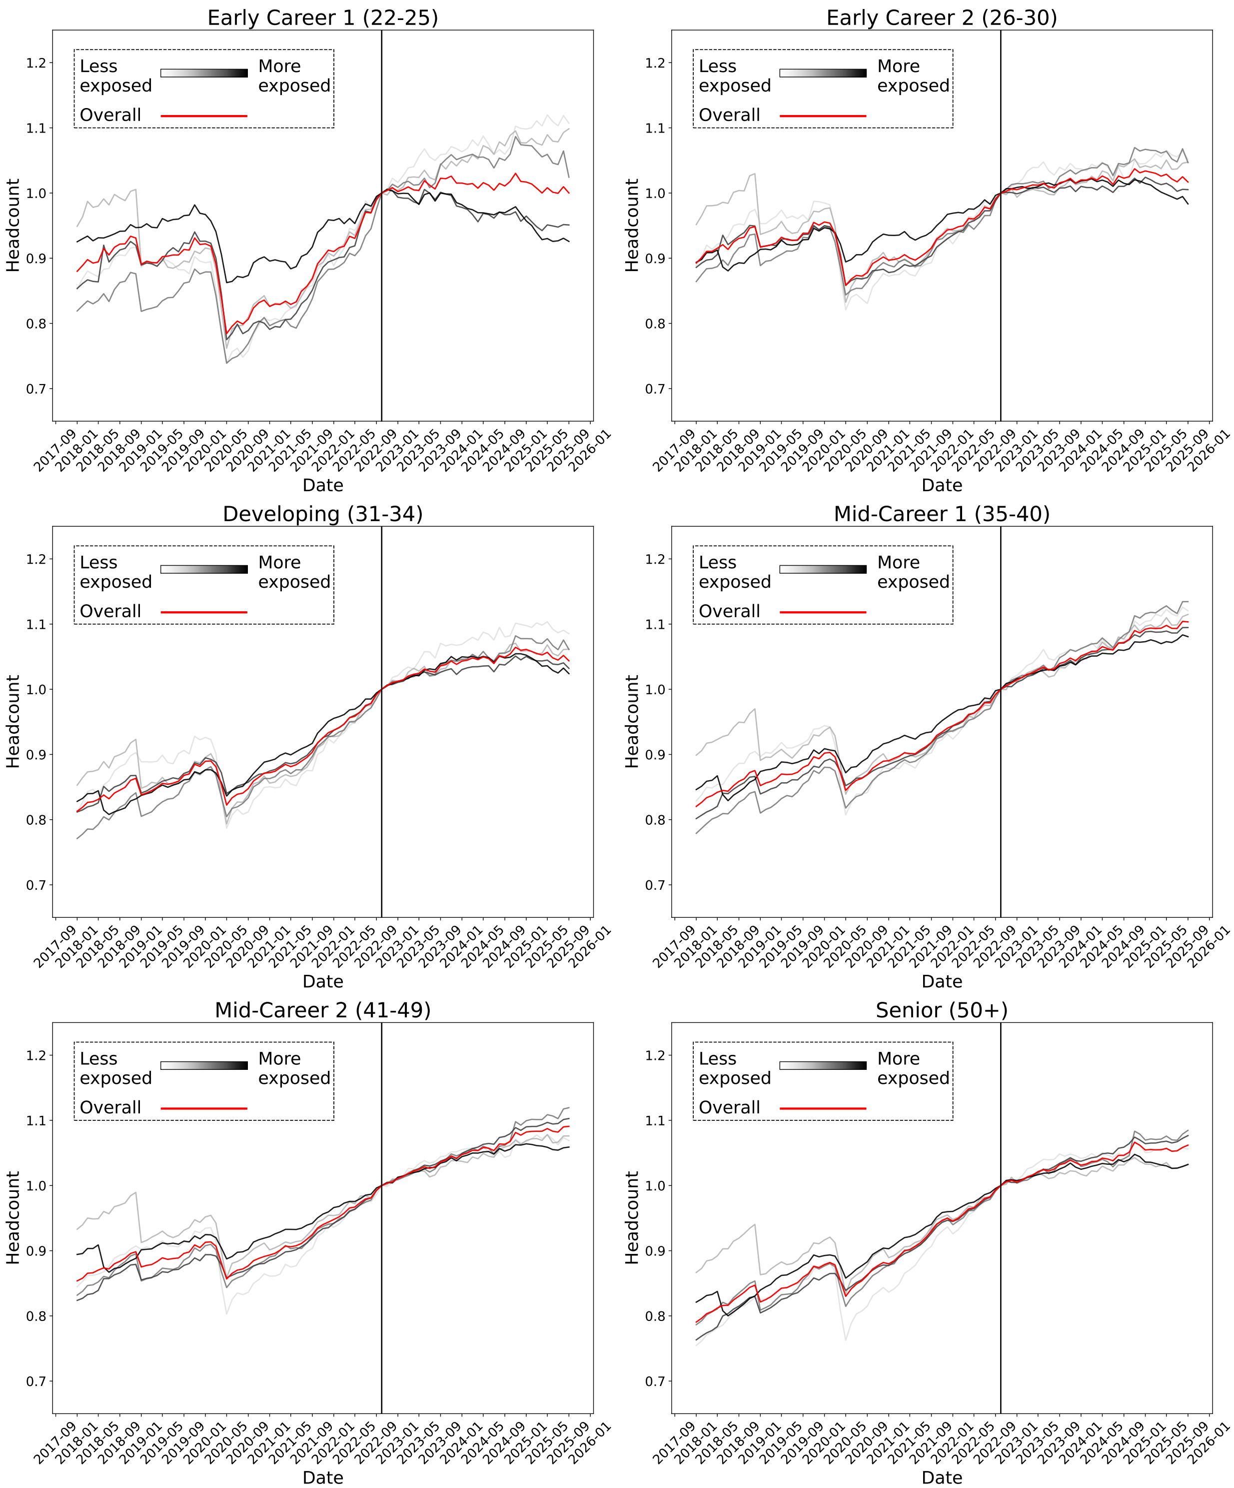

Figure A16: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Data is from 2018 to 2025.

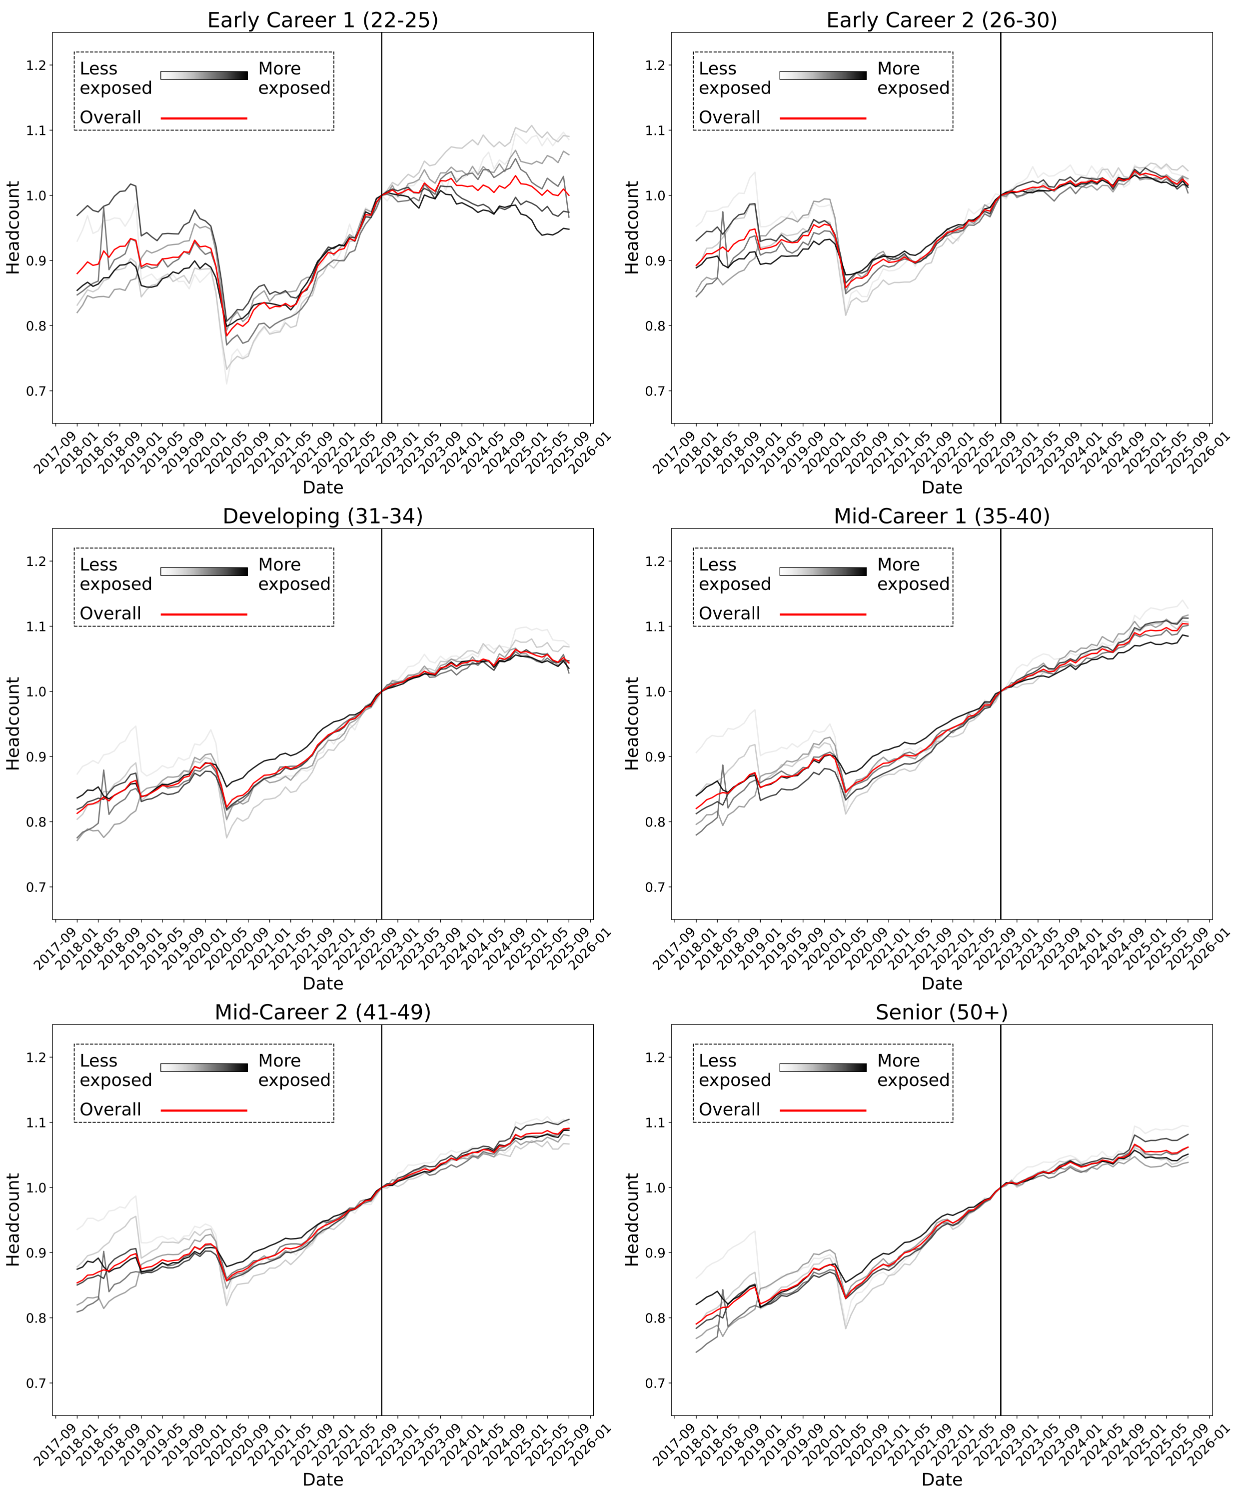

Figure A17: Employment changes by age and exposure quintile using Claude usage data from Handa et al. (2025). Occupations whose associated tasks all have fewer than the minimum number of queries to appear in the usage data are treated as a separate category, coded as 0. Data is from 2018 to 2025.

