You're (not) hired: Artificial intelligence and early career hiring in the Quarterly Workforce Indicators

Lee C. Tucker*

April 17, 2026

Abstract

Using detailed tabulations from matched employer-employee administrative data, I document evidence of an immediate, sizable, and persistent decrease in the level of early career (22-24 year old) hires following introduction of ChatGPT within the industry-state cells that are most exposed to AI. The decline in hires is the primary cause of large observed declines in employment over the subsequent period. Regression-adjusted employment of early career workers in the most AI-exposed quintile of industry-state cells declined by 12% over the 10 quarters following the introduction of ChatGPT, even as employment in less-exposed industries has remained stable. The rate of hiring largely recovered by early 2025, attributable to a smaller employment base. Earnings growth of early career workers in the most exposed industries slowed slightly relative to those in less exposed industries.

Although the most AI-exposed quintile of detailed industries is dominated by a handful of industry sectors, I find that the association of higher AI exposure with reduced early career employment and fewer hires is observed across most sectors of the economy.

Timing of effects in event studies is consistent with an immediate effect on hiring following introduction of ChatGPT. However, triple difference estimates provide some evidence of earlier trend shifts on employment, hiring, and separations around the onset of the COVID pandemic. I discuss potential explanations, including the increase in remote work and increased educational attainment among workers in AI-exposed occupations. Nonetheless, job gains to early career workers and backfill hires show evidence of discontinuous decline at the time of ChatGPT's release in comparison to older workers in the same industries.

A local projections analysis at the NAICS industry group level shows that industries with high AI exposure are not particularly sensitive to unexpected fluctuations in monetary policy on average relative to other industries in employment, hiring, or separations. A historical decomposition suggests that up to one quarter of relative early career employment declines through 2025q2 may be attributable to monetary policy shocks through 2023, but the analysis does not find evidence that these shocks can explain the rapid decline in hires at the most AI-exposed firms in comparison to others.

*Any opinions and conclusions expressed herein are those of the author and do not represent the views of the U.S. Census Bureau. This paper does not contain or use any confidential data. It benefited from discussions with and feedback from Emin Dinlersoz, Lucia Foster, John Haltiwanger, Raza Lamb, Cody Orr, Stephen Tibbets, and Larry Warren.

1 Introduction

It is widely acknowledged that generative artificial intelligence (AI) could have large impacts on labor markets in the future. And, survey evidence has shown that AI adoption in the workplace is both widespread and continuing to increase (Bonney et al., 2024; Bick, Blandin and Deming, 2026). There remains considerable uncertainty about AI's potential future impacts, in part because of uncertainty about the degree to which AI's capabilities could improve (Karger et al., 2026). But, an informative starting point is to identify and characterize the effects it has had to date. A number of recent analyses have taken on this task, using a variety of data sources to measure the effects of AI on a variety of outcomes such as employment, unemployment, job starts, occupation switching, and productivity (e.g., Brynjolfsson, Chandar and Chen, 2025; Massenkoff and McCrory, 2026; Humlum and Vestergaard, 2025; Johnston and Makridis, 2025; Hosseini Maasoum and Lichtinger, 2025). Despite the volume of current research activity and output, a consensus on AI's early impacts remains elusive.

Causal identification of new technologies' impacts is challenging. For digital technologies in particular, low marginal cost and broad distribution can limit the cross-sectional or geographic variation that researchers typically exploit for causal identification. However, AI poses some particular and unique challenges to attempts to measure its effects. The rapid rise of interest in ChatGPT and early consumer adoption following its release in November 2022 has led many researchers to study AI using difference-in-differences methods. These models have some advantages when used to study digital technologies (Goldfarb and Tucker, 2019), but they also make strong identifying assumptions about the lack of contemporaneous confounding factors. Because measured business adoption of AI has been comparatively gradual, many studies make comparisons over multi-year periods, where these assumptions require careful justification to rule out the role of other possible technological changes or trends. Yet, ChatGPT's release immediately followed a global pandemic that upended many labor market patterns. This poses a challenge to typical validation tests upon which researchers might otherwise rely.

In this paper, I document one statistic that appears to have declined 7in immediate proximity to ChatGPT's release, and whose magnitude and persistence is not readily explained by most alternatives: hires of early career workers in the most AI-exposed industries. I find that hires of these early career workers declined immediately by 9% in comparison with those in less exposed industries, and that they have not recovered over time. The overall reduction in hires fully explains a 15% decline in employment in the most AI-exposed industries since large language models became available, representing a loss of over 150,000 early career jobs. As the employment base has declined, the hiring rate has largely recovered, but there is no evidence of catch-up hiring that might cause the most AI-exposed industries to return to their previous size. Both job gains at growing firms and replacement hire activity declined. The rapidity of the decline in hires has no obvious labor supply driven explanation, but is consistent with reduced relative demand for early career workers.

Conversely, using a triple difference regression strategy, I show that more gradual shifts in the relative employment and job flows of young workers in the most AI exposed industries began earlier, at the onset of the COVID pandemic. This suggests that other new technologies such as remote work may be playing a role in reducing the demand for early career workers, or that supply side factors such as increased educational attainment may be delaying labor market entry. I show that working at home and educational attainment are both positively correlated with AI exposure, and in particular that growth in working at home from 2019 to 2021 is strongly associated with AI exposure. This poses a challenge for identifying changes in employment over longer periods – even if generative AI is continuing or complementing existing trends in early career employment, its effects are hard to disentangle from those alternative explanations. These confounding factors make the observation of a rapid reduction in hires more salient, because remote work or changes in labor supply are unlikely to explain such a discontinuous shift in firm behavior. Nonetheless, even in the triple difference setting, evidence shows that job gains, job losses, and backfill hiring all declined discontinuously at the time of ChatGPT's release. In other words, growing firms and firms who had an older worker separate appear to have become less likely to fill those jobs with an early career worker.

Finally, I consider the role of one other shock that might plausibly explain a rapid decline in early career hires leading to a subsequent decline in employment: the rapid tightening of monetary policy following the pandemic recession and recovery. To evaluate this possibility, I fit an impulse response to monetary policy shocks from Bauer and Swanson (2023) using a local projections method (Jordà, 2005). I construct separate response estimates for each detailed industry group and then aggregate them together to see whether differences in monetary policy sensitivity between the most AI exposed sectors and others can explain the large relative decline in hires and the subsequent gap in employment that I observe. I find limited support for this hypothesis. At best, differences in monetary policy sensitivity can explain about one fourth of the regression-adjusted gap in employment, but monetary policy sensitivity does not appear to predict that AI-exposed firms would reduce their early career hiring activity more than others.

A substantial difference of this paper from many recent papers on AI exposure is that it uses information about AI exposure of occupations crosswalked to detailed industry tabulations from the Census Bureau's Quarterly Workforce Indicators (Abowd et al., 2009). Several other researchers have performed such crosswalking exercises to link information about AI exposure to payroll survey tabulations (e.g Council of Economic Advisers, 2024) or QCEW employment and earnings data (Johnston and Makridis, 2025). And, recent research using private payroll processor data (Brynjolfsson, Rock and Syverson, 2017) or matched employer-employee datasets developed from job posting data (Hosseini Maasoum and Lichtinger, 2025) have evaluated employment and hiring patterns at the occupation level, reaching conclusions about weakness in early career hiring that align closely with the findings shown here. However, this is the first to take advantage of the comprehensive information about firm behavior linked to worker demographics in QWI. By crosswalking occupations to industries with public-use ACS microdata, I am able to map external measures of potential AI exposure to detailed industry, state, and age-group tabulations, based on a near-universe of private sector flows.

One goal of this paper is to demonstrate that despite its drawbacks, occupation-to-industry linkage has several benefits in this context. Although workers classified by occupation perform tasks, it is firms in industries that make most decisions about how and when new technologies are adopted in production. This paper demonstrates several ways in which observing the linkage of workers to firms is valuable in studying AI's impacts, even if and when we cannot observe those workers' occupations.

The paper proceeds as follows. Section 2 provides a description of the relevant data sources and linkages, including the crosswalking of occupations to industries; Section 3 provides baseline tabulations that replicate recent findings about early career workers, especially the finding of declining early-career employment associated with AI exposure demonstrated by Brynjolfsson, Chandar and Chen (2025); Section 4 formalizes these comparisons into an event study to compare how early career workers in the most AI-exposed industries differed in employment, labor flows and earnings in comparison to those in less-exposed industries; Section 5 develops a flow-based counterfactual analysis that establishes relative impacts of hires, separations, and compositional effects on changing employment; 6 develops a triple-difference framework that takes advantage of industry-specific scale effects, and provides brief discussion on the role of educational attainment and remote work in interpreting its results; Section 7 perform the local projections analysis to consider the role of monetary policy shocks, and Section 8 concludes.

2 Data sources and crosswalks

Since AI has the capability to perform many tasks that were previously performed using labor inputs, labor market exposure to generative artificial intelligence is typically assessed at the worker occupation level. Researchers have developed a variety of plausible measures to assess potential AI exposure at the occupation level, typically building on detailed information about tasks and work activities from O*NET and then linking this information to assessments on the subset of the tasks for which AI might serve as either a complement to or substitute for human task performance (e.g., Felten, Raj and Seamans, 2021; Eisfeldt, Schubert and Zhang, 2023; Eloundou et al., 2024; Handa et al., 2025; Tomlinson et al., 2025). The potential to directly link between tasks and outputs is an attractive feature of occupational AI exposure measures. However, most commonly used U.S. data sources can either link workers to their employers, or to their occupations, but not both. In this context, an industry-level measure of AI exposure based the typical task composition of the labor inputs they use (i.e., the observed distribution of workers' occupations within industries) is an appropriate alternative way to answer questions about how firms may be responding to the availability of generative AI tools for production.

To construct a detailed industry-level measure, I take an existing occupation-level measure and crosswalk it to detailed industries, largely following the methodology of Johnston and Makridis (2025). For this purpose I begin with the "GPT-4 beta" measure developed by Eloundou et al. (2024). This measure was developed by prompting OpenAI's GPT-4 model to evaluate the capacity of large language models to decrease the time required to complete specific tasks by at least 50%. The evaluated tasks are drawn from the Department of Labor's O*NET database, which is then used to correspond the information about potential AI exposure to individual SOC occupations.