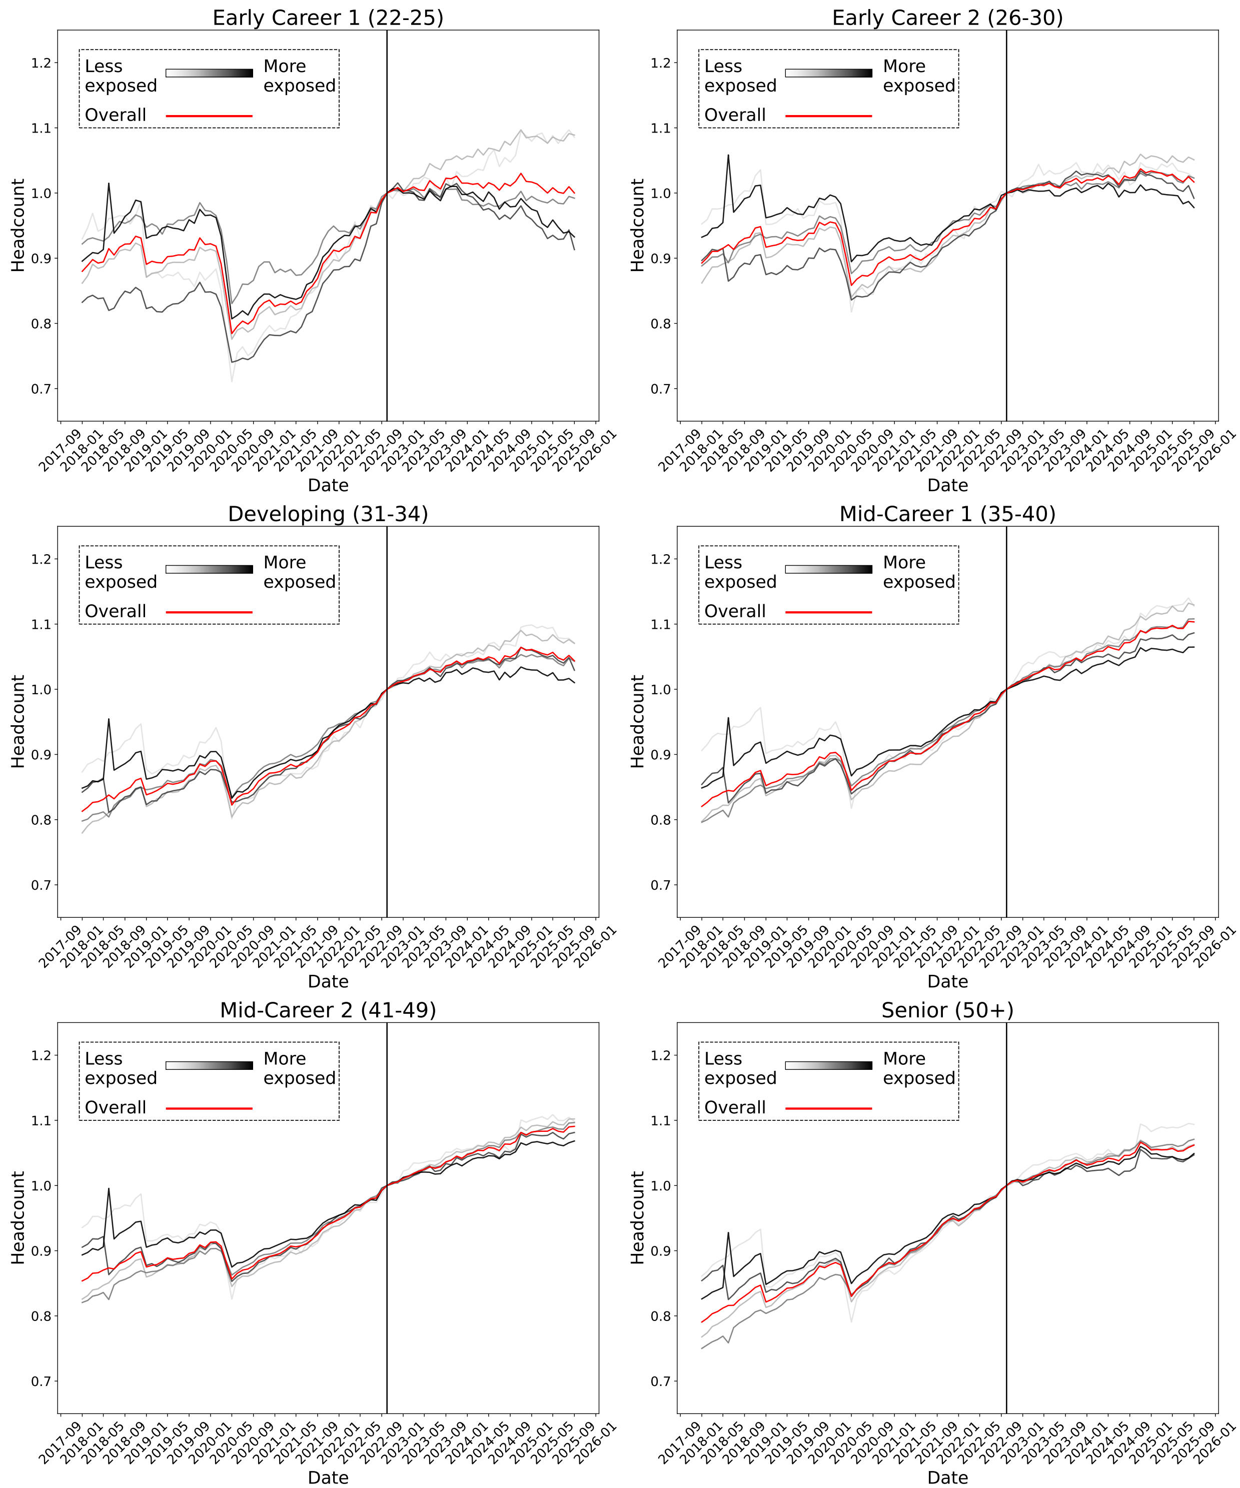

Figure A18: Employment changes by age and automation level using Claude usage data from Handa et al. (2025). Data is from 2018 to 2025.

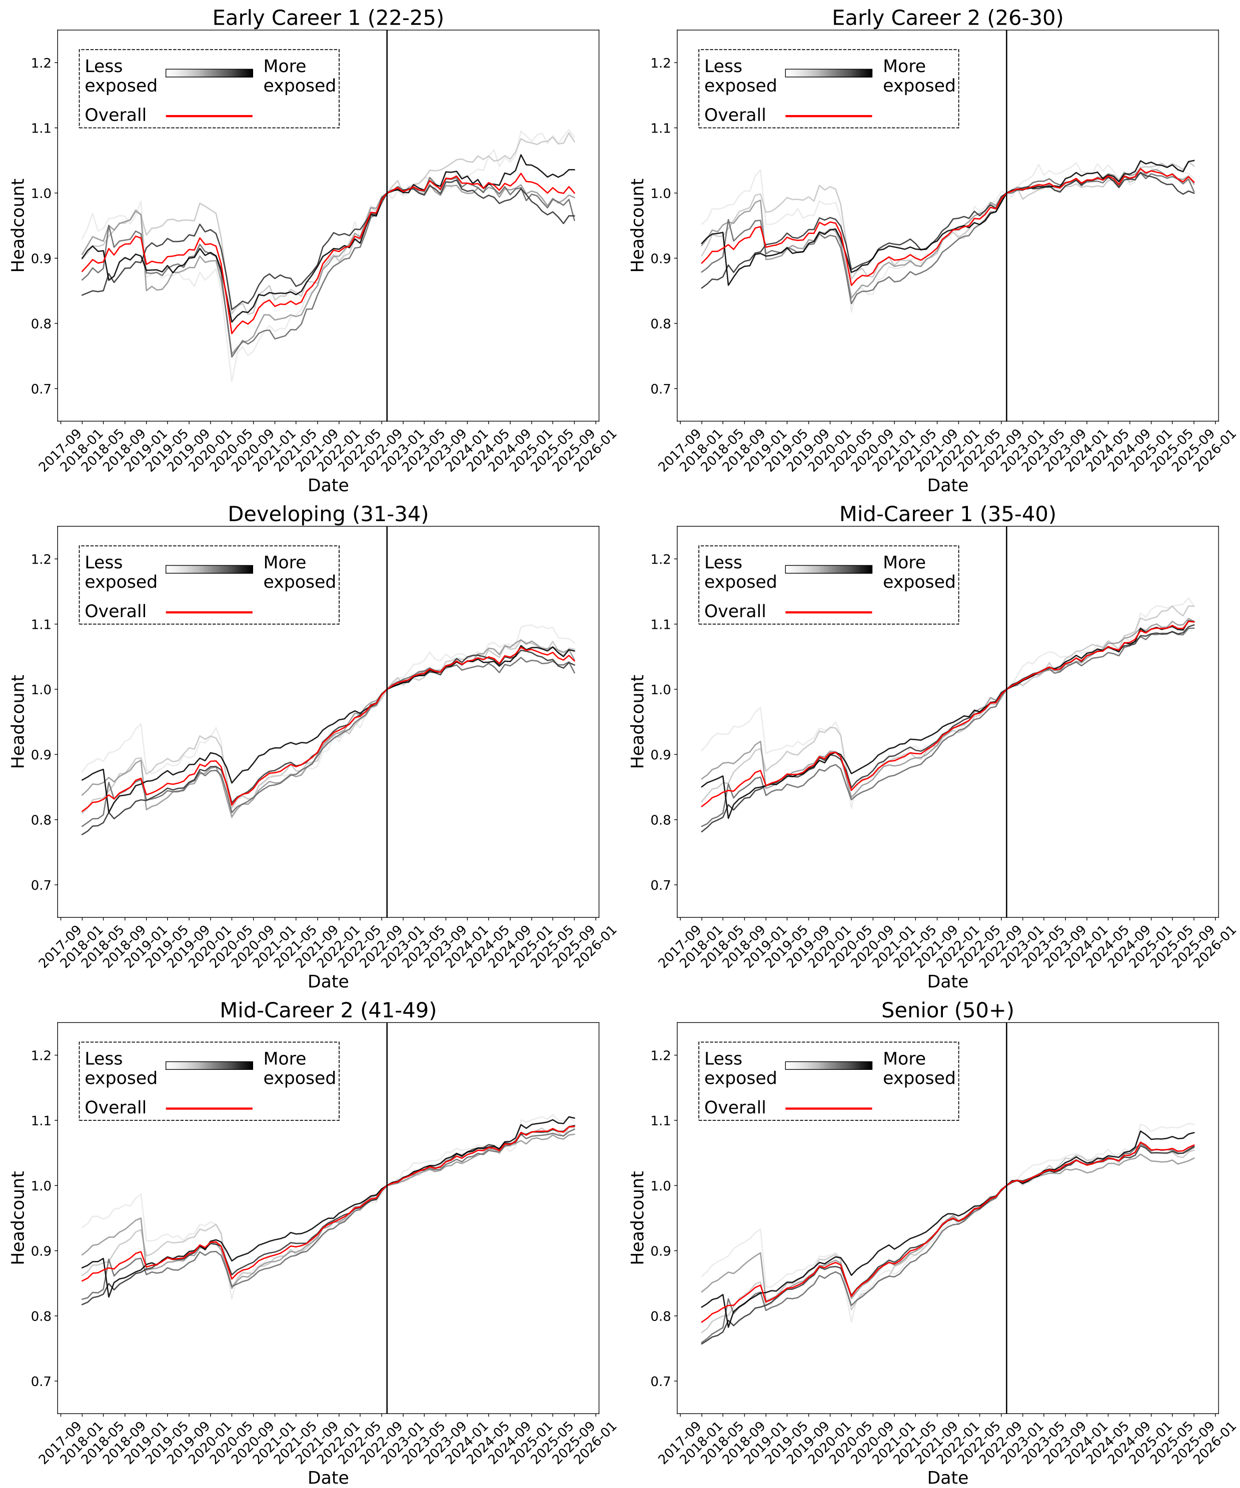

Figure A19: Employment changes by age and augmentation quintile using Claude usage data from Handa et al. (2025). Data is from 2018 to 2025.

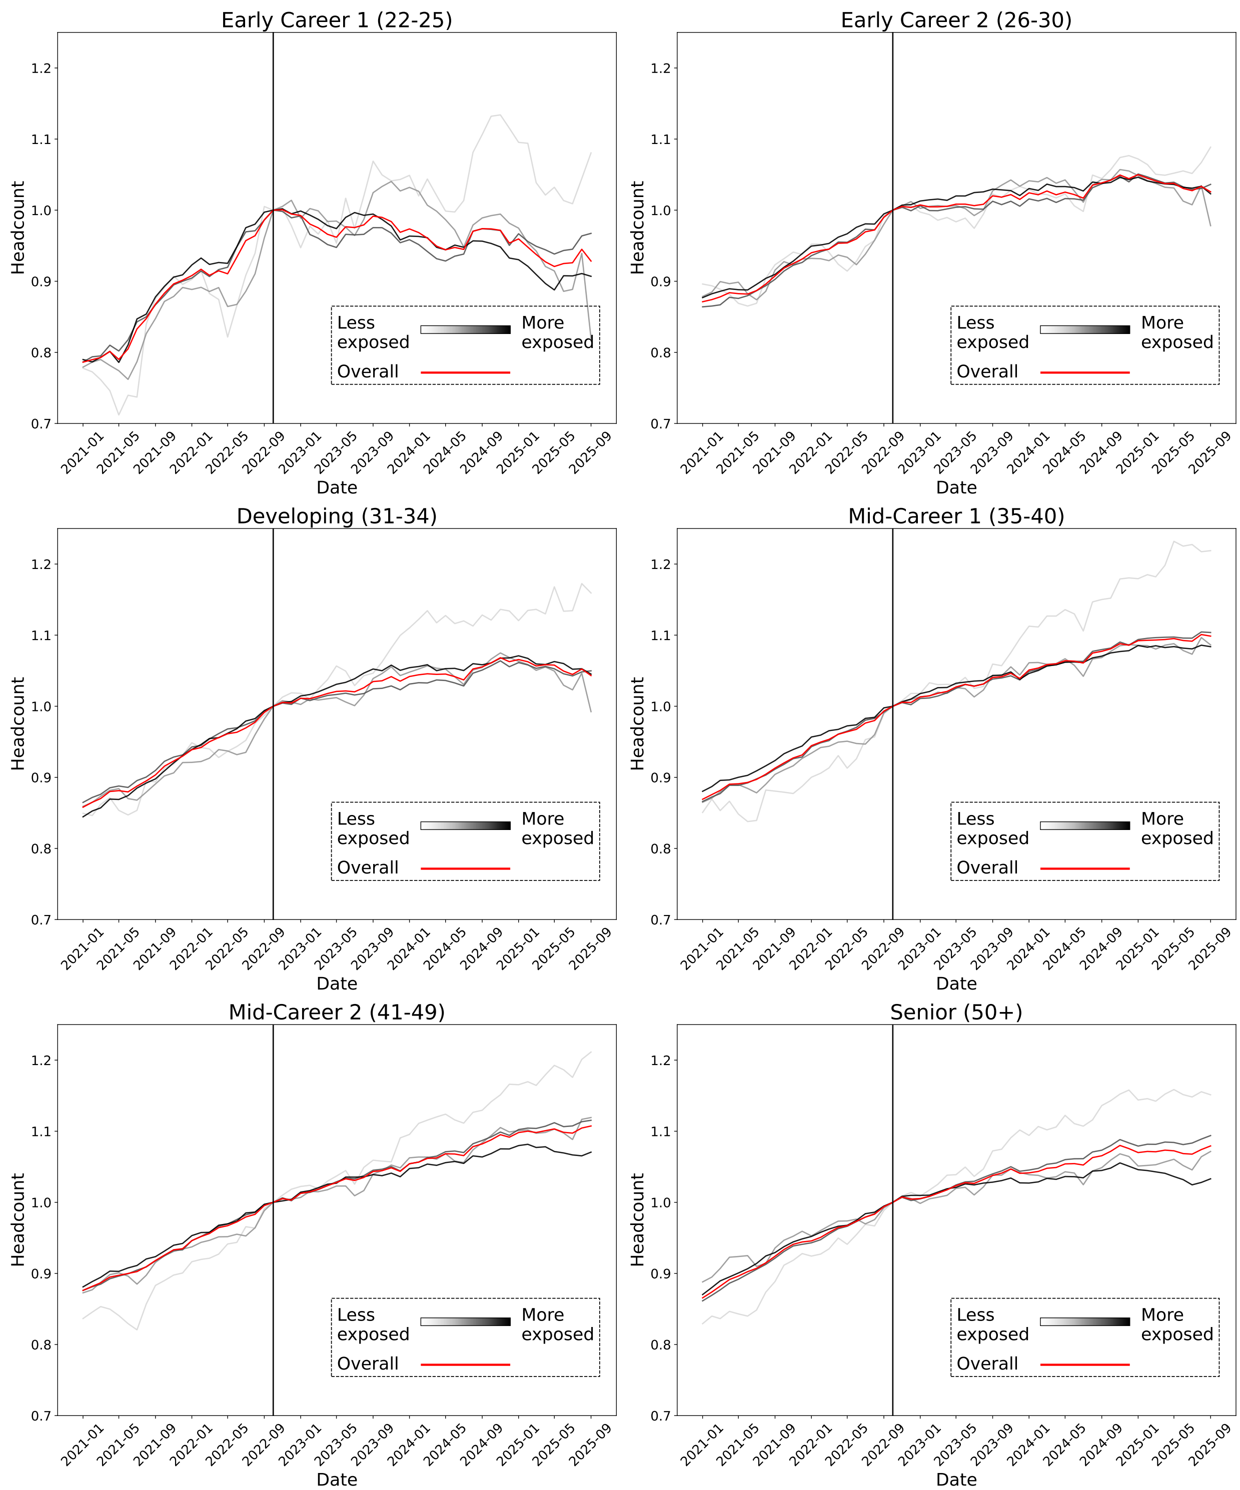

Figure A20: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only occupations in which at least 70% of workers have a college degree in the 2017 ACS. Note that no such occupations lie in quintile 1 of the GPT-4 β based exposure measure.

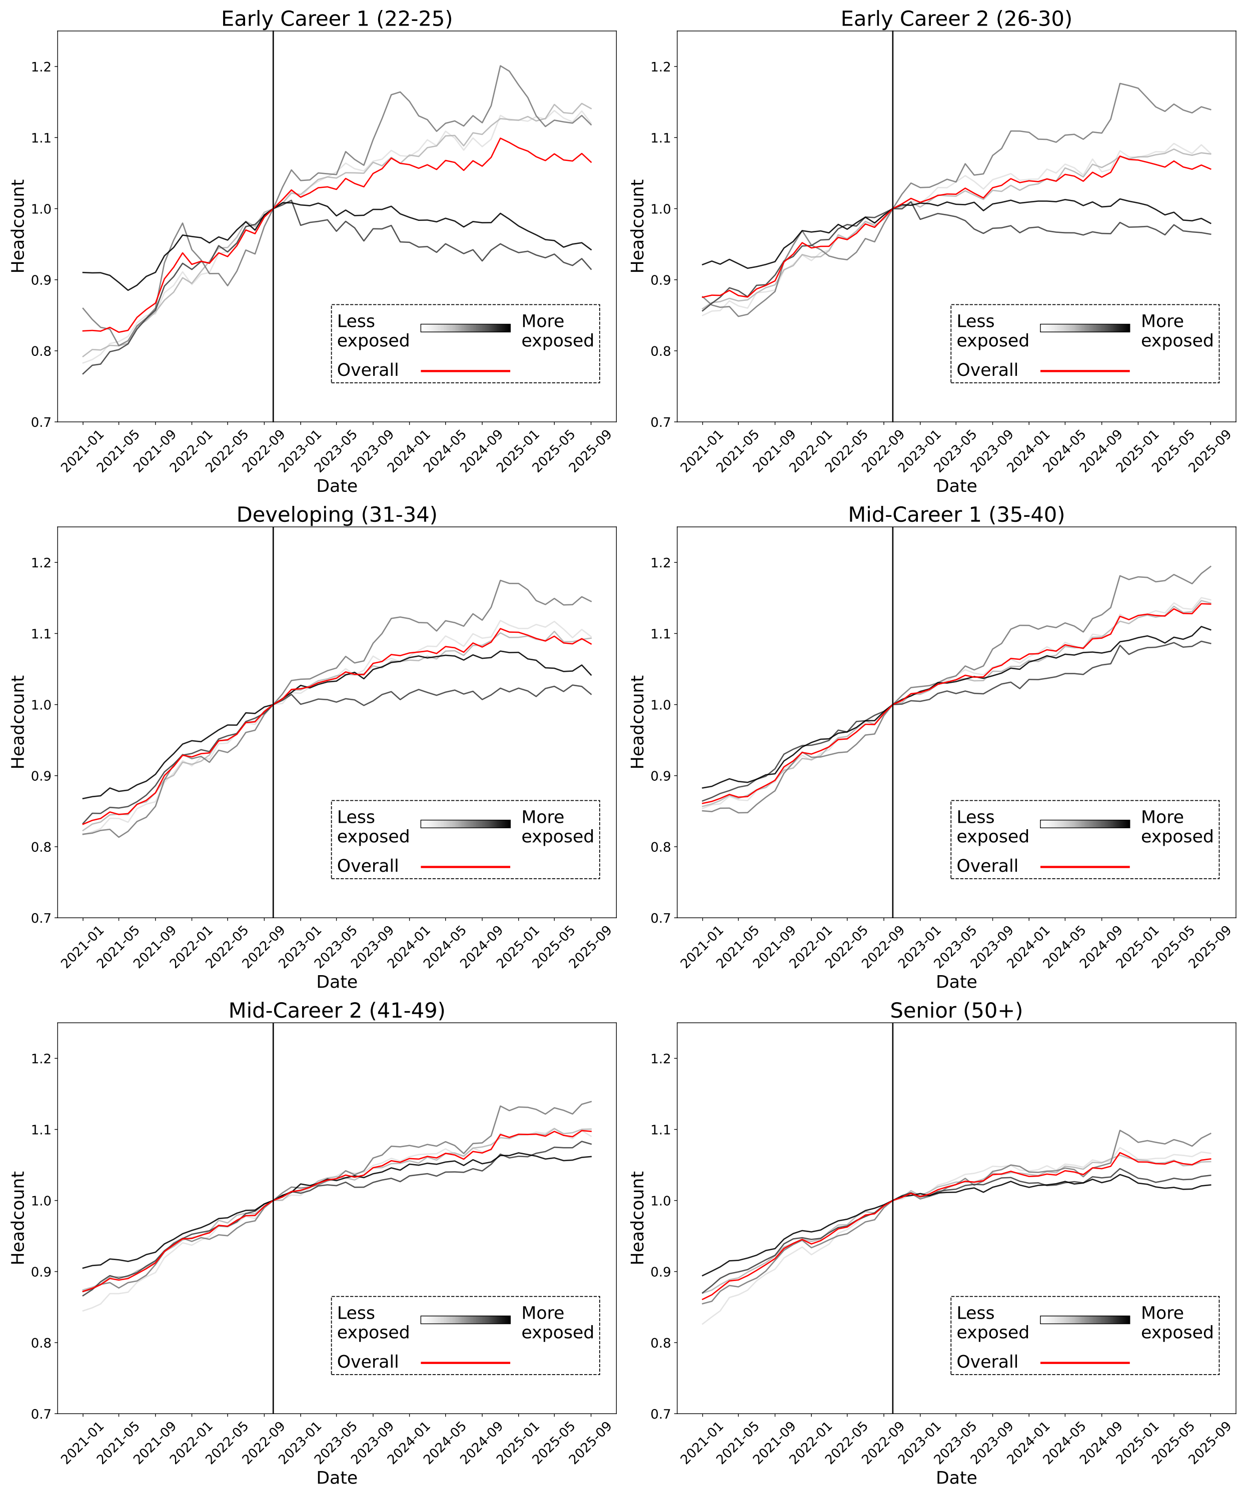

Figure A21: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only occupations in which at most 30% of workers have a college degree in the 2017 ACS.