To crosswalk occupations to industries, I use data from the 2015-2019 American Community Survey public-use microdata (ACS PUMS). I construct a person-weighted employment count matrix at the most detailed industry and occupation levels available, and then use this matrix to obtain an employment-weighted mean AI exposure measure for each detailed industry.[1]

One concern when crosswalking from occupations to industries is that much of the available information about potential AI exposure of workers must be discarded; all within-industry variation is lost. Table 1 provides a summary of the extent to which variation in AI exposure is retained or lost with each successive level of crosswalking and aggregation. Note that incremental statistics are not computed when crosswalking from ACS to QWI because these datasets use different employment weights constructed from their respective sources.[2] Most of the information loss associated with aggregation occurs in the initial crosswalk from ACS occupations to industry-states. Comparatively little information is lost by crosswalking from industry-states

[1]Eloundou et al. (2024) provides a measure of AI exposure at the SOC 2018 occupation level, while QWI provides data at the 6-digit NAICS 2022 level. So, SOC 2018 occupations are first crosswalked to ACS 2018 occupation codes using the published Census Bureau crosswalk. Then, ACS 2017 industry codes are crosswalked to ACS 2022 industry codes. Since this intermediate step involves many-to-many linkage, I constructed crosswalk weights by taking industry employment counts from the 2019-2023 5-year ACS PUMS (which is published using 2022 industries), subtracting out a fraction of total employment observed in the 2023 1-year ACS PUMS (also reported in 2022 codes), and comparing the residual employment to the total obtained by summing the 2019, 2021, and 2022 1-year ACS PUMS. Finally, ACS 2022 industries are linked to detailed NAICS 2022 industries from QWI using the published Census Bureau crosswalk; each NAICS 2022 industry is assigned the AI exposure value of the most detailed ACS industry to which it can be linked.

[2]The change in employment weights is also the source of the small weighted variance increase from ACS industry-states to QWI industry-states.

Table 1: Loss of AI exposure information in crosswalking and aggregation

	Weighted parameters		Incremental information loss	
Aggregation	Mean	Variance	MSE	ICC
ACS				
Original data	0.3419	0.0447		
ACS occupation to industry-state	0.3419	0.0148	0.0299	0.3307
QWI				
QWI 6-digit NAICS-state	0.3482	0.0150		
6-digit NAICS-state to 6-digit NAICS	0.3482	0.0146	0.0004	0.9723
6-digit NAICS to 5-digit NAICS	0.3482	0.0146	0.0000	0.9991
5-digit NAICS to 4-digit NAICS	0.3482	0.0145	0.0001	0.9925
4-digit NAICS to 3-digit NAICS	0.3482	0.0118	0.0026	0.8174
3-digit NAICS to 2-digit NAICS	0.3482	0.0102	0.0016	0.8655

Note: Each row shows the employment-weighted mean and variance of the Eloundou et al. (2023) GPT4 AI exposure measure in the specified dataset when exposure is assigned using the weighted mean value for the specified degree of aggregation. Aggregation steps also show incremental change in mean squared error (MSE) and incremental intraclass correlation coefficient (ICC). Employment weights in ACS are person weights of employed persons from 2015-2019 5-year PUMS. Employment weights in QWI are total beginning-of-quarter employment across all states from 2015-2019.

to industries, or by aggregating to NAICS more finely than the 4-digit level. This is likely attributable to two factors. First, the ACS industry classification typically does not map to NAICS industries more precisely than the 4-digit level. Second, the occupational compositions of industries across states may be relatively similar. The loss of information becomes more acute at aggregations coarser than 4-digit NAICS, such as subsector or sector.

In much of the remaining analysis, I map detailed industry and state cells into 5 quintiles of average AI exposure, weighting by total beginning-of-quarter employment (across all age groups) for the four quarters of 2022. Working with quintiles involves a further loss of information about exposure, but it also has some advantages. First, a binary classification simplifies the analysis, making possible simple comparisons among the quintiles and simple event study regression specifications. Second, it avoids relying on implicit assumptions about the functional form of the AI exposure measure being used. Comparisons across quintiles in this paper as well as in recent work (e.g., Brynjolfsson, Chandar and Chen, 2025) suggest that the effects of AI exposure may be concentrated among the most exposed workers; the use of quintiles helps to capture this nonlinearity. Also note that although the core methods of this paper all use the same quintile definitions based on detailed industry by state average exposure, I perform some analyses using coarser industry aggregations (such as the monetary policy analysis of Section 7). In these cases, I recompute AI exposure quintiles on the aggregated categories, again weighted by 2022 all-ages employment from QWI.

Although this analysis uses variation at the most detailed NAICS by state levels available, it can be more intuitive to think about broader industry sectors and how AI exposure maps to them. Figure 2 plots the weighted distribution of AI exposure across detailed industry by state cells within each two-digit NAICS sector. The boxes plot the inter-quartile range and median, while whiskers show the 5th and 95th percentiles of the employment-weighted distribution (using 2022 all-ages employment). The shaded regions represent the quintile thresholds of the AI exposure distribution outlined above. From the figure, one can identify that there are four main sectors of the economy where workers are most likely to be highly AI-exposed. These are

Figure 1: Distribution of AI exposure by NAICS subsector

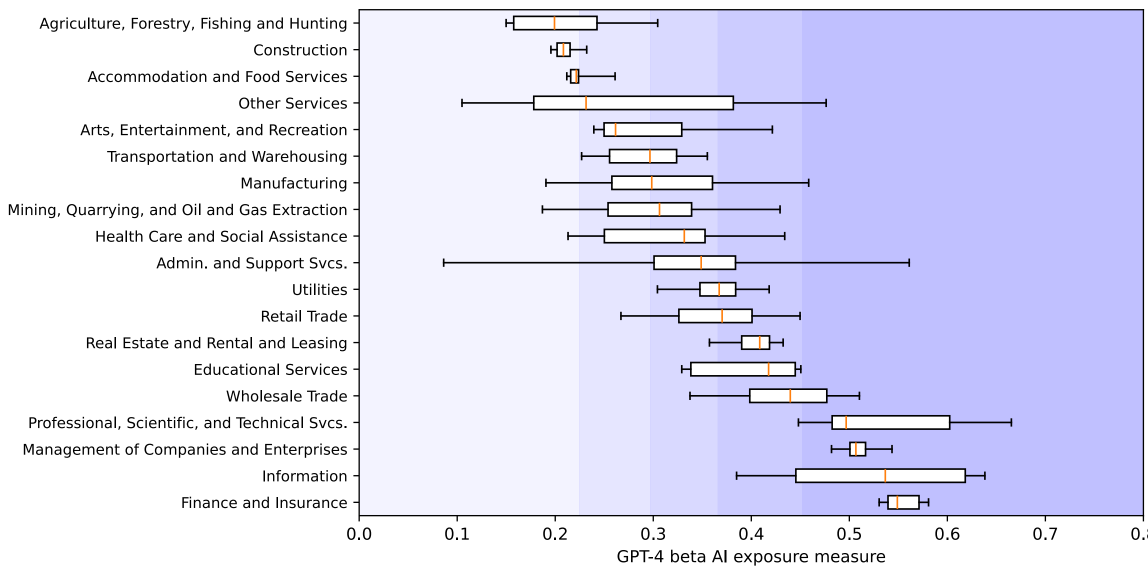

Note: Boxplots show the 5th, 25th, 50th, 75th, and 95th percentiles of 2022 employment-weighted AI exposure of industry-states within each sector. Shading represents employment-weighted quintile thresholds.

Finance and Insurance (NAICS 52), Information (NAICS 51), Management of Companies and Enterprises (NAICS 55), and Professional, Scientific, and Technical Services (NAICS 54). In these four sectors, the median worker is employed in an industry and state that is in the top quintile of industry AI exposure. In three of them, nearly all employment is in the most exposed quintile. However, one should also note that non-trivial fractions of employment are in the most AI-exposed quintile in several other sectors, such as Wholesale Trade (NAICS 42) and Administrative and Support and Waste Management and Remediation Services (NAICS 56).

One challenge inherent in most occupational AI exposure measures, including the GPT-4 beta measure, is that they observe AI exposure indirectly through the tasks that workers typically perform, rather than directly through their actions or place of employment.[3] One key advantage of crosswalking occupational AI exposure to industries is that it allows for simple comparisons of exposure and adoption that can help to validate the measures in use. Figure 2 shows that when crosswalked, the GPT-4 beta measure predicts observed subsector-level data on AI adoption from the Business Trends and Outlook Survey (BTOS) reasonably well. A one standard-deviation increase in subsector AI exposure is associated with a 6.7 percentage point increase in AI adoption. And, approximately 47% of the observed variation in adoption as of April 2026 can be predicted using the GPT-4 beta measure alone[4]. Also of note, the figure demonstrates that the relationship between AI exposure and appears to have steepened over time as subsectors with higher expected AI exposure appear to adopted AI more rapidly.[5]

[3]Work by Handa et al. (2025) has taken significant steps in addressing this concern by providing measures based on AI prompt usage. However, their measure of prompt activity is not based on an observation of user occupations, industries, or employers directly, only the mapping of tasks to occupations.

[4]Note that the BTOS data do suppress response information about AI adoption in subsectors that lack sufficient survey observations, so this finding is not based on a comprehensive set of subsectors. A regression on sector-level data with fewer suppressions obtains a slightly steeper standardized slope coefficient of 7.9 on the latest available data.

[5]BTOS changed its questions about AI adoption slightly beginning in late 2025; and this change may contribute to the increase in slope over time.

Figure 2: Subsector-level AI adoption vs. AI exposure

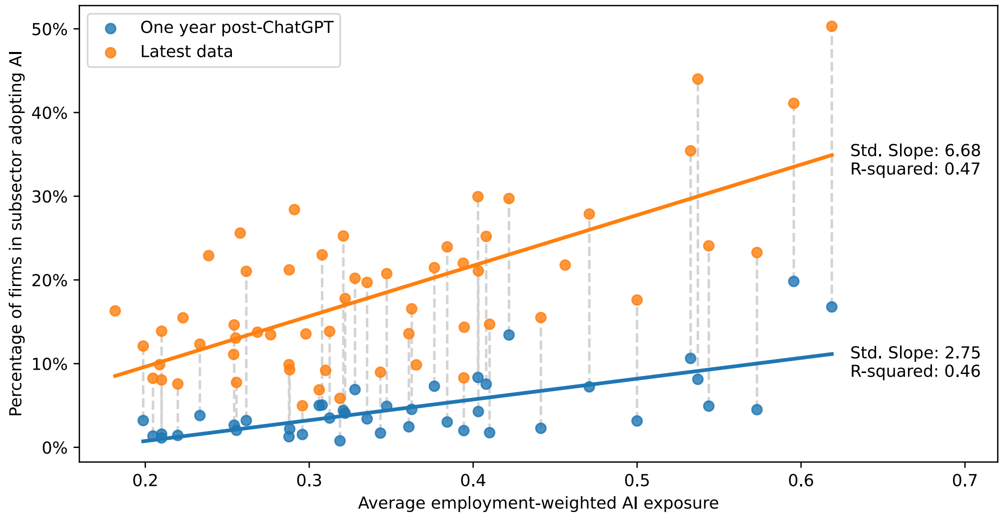

Source: Business Trends and Outlook Survey. AI adoption from BTOS is regressed without weights on subsector-state AI exposure. In both series, detailed industry-state exposure is aggregated to subsectors weighted by 2022 QWI employment. Dotted lines show the change in adoption by subsector between the two periods. The latest data series averages the four releases of BTOS data covering the reference period Jan. 26th though Mar. 22nd, 2026, while the one year post-ChatGPT series averages the four releases covering Oct. 9th through Dec, 3rd, 2023.