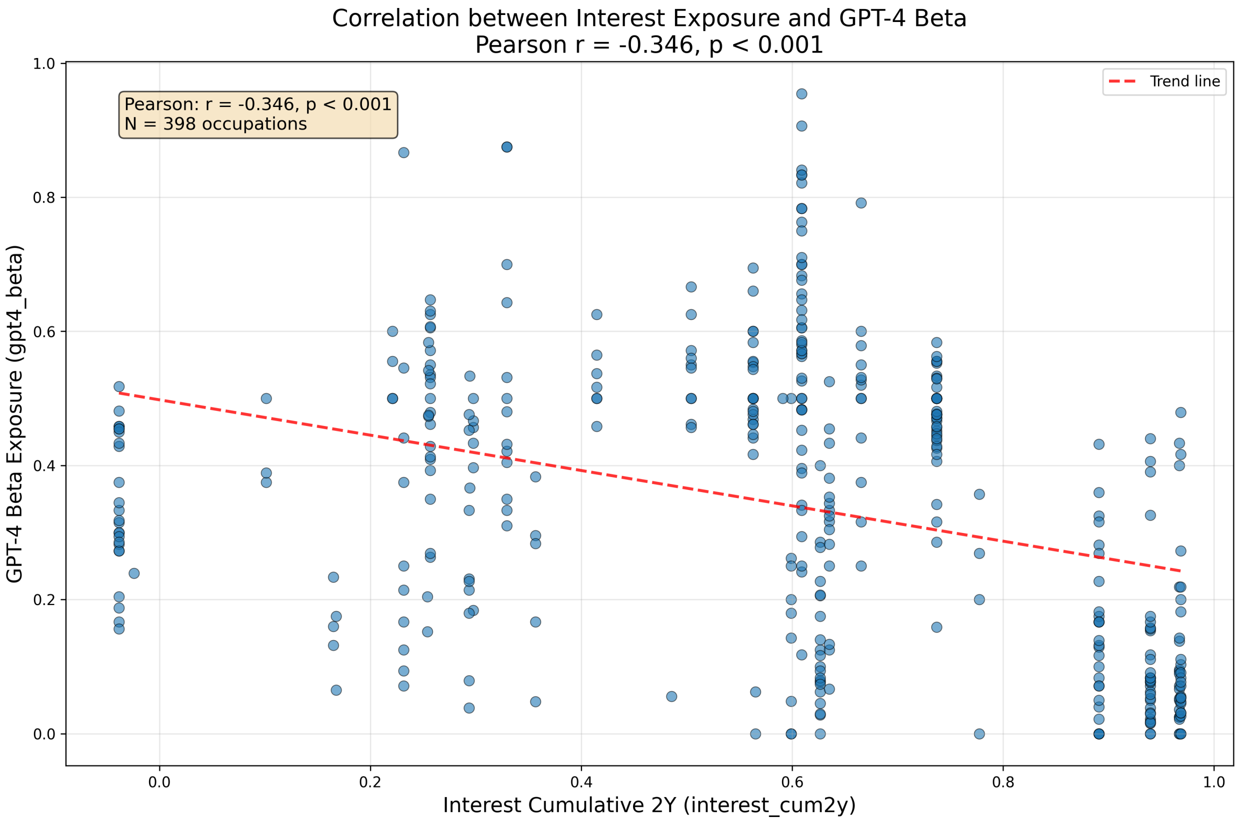

Figure A22: Occupational AI exposure according to Eloundou et al. (2024) compared to interest rate exposure according to Zens et al. (2020). We measure interest rate exposure using the 2-year cumulative impulse response function identified via Cholesky decomposition. We use the crosswalk from Autor (2015) 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, successfully merging to about 500 occupations.

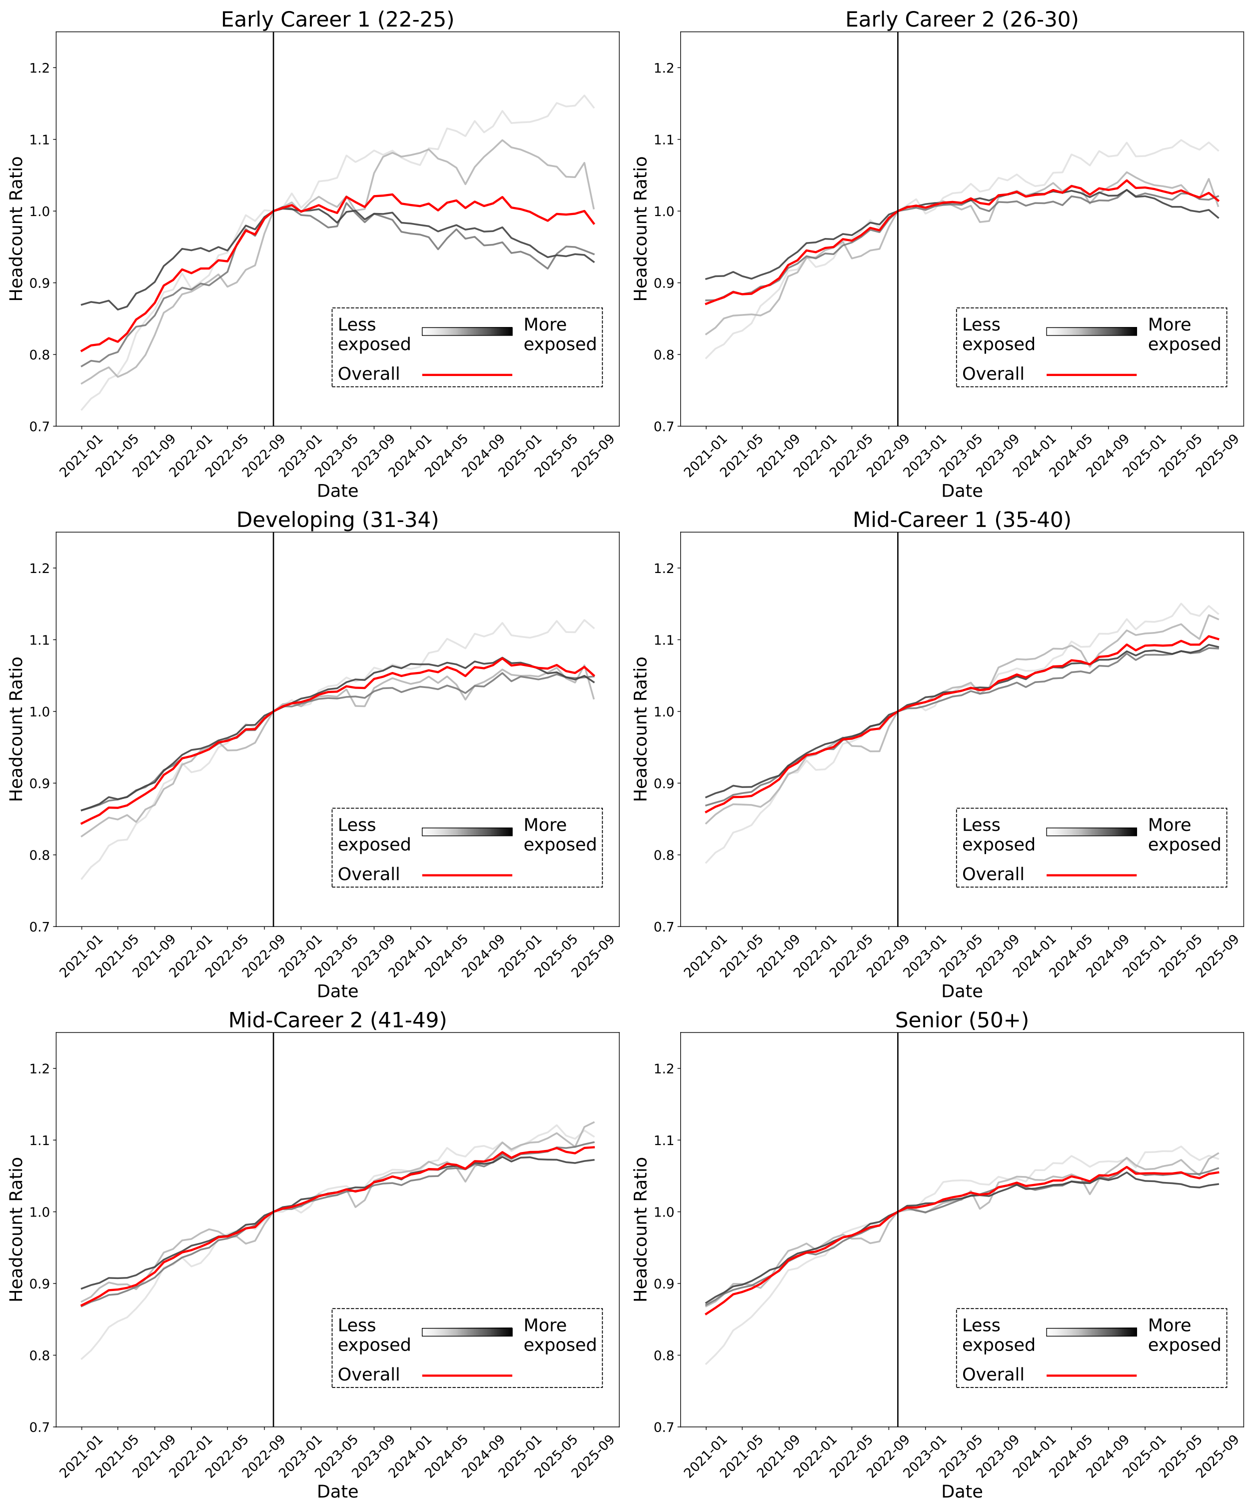

Figure A23: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only occupations with below median interest rate exposure using data from Zens et al. (2020). Only a few occupations with low interest rate exposure also have low AI exposure. All occupations 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.

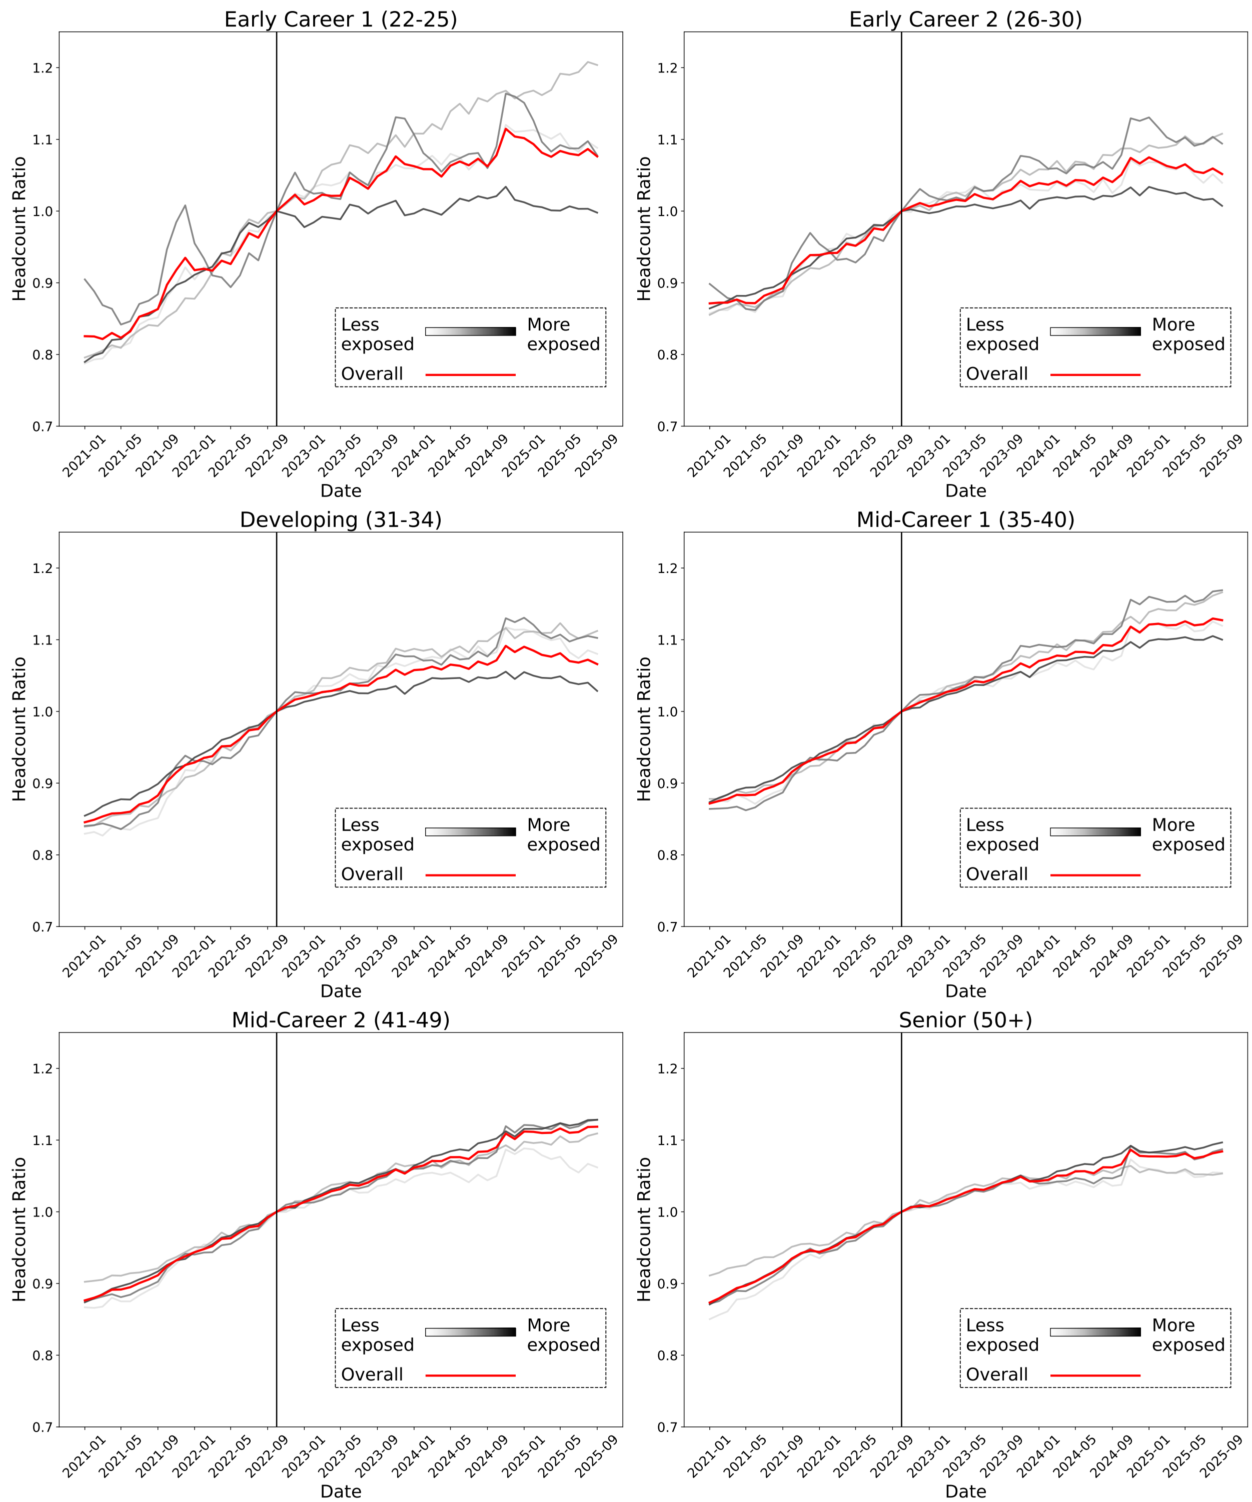

Figure A24: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only occupations with above median interest rate exposure using data from Zens et al. (2020). Only a few occupations with high interest rate exposure also have high AI exposure. All occupations in the fourth and fifth quintile are consequently grouped together in level 4. The remaining quintiles are coded as 1, 2, and 3.

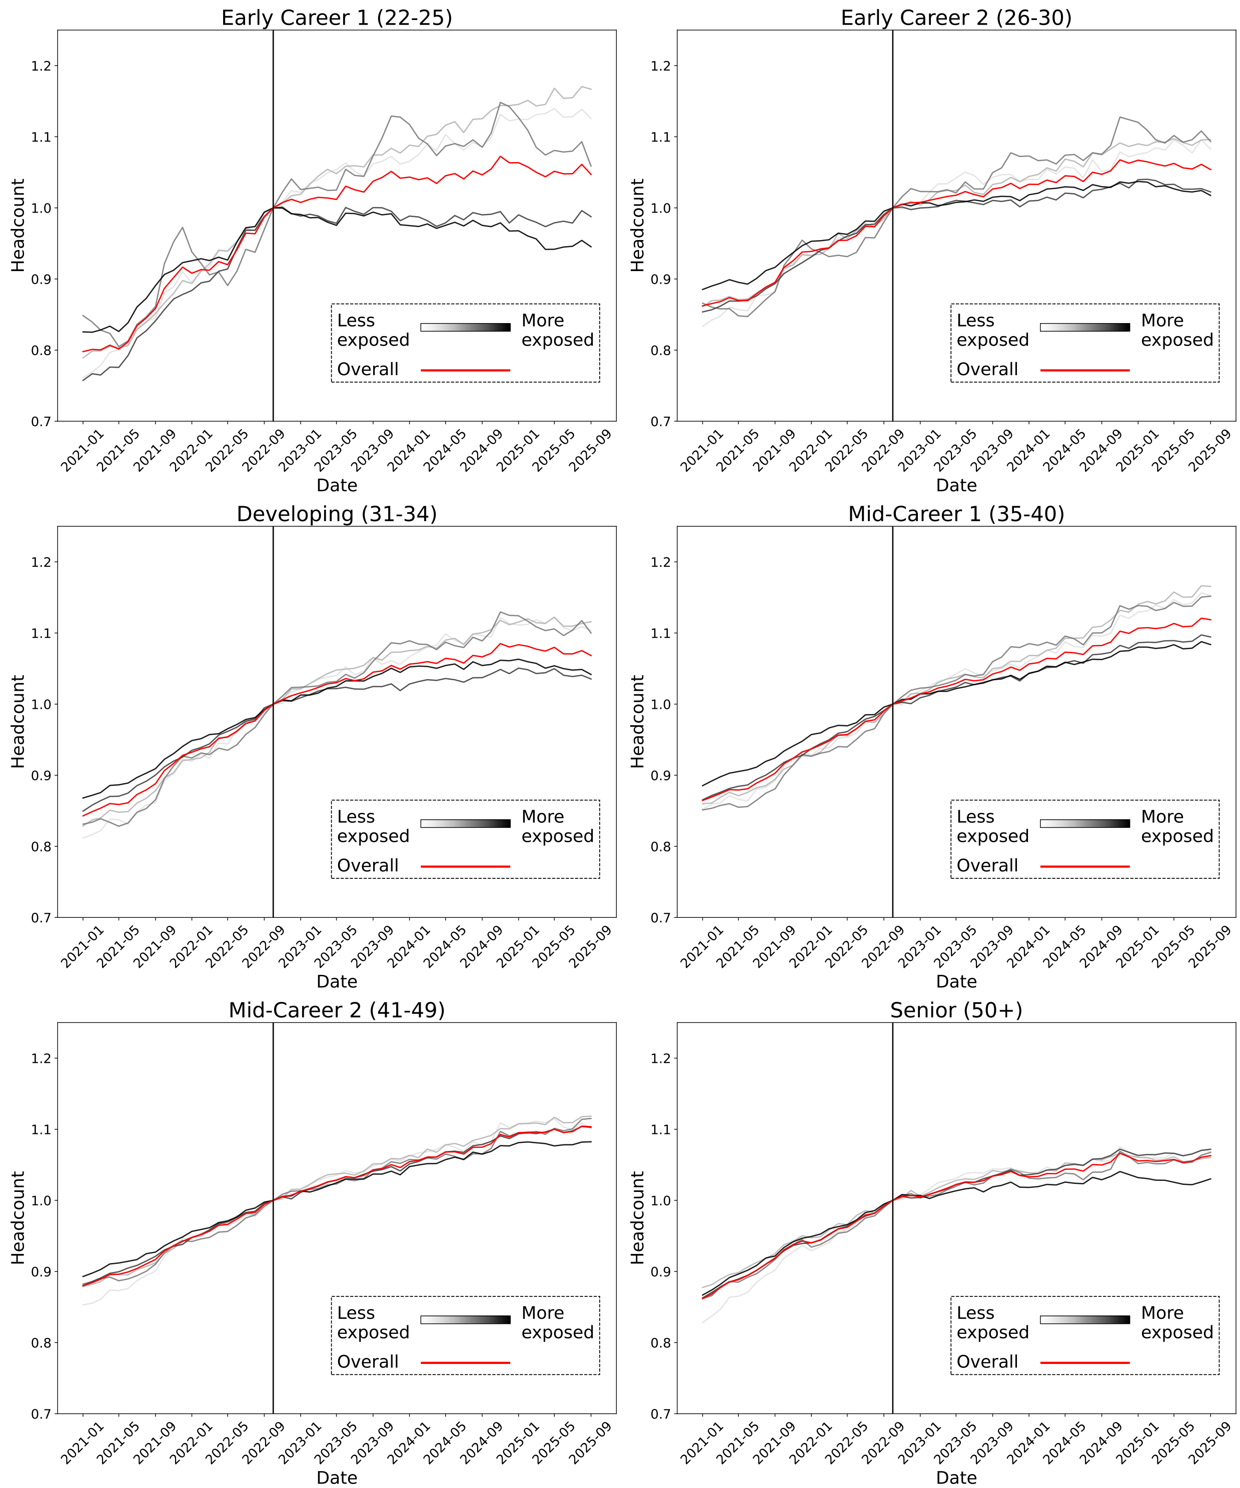

Figure A25: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only men.