In the Appendix, I provide a table showing results for other AI exposure measures and using other crosswalking assumptions. The GPT-4 beta measure appears to be among the best predictors of adoption; other measures developed by prompting LLMs also appear to perform well, including in particular the measure of Eisfeldt, Schubert and Zhang (2023). The range of predicted standardized coefficients is modest. This analysis may be updated as additional measures are compared and as new BTOS data become available. Additionally, ongoing work by Census Bureau researchers to link BTOS data directly to LEHD and other sources of employer and worker information may provide further insights into the relationship between predicted AI exposure risk, AI adoption, and subsequent labor market outcomes.

3 Baseline tabulations of early career employment

The recent work of Brynjolfsson, Chandar and Chen (2025) sparked considerable interest by demonstrating that employment of early career workers in the most AI exposed occupations had declined since the introduction of ChatGPT. Compellingly, their analysis suggests that a trend shift occurred in close proximity to the release of ChatGPT. Since existing analyses have not yet linked firm-level information about adoption to realized worker outcomes, most attempts to infer causal impacts of AI have relied on difference-in-differences or event study strategies. Establishing that the purported effects of generative AI did not begin before the technology became widely available helps to rule out alternative explanations.

In this section of the paper, I replicate and extend these findings in QWI data. Relative to the ADP payroll processing data used in the prior work, QWI has some disadvantages, as well as some advantages. In addition to the lack of occupation, QWI has a lower time frequency, which may mask variation of interest.

The age categories available in public-use tabulations are coarse and do not map exactly to those used by Brynjolfsson, Chandar and Chen (2025). In contrast, the primary advantages are its comprehensiveness – it draws on a near-universe of private-sector employment – as well as the breadth of information that it provides in public-use form.[6] Overall, the findings of this section are broadly consistent with those of prior research, but establishing that these patterns are also observed in QWI is an important first step toward taking advantage of its other features to evaluate mechanisms and robustness.

Figure 3: Indexed employment by AI exposure quintile, ages 22-24

Note: Total beginning-of-quarter employment for each industry exposure quintile from QWI is indexed to 2022q4 employment for the same quintile. Quintile 5 has the highest average AI exposure, and quintile 1 the least. Exposure quintile thresholds are obtained from 2022 all ages employment-weighted industry-state exposure.

Figure 3 provides a comparison of total employment among early career workers across quintiles of industry-state AI exposure. It shows that early career employment in the most AI-exposed quintile of employer declined by 15.2% from the reference period of 2022q4 to 2025q2. All other quintiles either had flat employment or exhibited modest gains over the same period. This result is considerably larger than the 6% decline in employment among the most exposed workers as reported by Brynjolfsson, Chandar and Chen (2025). However, since I also find weaker employment growth among less-exposed quintiles, the unadjusted gap in employment growth is similar in magnitude. For example, although the result of Figure 3 is not regression adjusted, it is highly consistent with the 16% regression-adjusted relative decline in employment that Brynjolfsson et al. observe among the most exposed occupations in ages 22-25.

Replicating the Brynjolfsson et al. result with a comprehensive data source is helpful in two ways. First, it helps to quantify the magnitude of the effect. In the reference period, I observe 6.55 million 22-24 olds employed, and 1.04 million of them (15.9%) are in industry-states that are classified as most exposed. So, the observed decline in employment represents a loss of 159,000 jobs, or 2.4% of total early career employment.[7]

[6]Prior analysis of ADP payroll processor data suggests that it is reasonably representative of the U.S. economy, although it does somewhat overweight large firms and firms in manufacturing and services (Cajner et al., 2018), which might moderately impact unweighted tabulations. The regression-adjusted methods of Brynjolfsson, Rock and Syverson (2017) use within-firm variation that is robust to this concern, but that also cannot capture some sources of employment change such as that occurring with entry and exit of firms.

[7]Since quintiles are based on all-ages employment in 2022, each one represents 20% of overall employment, but early career

Second, it provides additional information on the the timing of the discontinuity. The declining trend in most exposed quintile employment appears immediately between 2022q4 and 2023q1. Since the measure used here is beginning-of-quarter employment, the 2022q4 and 2023q1 observations correspond to observed employment on October 1st, 2022 and January 1st, 2023, respectively. There is some evidence of modest seasonality in these series, which can complicate inference. The regression results shown in the subsequent sections of the paper provide methods that are more robust to aggregate seasonality, industry or state-specific shocks, and similar concerns.[8]

Figure 4: Indexed top-quintile employment by age group

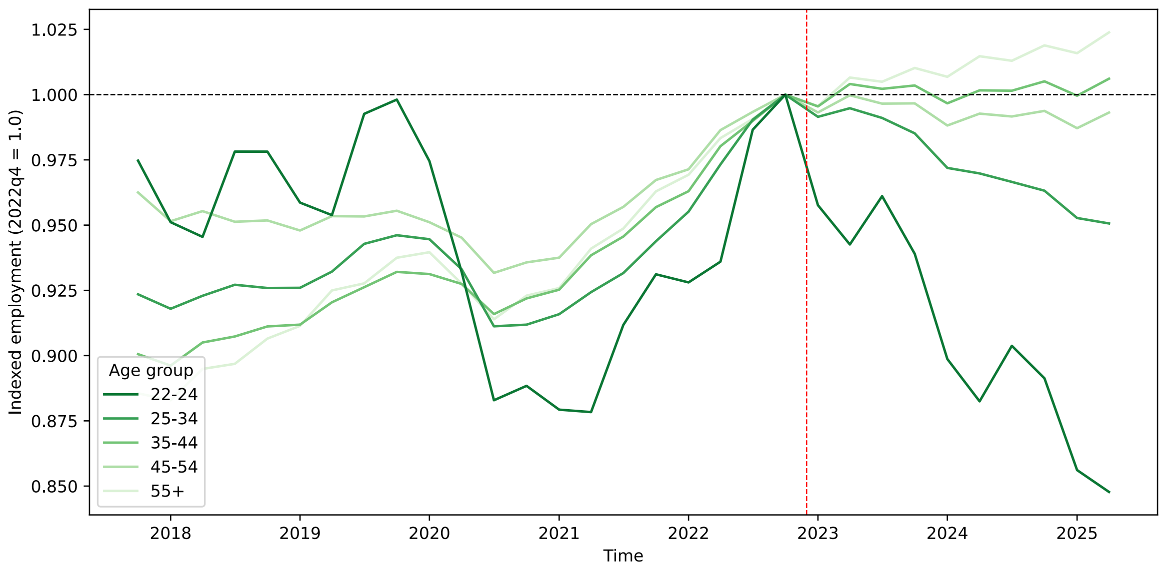

Note: Total beginning-of-quarter employment of the most AI exposed quintile of industry-states from QWI is indexed to 2022q4 employment for the same quintile and age group. Exposure quintile thresholds are obtained from 2022 all ages employment-weighted industry-state exposure.

Figure 4 compares most exposed quintile employment changes for early career workers against those of workers in four age bins of older workers. This figure demonstrates that the decline in employment of early career workers in AI-exposed industries is not generally observed among older workers. Only the 25-34 age group saw a substantial decline in employment over this period, but in percentage terms it was less than one third as large as that observed among ages 22-24. The pattern of employment changes across age groups is not monotonic with age, but it is suggestive of lower employment among early career workers defined broadly. This analysis is drawn from publicly available tabulations by coarse age bins, so I am unable to look more finely at the relationship between age and employment. However, future work with restricted-use microdata may be able to provide more fine-grained detail regarding this relationship.

Figure 4 also suggests that the choice of reference period may matter considerably for interpretation. This is largely because early career workers saw the largest percentage employment declines at the onset of the COVID pandemic. So, for example, if employment were indexed to the quarters immediately preceding the COVID pandemic, the relative employment decline of early career workers in these industries would

workers are somewhat underrepresented in the top exposure quintile and are overrepresented in the least AI-exposed quintile.

[8]Employment in QWI is inferred from UI wage records, which provide earnings in each quarter. Employment on quarterly boundaries is inferred from the existence of positive earnings in consecutive quarters. Use of this measure ensures that employment totals can be treated as a stock rather than a flow.

appear considerably larger. I will return to this point when discussing triple difference regressions in Section 6.

Figure 5: Distribution of changes in early career employment by sector

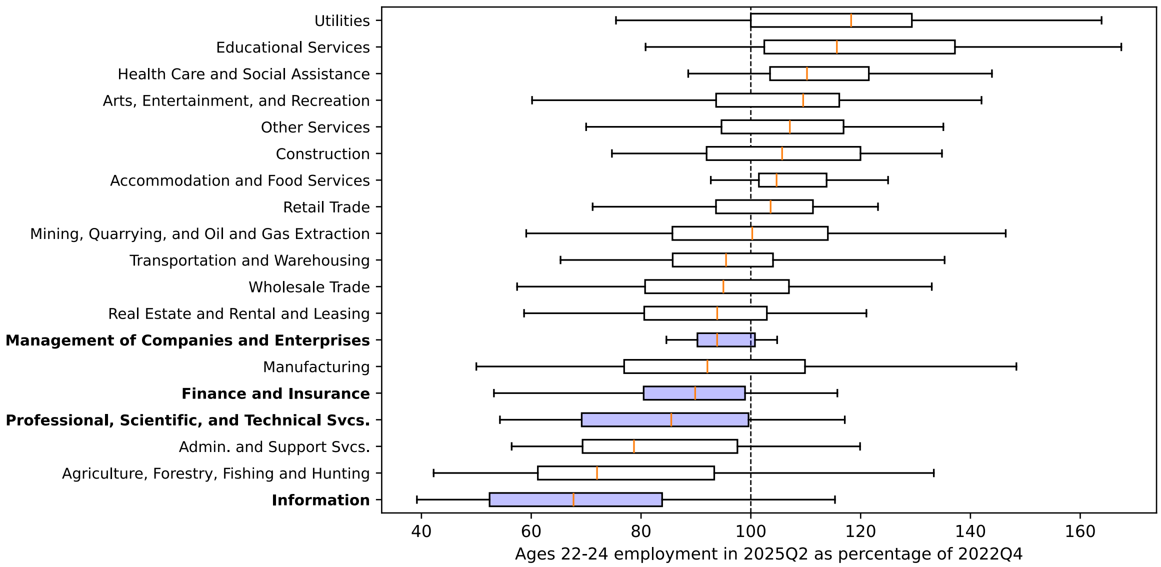

Note: Boxplots show the 5th, 25th, 50th, 75th, and 95th percentiles of employment-weighted indexed employment change between 2022q4 and 2025q2 within each sector, for ages 22-24.

One concern when looking at aggregate tabulations is that results may be driven by a small handful of observations (i.e., industries), obscuring what might be more informative or relevant patterns. To evaluate this possibility, Figure 5 provides boxplots of percentage employment declines by NAICS sector. As with Figure 3 above, the quintiles shown here are weighted by 2022 employment, So, for example, the industry median corresponds to the median change in employment observed by a worker for their detailed industry-state. The four sectors identified in Figure 1 are highlighted. In all of these sectors, roughly 75% or more of employment was in detailed industry-states whose employment declined over the period. Early career employment in Information declined by 30%, but Finance and Insurance and Professional, Scientific, and Technical Services saw double-digit percentage declines as well (-11% and -15%, respectively).

Differences in AI exposure are a reasonably good predictor of relative weakness in employment and hiring within sectors as well. Table 2 provides summary statistics by sector. The first three columns tabulate the contributions of each sector to most exposed quintile employment. A majority of the most AI-exposed employment is in two sectors – Finance and Insurance, and Professional, Scientific, and Technical Services – but a majority of sectors have some portion of detailed industry-states that are in the top quintile. The last four columns of the table show the result of within-sector regressions of the change in indexed employment or hires since the reference period on either a one standard deviation increase in the AI exposure measure, or on a binary indicator of most exposed quintile status. All regressions are weighted by 2022 employment. Because of seasonality in early career hires, the hires index compares the last four quarters of available hires data to the four quarters preceding the introduction of ChatGPT.

The table shows that even in industries with low percentages of top-quintile employment, high AI exposure is usually associated with reduced employment and fewer hires over this period. The Finance and Insurance sector is a notable exception; all of its detailed industries are in the most AI-exposed quintile, but the largest

Table 2: Sector-level summary and AI exposure relationships

	% of total top quintile employment	% of sector emp. in top quintile		Diff. in indexed employment		Diff. in indexed hires	
Sector	Ages 22–24	Ages 22–24	All ages	One S.D. exposure	Top quintile	One S.D. exposure	Top quintile
54: Professional, Scientific, and Technical Svcs.	40.0%	89.5%	93.9%	-0.073***	-0.106***	-0.070***	-0.087***
52: Finance and Insurance	20.3%	100.0%	100.0%	0.108*	None	0.120**	None
55: Management of Companies and Enterprises	9.4%	98.0%	98.1%	-0.014	-0.027	0.002	-0.079**
51: Information	8.8%	73.2%	71.1%	-0.110***	-0.110	-0.071***	-0.061
42: Wholesale Trade	7.3%	37.2%	42.0%	-0.130***	-0.103***	-0.063***	-0.055***
56: Admin. and Support Svcs.	5.0%	9.8%	9.7%	-0.062***	-0.140***	-0.056***	-0.128***
44-45: Retail Trade	4.4%	4.3%	3.9%	-0.106***	-0.077***	-0.007	-0.020
31-33: Manufacturing	2.3%	4.9%	6.2%	-0.042***	-0.204***	-0.046***	-0.150***
81: Other Services	1.2%	5.5%	8.0%	-0.018	-0.222***	-0.010	-0.079***
62: Health Care and Social Assistance	0.3%	0.4%	0.3%	-0.062***	-0.128***	-0.110***	-0.119***
71: Arts, Entertainment, and Recreation	0.3%	2.1%	3.1%	-0.051	-0.213***	-0.054**	-0.119
61: Educational Services	0.3%	2.6%	3.1%	0.275***	0.097	0.067	-0.067
53: Real Estate and Rental and Leasing	0.1%	1.5%	1.0%	0.019	0.053	0.045	0.096
48-49: Transportation and Warehousing	0.0%	0.0%	0.0%	-0.060	None	0.096	None
22: Utilities	0.0%	0.0%	0.1%	0.098	None	0.045	None
21: Mining, Quarrying, and Oil and Gas Extraction	0.0%	0.0%	0.0%	-0.029	None	-0.014	None
11: Agriculture, Forestry, Fishing and Hunting	0.0%	0.0%	0.0%	0.262***	None	0.061	None
72: Accommodation and Food Services	0.0%	0.0%	0.0%	-0.201***	None	-0.063	None
23: Construction	0.0%	0.0%	0.0%	-0.364***	None	-0.177*	None

Note: Each reported difference in employment associated with a one standard deviation increase in AI exposure or from top quintile status is from a regression of 2025q2 beginning-of-quarter employment as a percentage of 2022q4 on mean industry-state AI exposure or a binary indicator of top quintile status. Each reported difference in hires is from a regression of average 2024q2-2025q1 hires, expressed as a percentage of 2021q4-2022q3 hires, on the same exposure/quintile measure. Asterisks represent significance at the 90%, 95%, or 99% level with robust standard errors.

employment declines were in the least exposed portions of this sector. Educational Services also consistently counters the overall trend. Nonetheless, these results suggest that the association of higher exposure with lower employment is reasonably robust.

4 Event study

The initial tabulations shown in the preceding section show that a large decline in early career employment was concentrated within the most AI-exposed quintile of industries. In this section, I develop an event study analysis that formally tests the extent to which outcomes for the most exposed quintile have changed. Regressions allow for statistical inference, and they also allow for analysis of outcomes such as hiring, separations, and earnings that are difficult to evaluate in tabulations because of strong and heterogeneous seasonal patterns across exposure quintiles. Additionally, because this regression controls flexibly for common time-specific shocks, it nets out the baseline effects of the COVID pandemic before the reference period, simplifying the comparison with pre-trends.

To perform the event studies, I run regressions of the form:

$$y_{i,s,t} = \alpha_{i,s} + \gamma_t + \sum_t \beta_t \text{topquint}_{i,s} + \epsilon_{i,s,t} \tag{1}$$

Here, $y_{i,s,t}$ represents an outcome of interest for detailed industry i and state s at time t. Outcomes observed include logs of beginning-of-quarter employment, hires, separations, and average monthly earnings growth.[9] The fixed effects $\alpha_{i,s}$ and γ_t control flexibly for time-invariant industry-state characteristics (e.g., average industry-state employment) and for common shocks that affect all industry-states in each period, including common seasonal patterns shared by the age group, the shared effects of the COVID pandemic, and the shared impact of economic shocks. The coefficients of interest β_t measure the extent to which the patterns of the outcome variable in the most exposed industry-states differed from overall trends, net of these cells' time-invariant characteristics. Since all outcomes are measured in logs, the estimated coefficients can be roughly interpreted as a percent of decline or growth in the measured outcome in comparison to the reference period.

Regression samples in this section are specific to a particular age group, typically ages 22-24. The reference quarter is either 2022q4 or 2022q3, depending on the measure – it is chosen such that the last quarterly boundary implied by the covariate in the reference quarter is always the one spanning quarters three and four. So, 2022q4 is used as the reference quarter when beginning-of-quarter employment or separations are measured, because these imply employment on October 1st, 2022. Conversely, when measuring hires that result in end-of-quarter employment or monthly earnings, 2022q3 is used as the reference quarter so that the relevant quarterly boundary occurs on September 30th, 2022, before ChatGPT's release. All quarterly figures are plotted so that the points land on the associated boundary dates as well. Standard errors are two-way clustered by detailed industry and state.

Figure A.6 shows the results of a regression on beginning of quarter employment. The results are consistent with the earlier tabulation results shown in Figure 3, with a clear trend shift that appears to begin immediately following the reference quarter. Despite the sizable effects of the COVID pandemic

[9] Specifically, I use measures of "end of quarter" hires (hires within the quarter that remained employed as of the next quarterly boundary) and "beginning of quarter" separations (separations that were employed as of the last quarterly boundary) to limit to hires and separations that result in measured employment. For earnings growth, I use the four quarter change in log average monthly earnings among stable (full quarter) employed workers.

Figure 6: Event study regression: Early career log employment

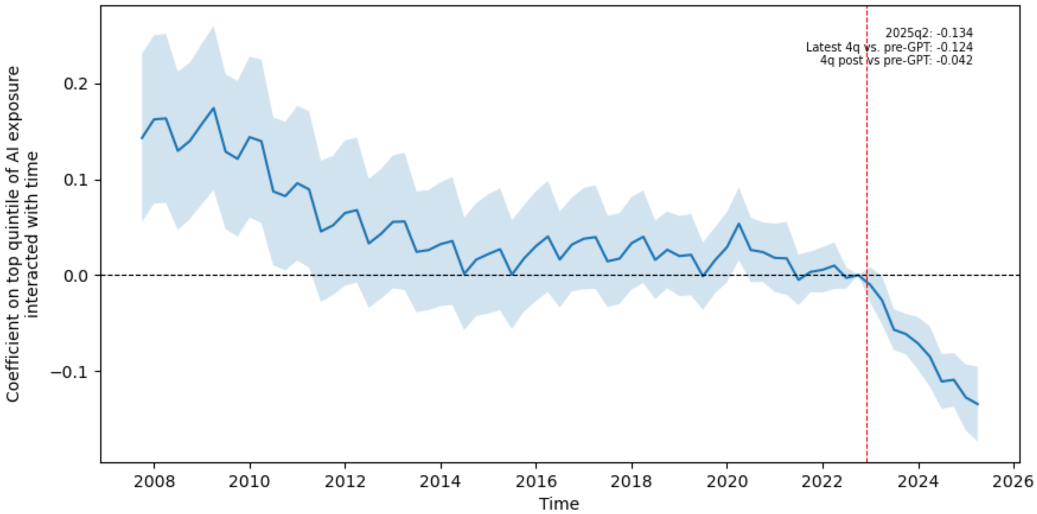

Note: Each plotted coefficient is estimated from an event study regression of log beginning-of-quarter employment of ages 22-24 on a period by most-exposed quintile indicator. The reference period is 2022q4. Text at top right shows the coefficient for the latest available period, and the last four quarters compared to the four quarters ending in the reference period. Standard errors are two-way clustered by detailed industry and state.

on employment, the pre-trend here is largely flat, implying that the pandemic's effects on young worker employment in top-quintile industries was similar to its effect on young workers in other sectors. Pre-trends alone are not sufficient to establish causality, since we must also rule out possible contemporaneous shocks. Nonetheless, on its face, the timing and strength of this shift appears to be consistent with an immediate causal impact of generative AI.

This event study specification controls for commonly-felt seasonal shocks, but the figure shows evidence of residual seasonality (i.e., seasonal patterns in the most exposed quintile that differ from the overall average pattern among young workers).[10] And, the last quarter of observation is not the same calendar quarter as the reference period. For these reasons, my preferred estimate of the current effect compares the point estimates for the last four quarters to the four-quarter window ending with the reference quarter. This method implies a 12.4% decline in employment from the release of ChatGPT to 2025q2, which is a bit smaller in magnitude than that suggested by Brynjolfsson et al., and also a bit smaller than the unadjusted 15% decline implied by the unadjusted tabulation.