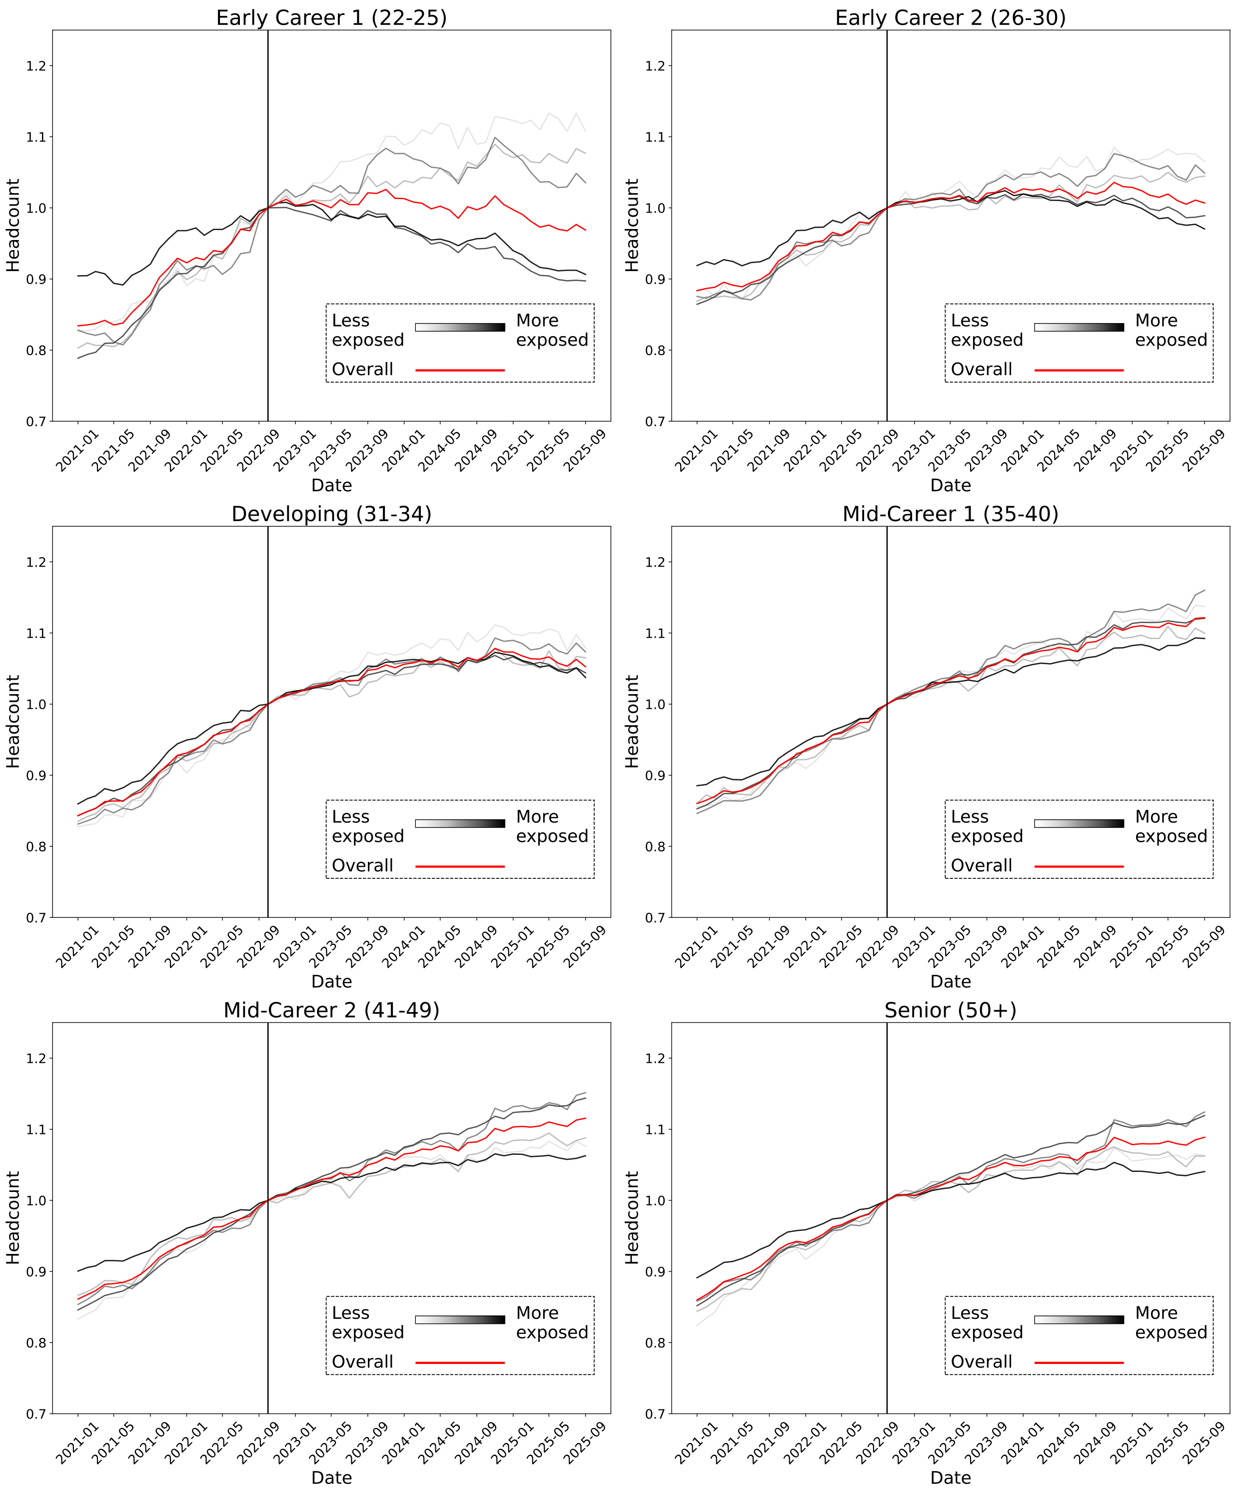

Figure A26: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Considering only women.

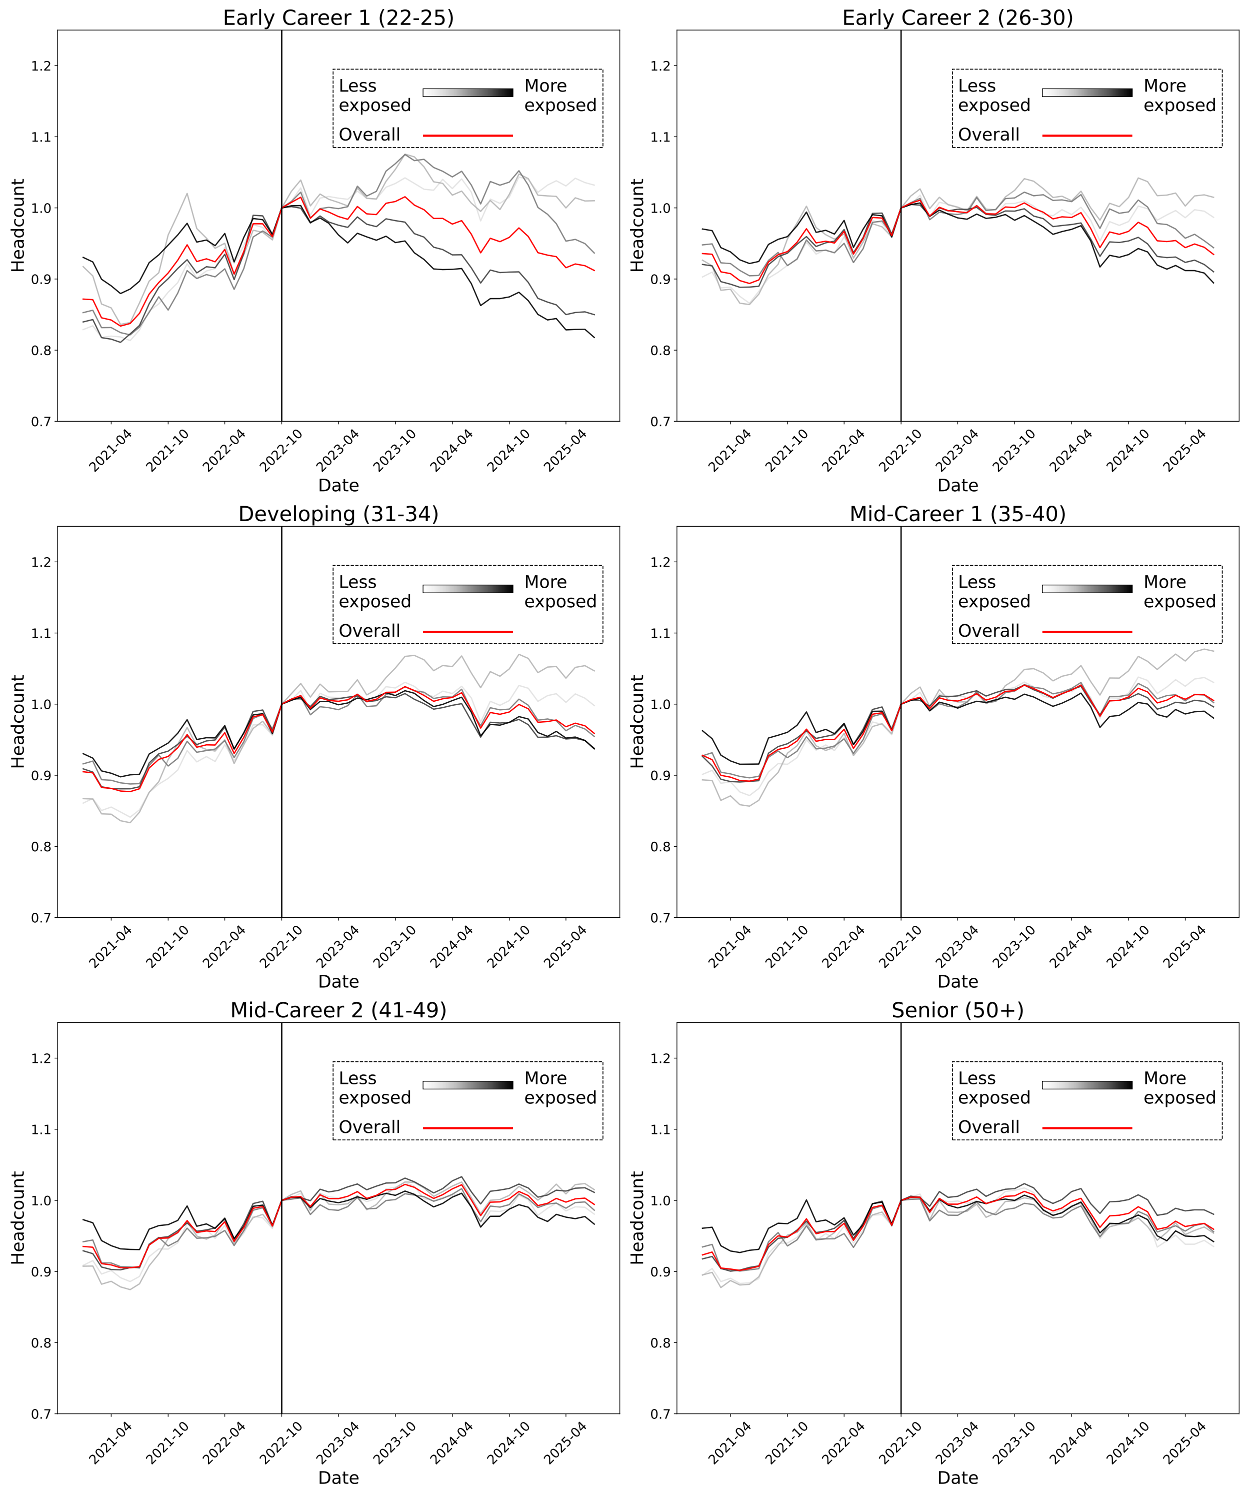

Figure A27: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Using the full sample of firms.