In Figure 7, I look at the two primary mechanisms by which employment changes over time; hires and separations.[11] Despite considerable residual seasonality in the series, there is evidence of a sharp, immediate and persistent decline in hires at the time of ChatGPT's release of roughly 9 percent. This decline in hires (top left panel) has not recovered over time; hires over the last four quarters of available data are 12.7% below the corresponding reference quarters on average. A regression on the log of the hiring rate (top right panel) shows a similar sharp and immediate decline, followed by a gradual increase as the employment denominator

[10] Interestingly, the figure also appears to show that residual seasonal patterns in employment have declined in magnitude over time.

[11] Because we are looking at employment within age group, state, and industry bins, there are additional mechanisms in the form of aging and industry changes; I will discuss their impact later in the paper. Because of how QWI data are collected, firm relocations across state boundaries are coded as separations and hires.

Figure 7: Event study regressions: Hires and separations of early career workers

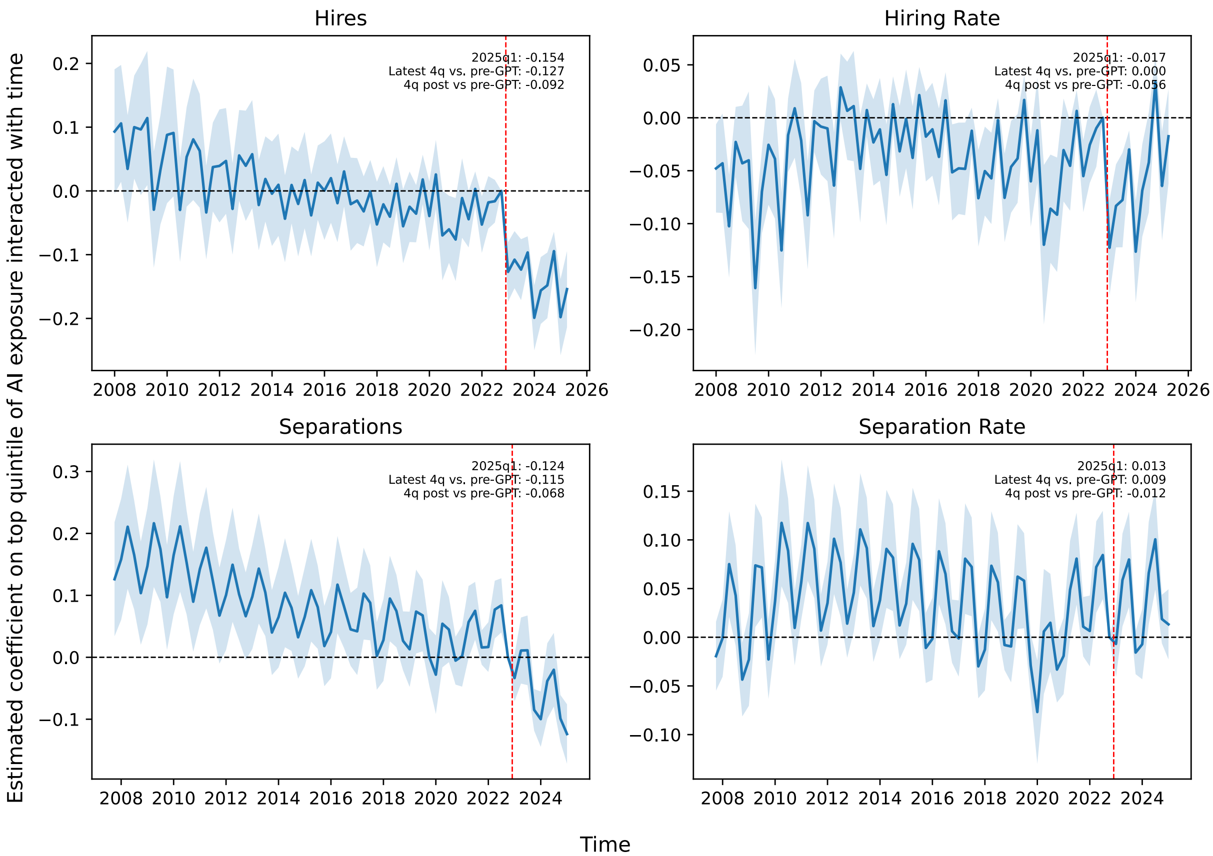

Note: Each plotted series is estimated from a separate event study regression of the indicated outcome for ages 22-24 on a period by most-exposed quintile indicator. The reference period is 2022q3 for end-of-quarter hires and hiring rate, and 2022q4 for beginning-of-quarter separations and separation rate. Text at top right shows the coefficient for the latest available period, and the last four quarters compared to the four quarters ending in the reference period. Standard errors are two-way clustered by detailed industry and state.

shrinks.[12] As of 2025q1, hiring has nearly recovered to its previous rate, but the number of hires has not.

The series for separations exhibits even more residual seasonality, but there is again evidence of a decline immediately following introduction of ChatGPT. As of 2025q1, separations have declined 11.5% in comparison with the four quarter reference. Note that there is a gradual downward trend in separations before the COVID pandemic, followed by a modest uptick before the reference period. One plausible explanation for the latter trend is that workers in the most exposed quintile were slightly more likely to resign during the supposed "Great Resignation" period. These two offsetting trends could cause the standard point estimate to overstate the magnitude of a presumed negative effect of AI on job separations. Expressed as a separation rate, the immediate effect is negative but less pronounced, suggesting that the decline in the level of separations may be in large part driven by the decline in employment.

In the Appendix, I provide figures showing event studies of employment, hires, and separations separately by bins of firm age and size. Because of concerns about small cell size suppression in QWI, these analysis are performed at the NAICS industry group (4-digit) level, but all-ages and all-sizes results are very consistent with the baseline analysis. Overall, it appears that declines in hiring and employment were somewhat more pronounced among medium and large firms, and a bit smaller in the newest firms. However, the pattern of immediate decline beginning at the reference period is broadly observed.

One valuable feature of the QWI published tabulations is that they link hiring and separation activity of workers to the corresponding changes in firms. This allows us to take advantage of some information that exists in the underlying linked microdata without requiring access to firm-level records. The three measures shown in the panels of Figure 8 provide additional information about the types of hiring the separation that changed over time. Panel A shows job gains, or employment increases at firms that grew over the course of the quarter due to hiring. Job gains represent about two thirds of early career hires in the most exposed quintile, with the remainder being categorized as replacement hires. As with overall hires, the series shows evidence of a sharp discontinuity aligning with the release of ChatGPT. The 9.1% immediate reduction in job gains in the first four post-ChatGPT quarters is similar in magnitude to the decline in overall hires over the same period.

The second panel shows job losses, or employment declines at firms that shrank in size over the course of the quarter. These represent just over half of early career separations in the most exposed quintile. This series shows evidence of decline over time, but not an immediate discontinuity, similar in magnitude and shape to the decline in employment over time. Since these series estimate counts and not rates, this implies that workers in the most exposed quintile have had relatively little change in their probability of separation due to job destruction in comparison with less exposed workers. That is, the decline in job losses is mostly the result of fewer jobs to lose, arising from less hiring.

The bottom panel of Figure 8 shows the decline in replacement hires, again expressed as a level. As with job gains and hires overall, this shows evidence of an immediate decline of about 9%. However, in more recent quarters, replacement hiring has declined a bit more than hiring overall; the last four available quarter suggest that that replacement hires have declined by about 15%.

To complete the event study analyses of this section, Figure 9 compares how average monthly earnings of early career workers in the most AI-exposed industries grew in comparison with those of other early career workers. It is clear from the figure that AI-exposed workers saw stronger relative earnings growth at the onset of the COVID pandemic, and weaker earnings growth during the subsequent "Great Resignation" period.

[12] I follow published QWI statistics in calculating the hiring rate based on the average of beginning- and end-of-quarter employment.

Figure 8: Event study regressions: Job gains, job losses, and replacement hires of early career workers

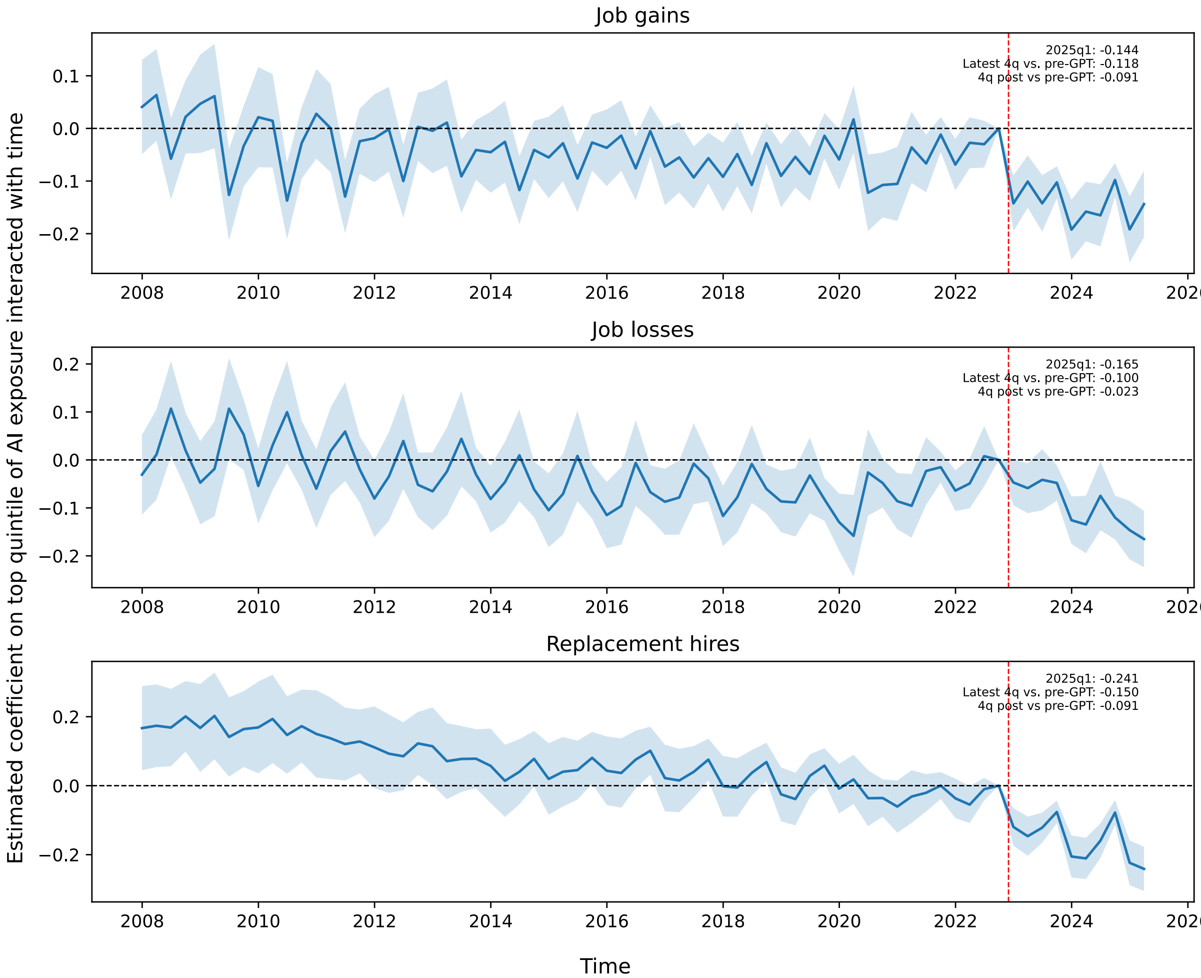

Note: Each plotted series is estimated from a separate event study regression of the indicated outcome for ages 22-24 on a period by most-exposed quintile indicator. Job gains, losses, and replacement hires are measured over a quarter, so the reference period is 2022q3 for each series. Text at top right shows the coefficient for the latest available period, and the last four quarters compared to the four quarters ending in the reference period. Standard errors are two-way clustered by detailed industry and state.