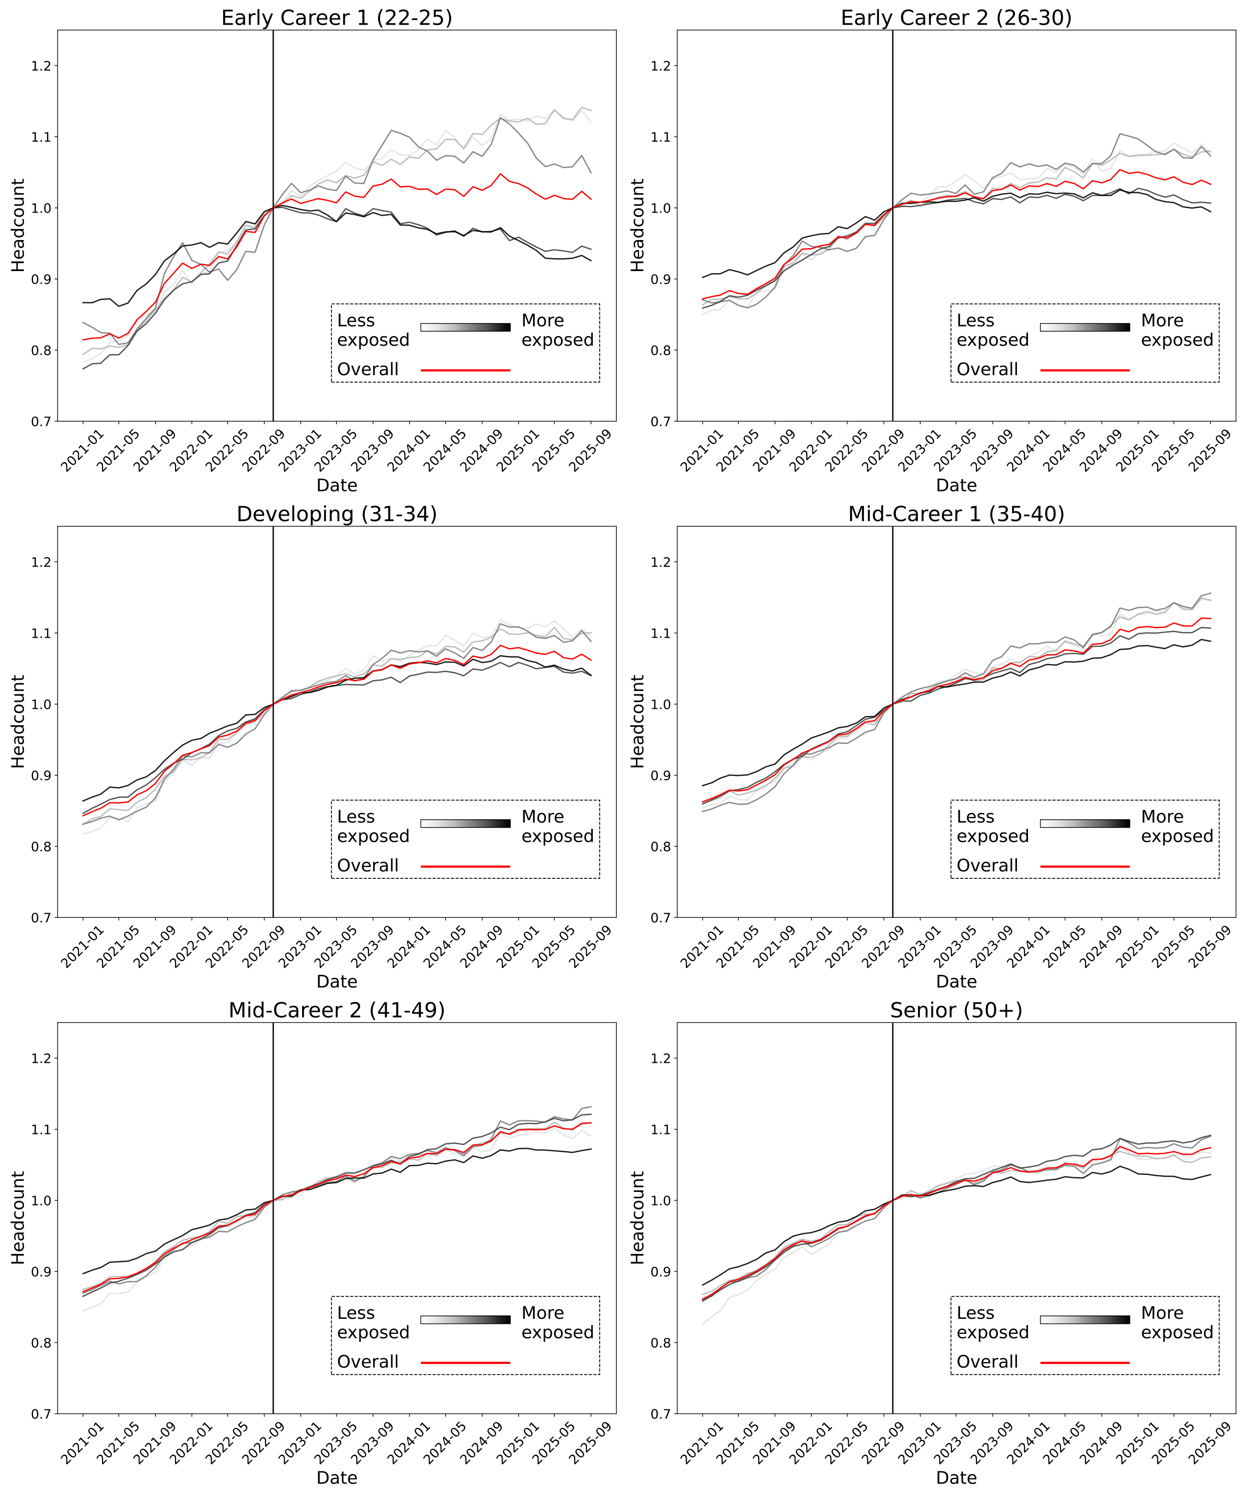

Figure A28: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Including part-time and temporary workers.

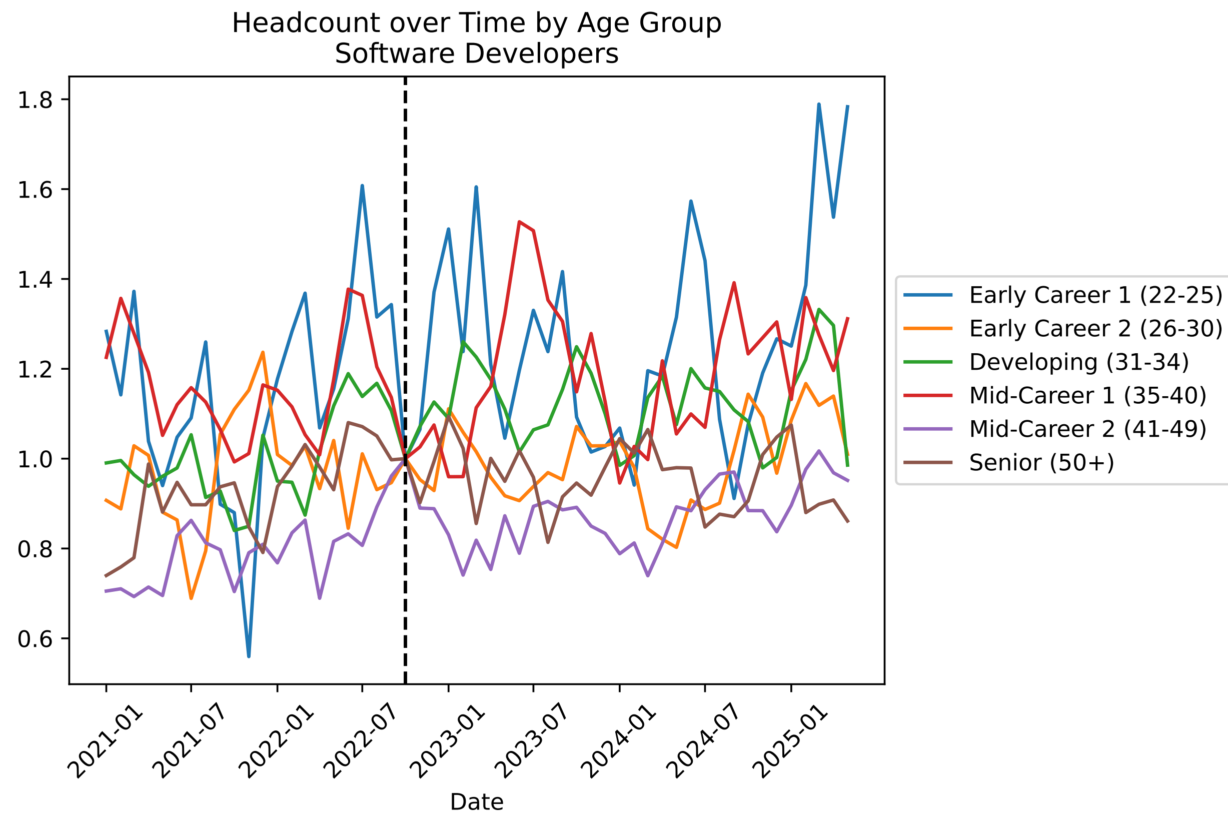

Figure A29: Employment changes for software developers by age, normalized to 1 in October 2022. Data come from the monthly CPS.

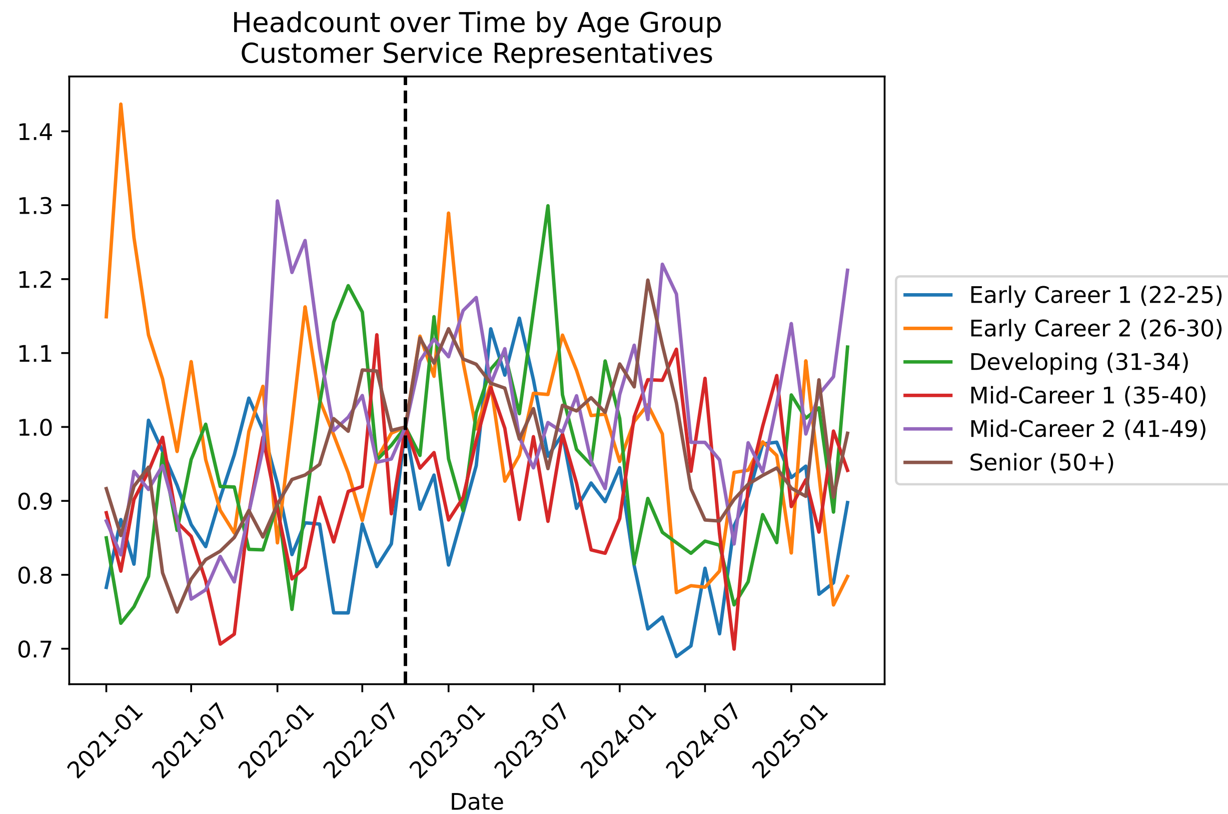

Figure A30: Employment changes for customer service representatives by age, normalized to 1 in October 2022. Data come from the monthly CPS.