Figure 9: Event study regressions: Average earnings growth of early career workers

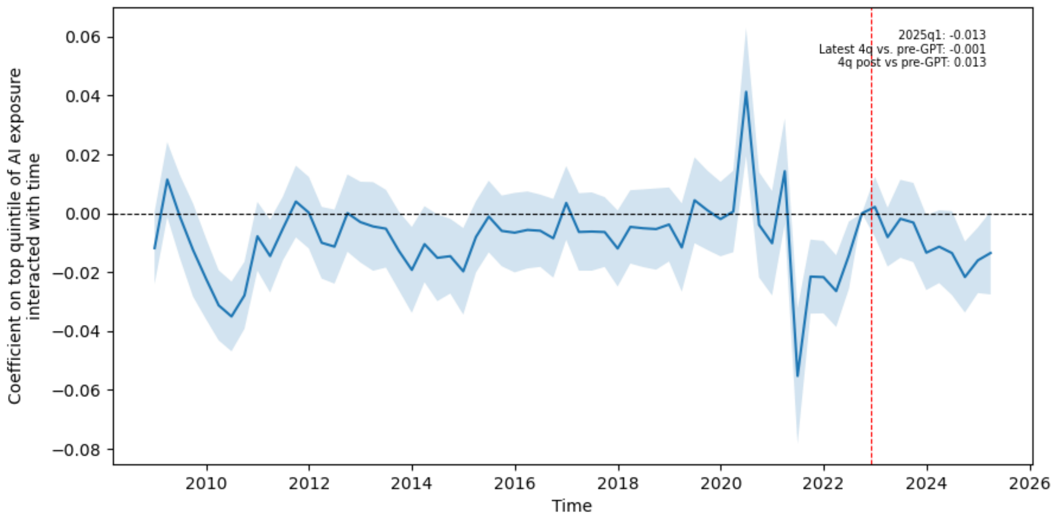

Note: Each plotted coefficient is estimated from an event study regression of the four-quarter change in log average monthly earnings for full-quarter employed workers aged 22-24 on a period by most-exposed quintile indicator. The reference period is 2022q3. Text at top right shows the coefficient for the latest available period, and the last four quarters compared to the four quarters ending in the reference period. Standard errors are two-way clustered by detailed industry and state.

Notably, average earnings growth of early career workers in these industries also lagged for several quarters following the Great Recession, before experiencing a period of stability during the subsequent expansion. Since ChatGPT's release, it appears that earnings growth has slowed modestly among early career workers in comparison to their peers. However, the decline is small relative to cyclical patterns.

Although the comparison to other early career workers is informative, a natural interpretation of Figure 9 is that workers in the most AI-exposed industries may experience cyclical fluctuations differently than those employed elsewhere. This also suggests that event study estimates might be capturing the effects of correlated industry-specific shocks. I will return to this concern in Section 6.

5 Counterfactual flows-based analysis

The previous section showed evidence of a substantial decline in hires since the introduction of ChatGPT, as well as a corresponding decrease in separations. These have offsetting impacts – fewer hires reduces total employment, while fewer separations increases it. A counterfactual analysis makes it possible to quantify the magnitude of these countervailing changes to assess the contribution of each, as well as the contribution of other compositional factors that may have an impact on the path of overall employment.

At the firm level, the relationship between employment, hires, and separations is a simple identity. However, when tabulating employment across industry and age group cells over time, the relationship must also take into account the aging of workers, changes to the NAICS codes reported by establishments, and imputations that could lead employment totals to change for reasons other than hiring or separation. That is, the change in employment is:

$$Emp_{g,t+1} - Emp_{g,t} = Hires_{g,t} - Separations_{g,t} + \underbrace{AgeIn_{g,t} - AgeOut_{g,t} + Misc_{g,t}}_{Composition_{g,t}} \quad (2)$$

With data on employment, hiring and separations in each period, we can infer the collective contribution of compositional effects.[13] And, by tabulating all ages employment and flows in QWI, it is easy to verify that on average the collective contributions of $Misc_{gt}$ are typically small and mean 0. So, we can infer that the compositional effect is mostly the result of net aging of workers into or out of the age bin.

Figure 10: Average magnitude of compositional adjustments over time

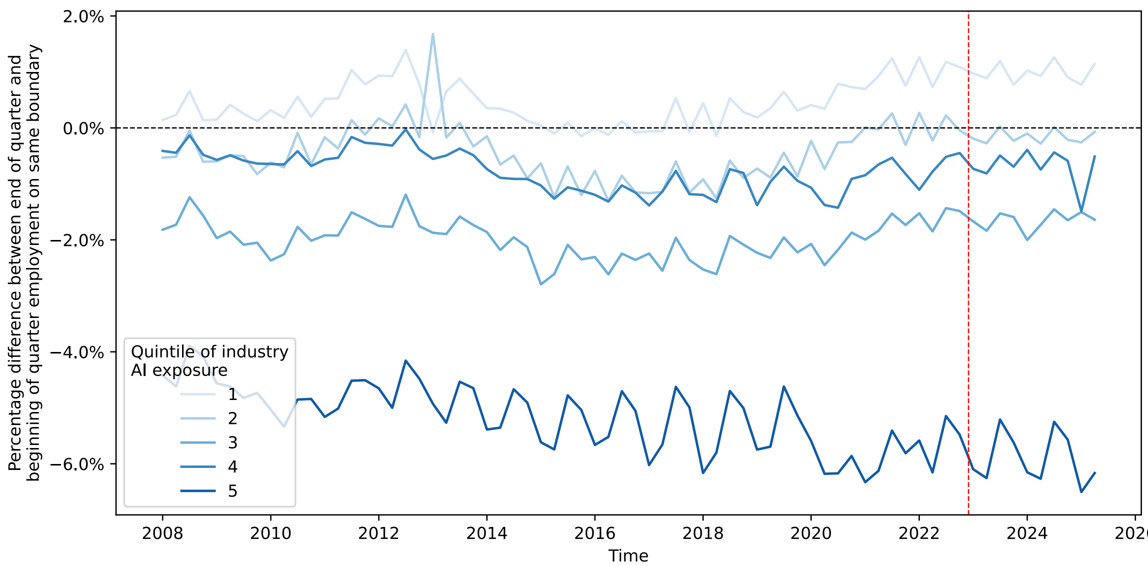

Note: Each series shows the difference between end-of-quarter employment in one quarter and beginning-of-quarter employment in the next, expressed as a percentage of end-of-quarter employment of for ages 22-24, separately for each quintile of industry-state average AI exposure.

Figure 10 shows the magnitude of compositional effects for early career workers in each AI exposure quintile, by plotting the average percentage by which beginning-of-quarter employment differs from what hire and separation flows would predict. Positive values are associated with net aging into the age bin (i.e., more 22-year-old workers whose jobs continued from the previous quarter), while negative values are associated with net aging out (i.e., more workers turning 25 and aging out). The most AI-exposed quintile stands out; the gap between workers aging in and workers aging out is five percent or more of employment in each quarter. There is also evidence of a gradual long run trend toward more aging out than aging in in the most exposed quintile, but not in the other quintiles. The most likely explanation for this is delayed labor market entry.

Whether delayed labor market entry is reflective of supply-led or demand-led changes is less clear. One plausible explanation is that early career individuals who might work in AI-exposed industries are increasingly delaying their entry into the labor market to obtain additional education, or alternatively that age at college graduation has increased. As discussed later in Section 6, workers in the most AI-exposed occupations have

[13] Throughout this paper, I use the published QWI measures *HireAEnd* and *SepBeg*, which are hires in the quarter resulting in end-of-quarter employment and separations after beginning-of-quarter employment. QWI also publishes *EmpEnd*, or end-of-quarter employment, which incorporates hires and separations over the period but not compositional effects on the quarterly boundary. I use the ratio of $EmpEnd_t$ to Emp_{t+1} to measure the magnitude of compositional effects.

higher average educational attainment, and the most exposed quintile has also seen its rate of educational attainment grow more quickly over time. An alternative explanation is that workers are delaying labor market entry due to weak labor demand.

Delayed labor market entry has both an immediate effect on tabulated hires and employment and an offsetting deferred effect in either future hires or foregone future compositional adjustments. With coarse age bins and no direct information about educational attainment, my ability to evaluate the role of changes in educational attainment or other labor supply factors is limited in these data. However, there is no evidence of a discontinuity or trend shift that might explain a large portion of the change in employment on its own. For the purposes of the counterfactual analysis, I assume that the magnitude of composition adjustments, expressed as a percentage of employment, remains constant over time. Whether such a counterfactual compositional adjustment reflects conditions of greater labor supply or stronger labor demand is less clear.

Figure 11: Flows-based counterfactual analysis

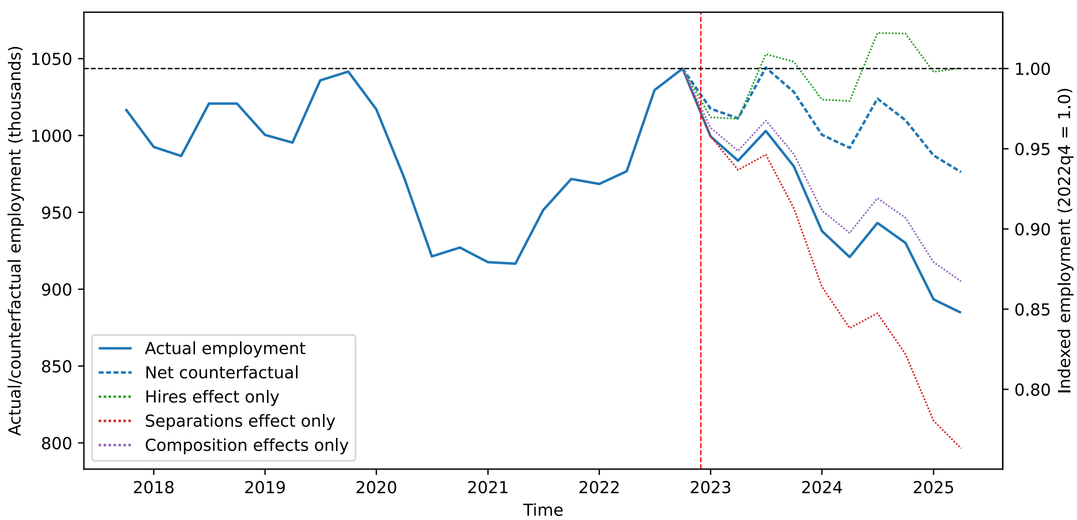

Note: Each dotted line shows the net effect of one counterfactual channel on the path of realized employment for ages 22-24 workers in the most AI-exposed quintile of industry-states. The net counterfactual series enables all three channels.

Figure 11 shows the result of the counterfactual exercise. The dashed blue line represents the net counterfactual incorporating all three effects (counterfactual hires, separations, and compositional effects). To account for residual seasonality in regression estimates, I use the gap between each quarterly regression point estimate and the point estimate for the same calendar quarter that immediately precedes the reference quarter. The net counterfactual series shows that counterfactual trends in hires, separations, and compositional effects would have led to a smaller employment decline of 67,000 jobs, or about 40 percent of the size of the 159,000 job decline actually observed in the data. That is, to the extent that weaker labor market conditions were shared among the most AI-exposed firms and others, they would have predicted a modest but substantially smaller decline in early career employment.

The thinner dotted lines represent the effect of enabling only one counterfactual channel at a time, while allowing the other two series to take on their true value. Most notable of these is the series on hires. If the counterfactual level of hires had been maintained, early career employment in early 2025 would be about the same as it was before the ChatGPT's release, at 1.04 million. The decline in hires explains 100% of the loss in employment. If separations had followed their previous (higher) counterfactual, then employment would have declined by an additional 88,000 jobs. Compositional effects had a downward effect on measured employment over this period, presumably because of increased net aging, which may be the result of delayed labor market entry. However, the magnitude of this effect is small in comparison to the direct effects of fewer hires and separations.

Overall, this counterfactual exercise confirms the intuition of the preceding regression estimates. Employment declines in the most AI-exposed industries were attributable to the decline in hires. If not for an offsetting decline in separations, employment losses would have been even larger.

6 Triple-difference event study

The event study figures of the preceding section make it clear that employment, hires, and separations, and earnings of early career workers all declined in top-quintile industries in comparison to less-exposed industries. Because of time-period fixed effects, these regressions control for the common component of economic shocks such as the COVID pandemic, changes in monetary policy, and changes in aggregate labor supply of early career workers. However, these regressions do not account for the possible role of time-specific shocks to industries or firms that might explain a decline in hires and subsequent employment. If, for example, the AI-exposed industries that hire the most early career workers were affected by a negative product demand shock, then we would expect to observe a decline in hiring of early career workers in the most exposed quintile, even if that demand shock was unrelated to the availability of generative AI.

To evaluate this possibility, I extend the regression analysis using a triple-difference strategy that leverages a simple observation: demand-driven increases and decreases in employment, hiring, and separations tend to affect both early career workers and older workers through scale effects. At times of strong demand, firms increase their hiring and employment of both early career workers and older workers, and conversely in times of weak demand.[14] A modified event study can net out the correlated component of industry-state specific shocks (i.e., the scale effects), allowing us to evaluate how the relationship between early career and older worker outcomes is changing over time. A change in the estimated coefficients over time may capture changing patterns of labor substitution. It cannot, however, distinguish on its own whether these changing patterns are the result of supply-side or demand-side forces.

The primary advantage of this triple-difference design is that it weakens the identifying assumptions needed for the original event study (Olden and Møen, 2022). Industry-specific shocks do not pose a threat to a identification in and of themselves. Instead, the identifying assumption of this design is that, absent the introduction of generative AI, the differential relationship between early career and older worker outcomes would have evolved similarly across industries regardless of AI exposure. This assumption would be violated if correlated shocks affected early career and older workers differently within high-exposure industries than in less exposed ones.

To implement this strategy, I estimate regressions of the form:

$$y_{i,s,g,t} = \alpha_{i,s} + \gamma_t + \sum_t \beta_t y_{i,s,-g,t} + \sum_t \delta_t \text{topquint}_{i,s} y_{i,s,-g,t} + \epsilon_{i,s,t} \quad (3)$$

Here, the outcome for age group g is compared to the same covariate or another indicator for a different

[14] Although employment growth and decline is strongly correlated across demographic groups, there is evidence that young workers are disproportionately impacted by recessions in unemployment risk and employment transitions (e.g., Hoynes, Miller and Schaller, 2012; Haltiwanger, Hyatt and McEntarfer, 2018).

(older) age group. The covariates of interest δ_t estimate the extent to which the relationship between the two age groups is changing differently over time in the most exposed quintile. This estimate can be scaled by the baseline effect β_0 to obtain the degree to which the underlying scale relationship has changed.

Figure 12: Triple difference regressions: Early career worker outcomes regressed on those of prime-age workers

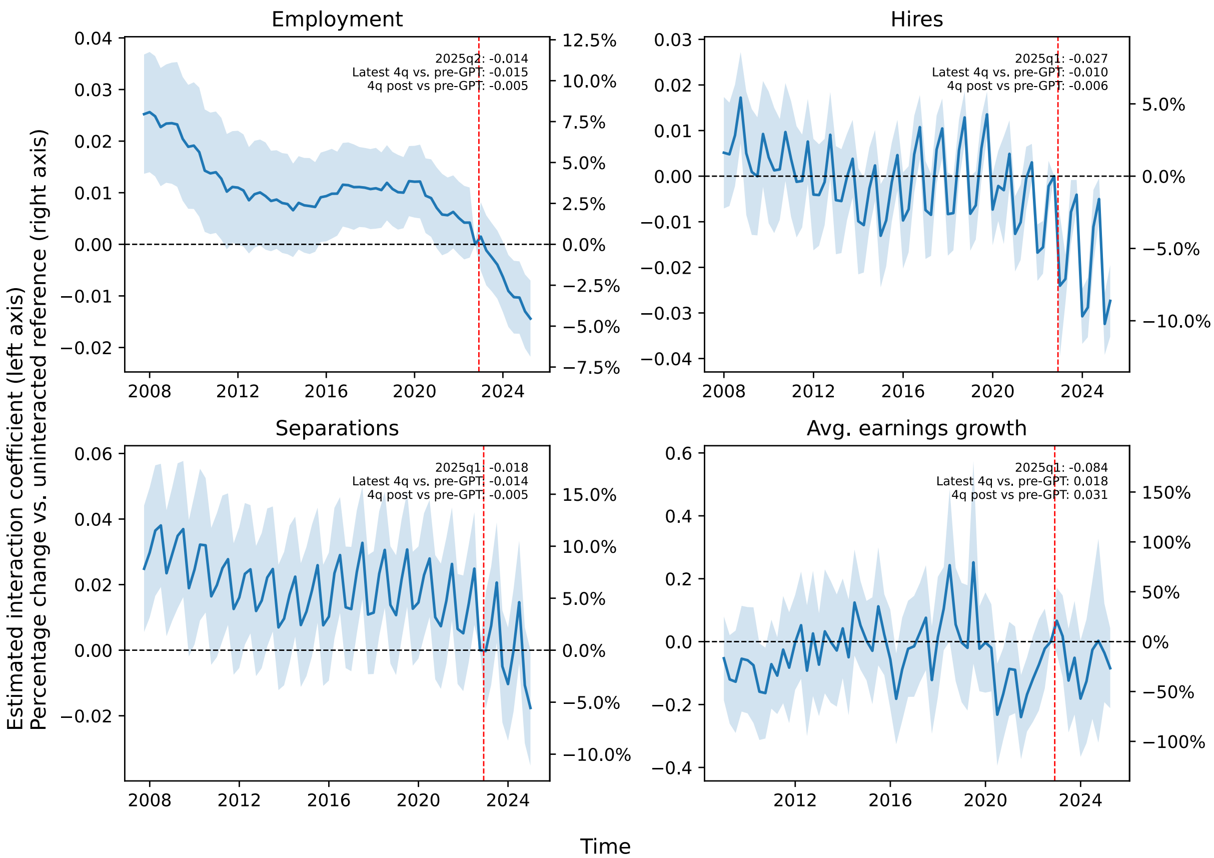

Note: Each plotted series is estimated from a separate regression of the indicated outcome for ages 22-24 on that of ages 25-54, interacted with a period by most-exposed quintile indicator. The reference quarter for beginning-of-quarter employment and separations is 2022q4; it is 2022q3 for end-of-quarter hires and growth in average full-quarters. Text at top right shows the coefficient for the latest available period, the last four quarters compared to the four quarters ending in the reference period, and the first four post-reference quarters compared to pre-reference quarters. Standard errors are two-way clustered by detailed industry and state.

Figure 12 shows the results of regression of early career log employment, hires, separations, and earnings growth on that of prime age (age 25-54) workers. Despite strong residual seasonal patterns in the hires and separation series, the first three panels tell a similar story. The relationship between early career employment and flows and prime age employment and flows has weakened over time, and was significantly weaker as of early 2025 than it was when ChatGPT was released. However, none of these weakening trends began with introduction of ChatGPT. Rather, all three series appear to begin declining at roughly the time of the COVID pandemic.[15] The earnings plot has wide error bounds and is harder to interpret. There is some evidence that relative earnings growth of early career workers may have declined in the most exposed

[15]This pattern also aligns with the earlier observation of Figure 4 that the choice of index period on early career AI-exposed employment has implications for interpretation.

industries, but the relationship between earnings growth of early career workers and prime age workers is now closer to its historical average.