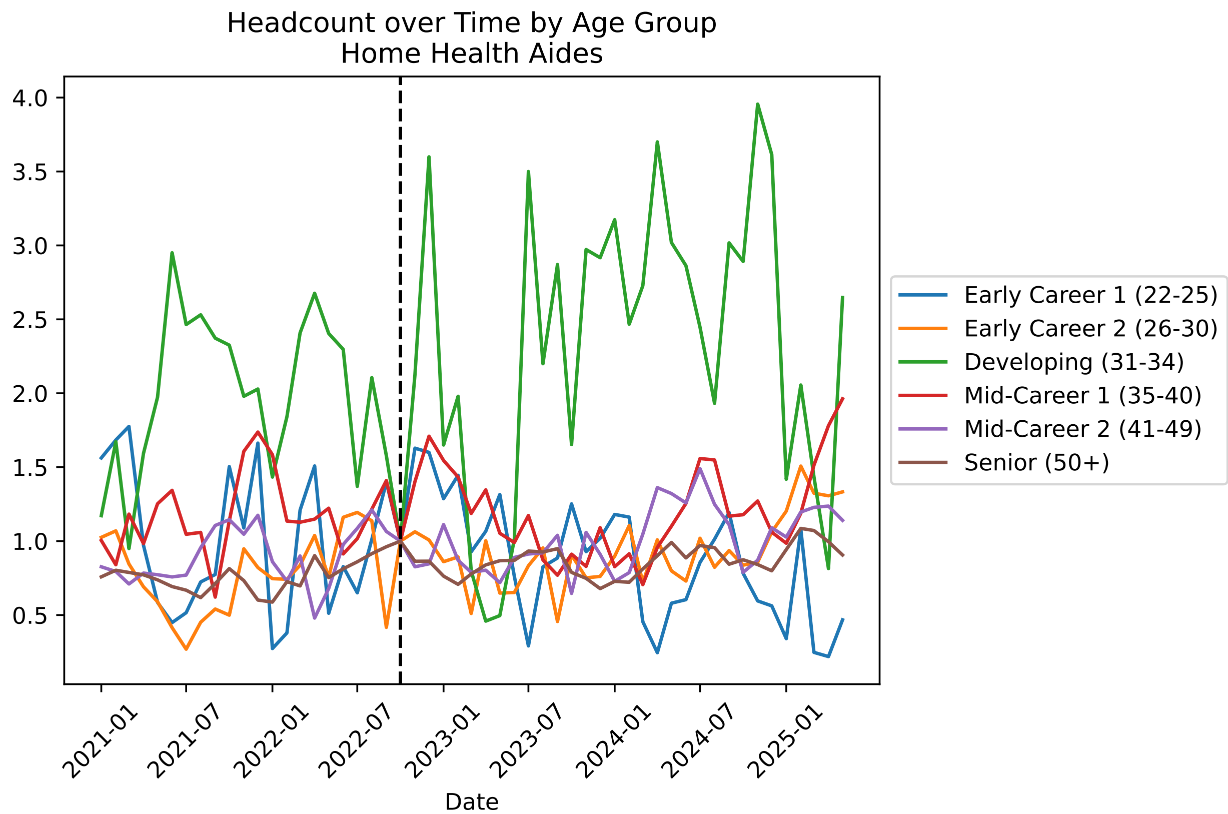

Figure A31: Employment changes for home health aides by age, normalized to 1 in October 2022. Data come from the monthly CPS.

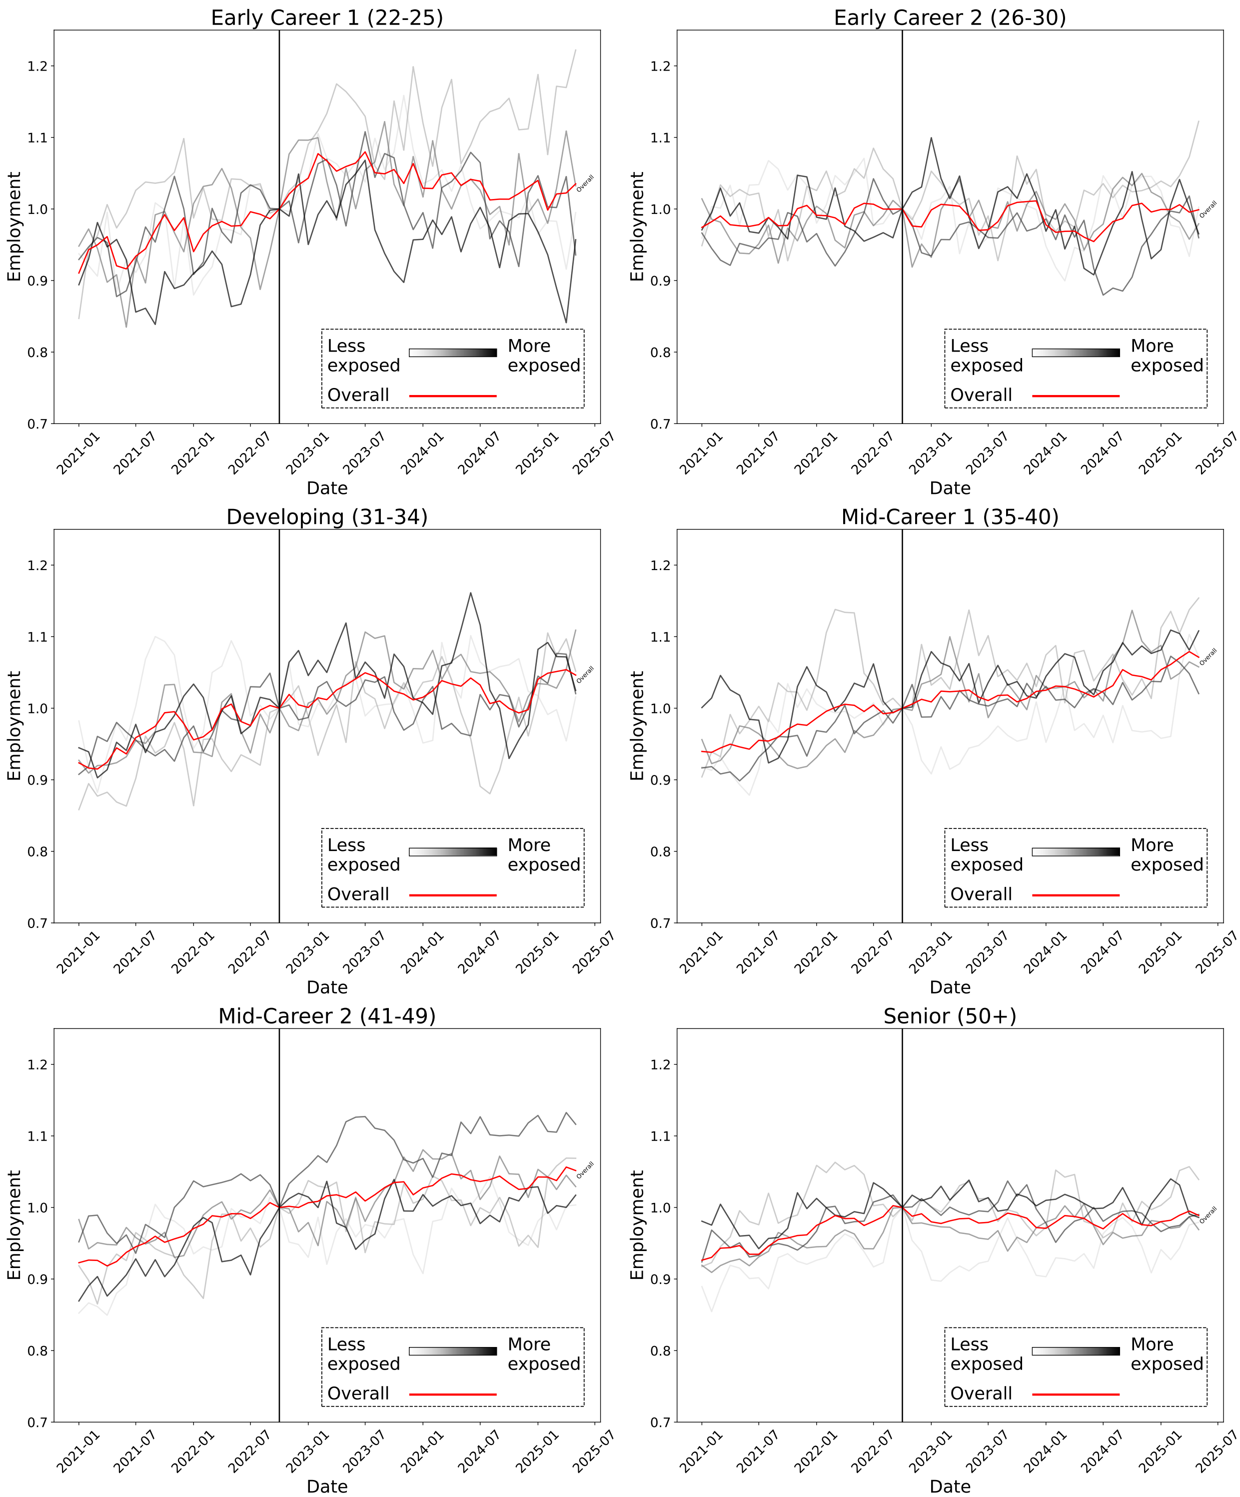

Figure A32: Employment changes by age and exposure quintile using measures from Eloundou et al. (2024). Data come from the monthly CPS.

Table A1: Example occupations by exposure category

Metric	**Least exposed (examples)**	**Most exposed (examples)**
Eloundou et al. (2024) GPT-4 β	• Maintenance and Repair Workers, General • Nursing, Psychiatric, and Home Health Aides • Laborers and Freight, Stock, and Material Movers, Hand • Maids and Housekeeping Cleaners	• Customer Service Representatives • Accountants and Auditors • Software Developers, Applications and Systems Software • Secretaries and Administrative Assistants
Handa et al. (2025) (Overall)	• Taxi Drivers and Chauffers • First-Line Supervisors of Production and Operating Workers • Laborers and Freight, Stock, and Material Movers, Hand • Maids and Housekeeping Cleaners	• Computer Programmers • Financial Managers • Accountants and Auditors • Sales Representatives, Wholesale and Manufacturing
Handa et al. (2025) (Automation)	• Maintenance and Repair Workers, General • Managers, All Other • Nursing, Psychiatric, and Home Health Aides • Driver/Sales Workers and Truck Drivers	• General and Operations Managers • Accountants and Auditors • Software Developers, Applications and Systems Software • Receptionists and Information Clerks
Handa et al. (2025) (Augmentation)	• Cooks • Welding, Soldering, and Brazing Workers • Tellers • Drafters	• Chief Executives • Maintenance and Repair Workers, General • Registered Nurses • Computer and Information Systems Managers

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 A2: 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.

B Additional Literature

Numerous media articles have highlighted the potential employment effects of AI.[28] Some work has noted that the unemployment rate for college graduates has risen above the rate for non-graduates, suggesting this as evidence of employment disruptions from AI (Thompson, 2025). Others have noted that these trends long preceded the spread of AI and have noted that publicly available data such as the Current Population Survey (CPS) show mixed evidence on employment changes in AI-exposed occupations (Lim et al., 2025; The Economist, 2025; Smith, 2025; Eckhardt and Goldschlag, 2025; Frick, 2025).[29] A number of technology executives have also warned of potential job loss from AI (Allen, 2025; Sherman, 2025; Bacon, 2025) or laid off workers with the aim of incresing AI investments (Jamali, 2025).

C Comparison to CPS 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 data to estimates from the CPS.

Figures A29 through A31 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 A32 shows estimated employment changes by age and exposure quintile using the CPS, also suggesting a high degree of volatility in the estimates.

[28] See Horowitch (2025); Ettenheim (2025); Peck (2025); Hoover (2025); Milmo and Almeida (2025); Wu (2025).

[29] Reports from industry have also shown mixed findings. Job posting platform TrueUp suggests a recent increase in postings in the tech sector (Lenny Rachitsky, 2025). On the other hand, Revelio Labs finds a decline in job postings, with the decrease steeper for entry-level workers (Simon, 2025). Indeed job posting data suggest declines in postings for new graduates but find these declines for less AI-exposed occupations as well (Lim et al., 2025). Chandar (2025b) notes that the correlation between job postings and employment has been weak over recent years. SignalFire finds steep declines in new graduate hires in the tech sector compared to pre-Pandemic levels (Doshay and Bantock, 2025), consistent with the findings in this paper. Data from Gusto also suggests a decline in new graduate hiring (Bowen, 2025).

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.[30] 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 data, though the ACS is released with a significant lag compared to the data from ADP. We encourage comparison of our findings to results from other data sources such as the ACS upon their release.[31]

[30]The CPS includes between 26 and 53 young software developers aged between 22 and 25 per month over our sample period. It includes between 49 and 95 young customer service representatives, and between 2 and 14 young home health aides.

[31]The 2024 ACS 1-Year Public Use Microdata Sample is scheduled to be released on October 16, 2025.