Figure 13: Growth in work at home status for ages 21-29

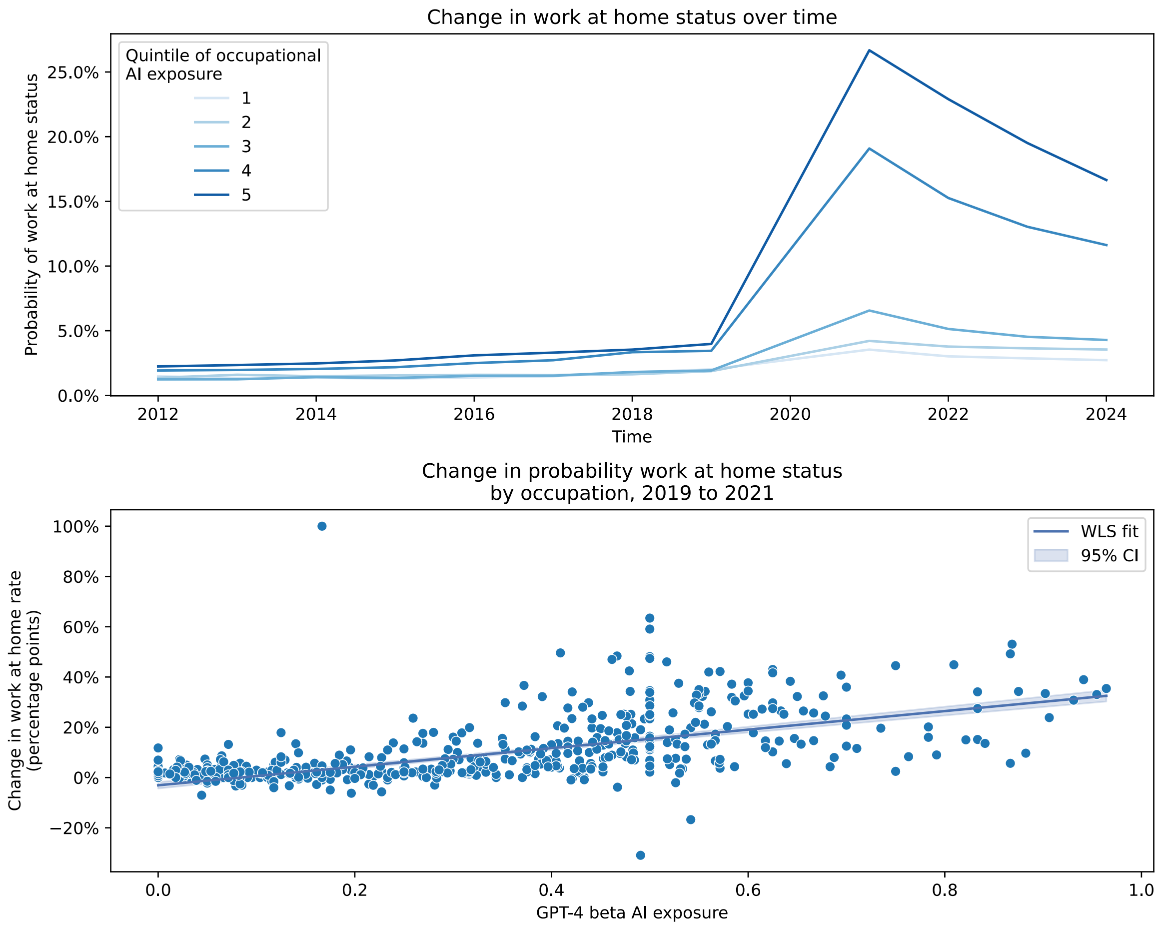

Source: American Community Survey. Each series in the top panel plots the percentage of employed individuals ages 21-29 who report that they worked at home in each year, separately by quintile of occupational AI exposure. Quintile 5 is the most exposed, and quintile 5 the least. The bottom panel plots occupational AI exposure for each ACS occupation against the change in the percentage of individuals ages 21-29 who report that they work at home.

What should one make of these results? Although AI tools are much broader than large language models, there is not a clear case for their impact to have begun as early as 2020. One of the most likely alternative explanations is remote work. Working at home increased at the onset of the pandemic, but it increased more for the most AI-exposed occupations (Figure 13, top panel). ACS data show that a one standard deviation increase in employment-weighted AI exposure among workers ages 21-29 was associated with a 7.8 percentage point increase in the probability of working at home between 2019 and 2021 (bottom panel). Since remote work and AI exposure are quite positively correlated, one plausible explanation for the triple-difference results of Figure 12 is that increased remote work has led to a decline in relative demand for early career workers. Firms might choose to employ fewer young workers if, for example, remote work makes mentorship or monitoring more costly (Emanuel, Harrington and Pallais, 2023). If true, then fully distinguishing the effects of AI from the effects of remote work will likely require better data on remote work linked to individual firms. Ongoing efforts to link BTOS response data to employment records in LEHD may be one avenue for further insights, because it includes supplements that ask about both remote work policies and AI adoption.

Figure 14: Indexed increase in undergraduate degree holding for ages 21-29

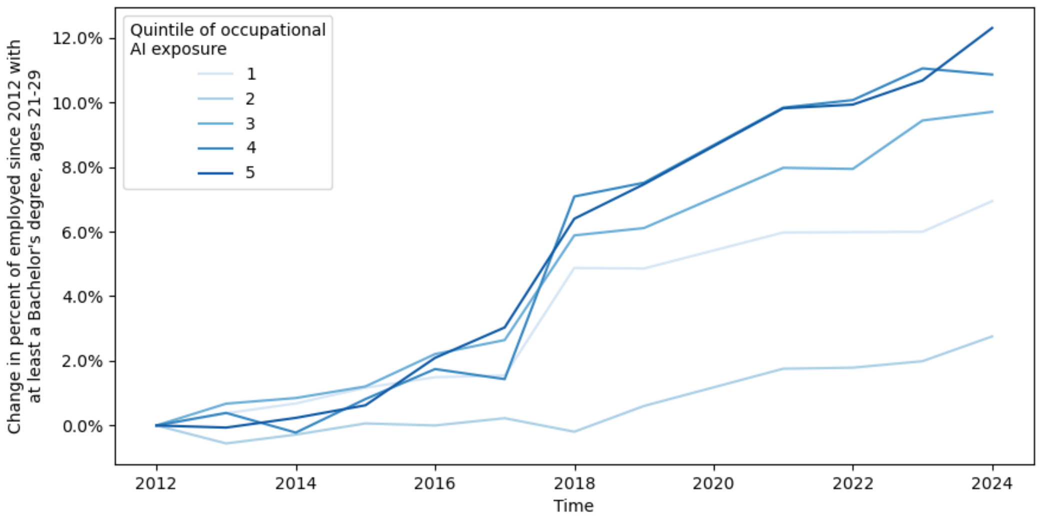

Source: American Community Survey. Each series in the top panel plots the percentage of employed individuals ages 21-29 who report that they have obtained at least a Bachelor's degree separately by quintile of occupational AI exposure. 2012 Bachelor's degree holding for each quintile is indexed to 0.

An alternative explanation is that labor supply of early career workers has declined because individuals are increasing their educational attainment or spending more time in school. The validity of this concern is borne out by Figure 14, which shows using ACS data that workers ages 21-29 in the most AI-exposed quintile of occupations saw more rapid growth in educational attainment than less exposed workers. With narrowly defined age groups, relatively small increases in age at labor market entry could plausibly cause cause large declines in measured employment (Iscenko and Millet, 2026).

The counterfactual analysis shows only a modest contribution of changing compositional effects to trends in employment, but it is important to consider how a delay in labor market entry shows up in the tabulated data. An age 22-24 individual who defers entry into employment – whether because of increased time in school or other supply or demand factors – is immediately observed as a foregone hire. If they are later hired before aging out of the age bin, they are then observed as a hire, but their contribution to total employment of the age bin occurs over fewer quarters. If they are never hired before aging out, then at the time they turn 25 they have an offsetting positive effect on the compositional adjustment. Conversely, a person who would have been hired before turning 22 but defers entry until they are ages 22-24 would have a negative impact on the compositional adjustment first at the time they turn 22, and then might show up as a net addition to hiring within the age group at time of entry. So, only limited information can be obtained about the contribution of delayed labor market entry from the published tabulation. Future work linking college graduates from the Postsecondary Employment Outcomes (PSEO) data by major to their subsequent labor market outcomes may be one way in which the role of education in labor supply is distinguished from the possibility of declining relative demand.

Although residual seasonality complicates inference, it appears that like employment and separations, hires of early career workers began declining in comparison to those of older workers before ChatGPT's release, and a discontinuous drop in hires is not apparent in the triple difference regression specification. Does that imply that the steep decline in hires might be solely the result of industry-specific shocks in industries that hire many young workers, rather than a decline in relative labor demand?

Two additional pieces of evidence suggest that the answer is no. The first comes from the relationship between job gains and losses to early career workers and job gains and losses to prime age workers. Job gains are a subset of hires – specifically, the ones at growing firms, while job losses are separations at shrinking firms. Even more clearly than with overall employment and hiring, job gains and losses of younger and older workers are linked by scale effects. In contrast, replacement hires are not driven by scale relationships, because by definition the scale of the firm is not changing. Figure 15 shows that both firm job gains and firm job losses of early career workers appear to have declined immediately following ChatGPT's release when compared with those of prime age workers in the same top quintile firms. That is, when firms in these industries were growing, they immediately became less likely to hire an early career worker, and relatively more likely to hire a prime age worker. Conversely, when firms in these industries were shrinking, it immediately became more likely that this worker was a prime age worker and less likely that the worker was early career. QWI data do not record whether job losses are voluntary or involuntary, so I cannot distinguish whether the pattern in job losses is because of fewer layoffs specifically. However, these series make clear that the gradual pre-ChatGPT decline of the overall triple difference estimates in hires and separations are a result of changing patterns for replacement hires and for separations unrelated to firm job losses.

The second piece of evidence demonstrating that hiring patterns shifted discontinuously when ChatGPT was released comes from regressions on backfill relationships. That is, many early career hires may be precipitated by the need to backfill a position for someone who separated. True backfill hires are a subset of replacement hires. The QWI data do not indicate which specific separations were backfilled and which were not, so this analysis should be treated with some caution. However, the strong positive association between early career replacement hires and prime age or age 55+ separations suggests that many of these replacement hires are a result of younger workers replacing older ones, rather than early career workers replacing other early career workers. Figure 16 shows that these backfill relationships also appear to have declined discontinuously following ChatGPT's release. That is, firms in the most AI-exposed quintile whose older workers separated became less likely to replace that perspm with a new early career worker.

Future research with restricted use microdata will likely be able further investigate these mechanisms. With firm-level microdata, one can more precisely identify backfill hires, and also do even more to rule out the role of correlated shocks. Nonetheless, the discontinuous changes in job gains, job losses, and backfill suggest that the most AI-exposed firms immediately changed their hiring practices in period following ChatGPT's release.

7 The role of monetary policy

One concern that has been raised in response to event study analyses of generative AI is that ChatGPT's release was immediately preceded by rapid monetary policy tightening in response to a rapid rise in post-pandemic inflation. Moreover, these events also corresponded to a rapid rise in job vacancies and turnover commonly termed the "Great Resignation," ostensibly spurred by the period of historically low interest rates and accommodative fiscal policies that followed the onset of the pandemic. In this telling of events, recent declines in early career employment might be either a direct response to tighter monetary policy, or

Figure 15: Triple difference regressions: Early career job gains, losses, and replacement hires on those of prime-age workers

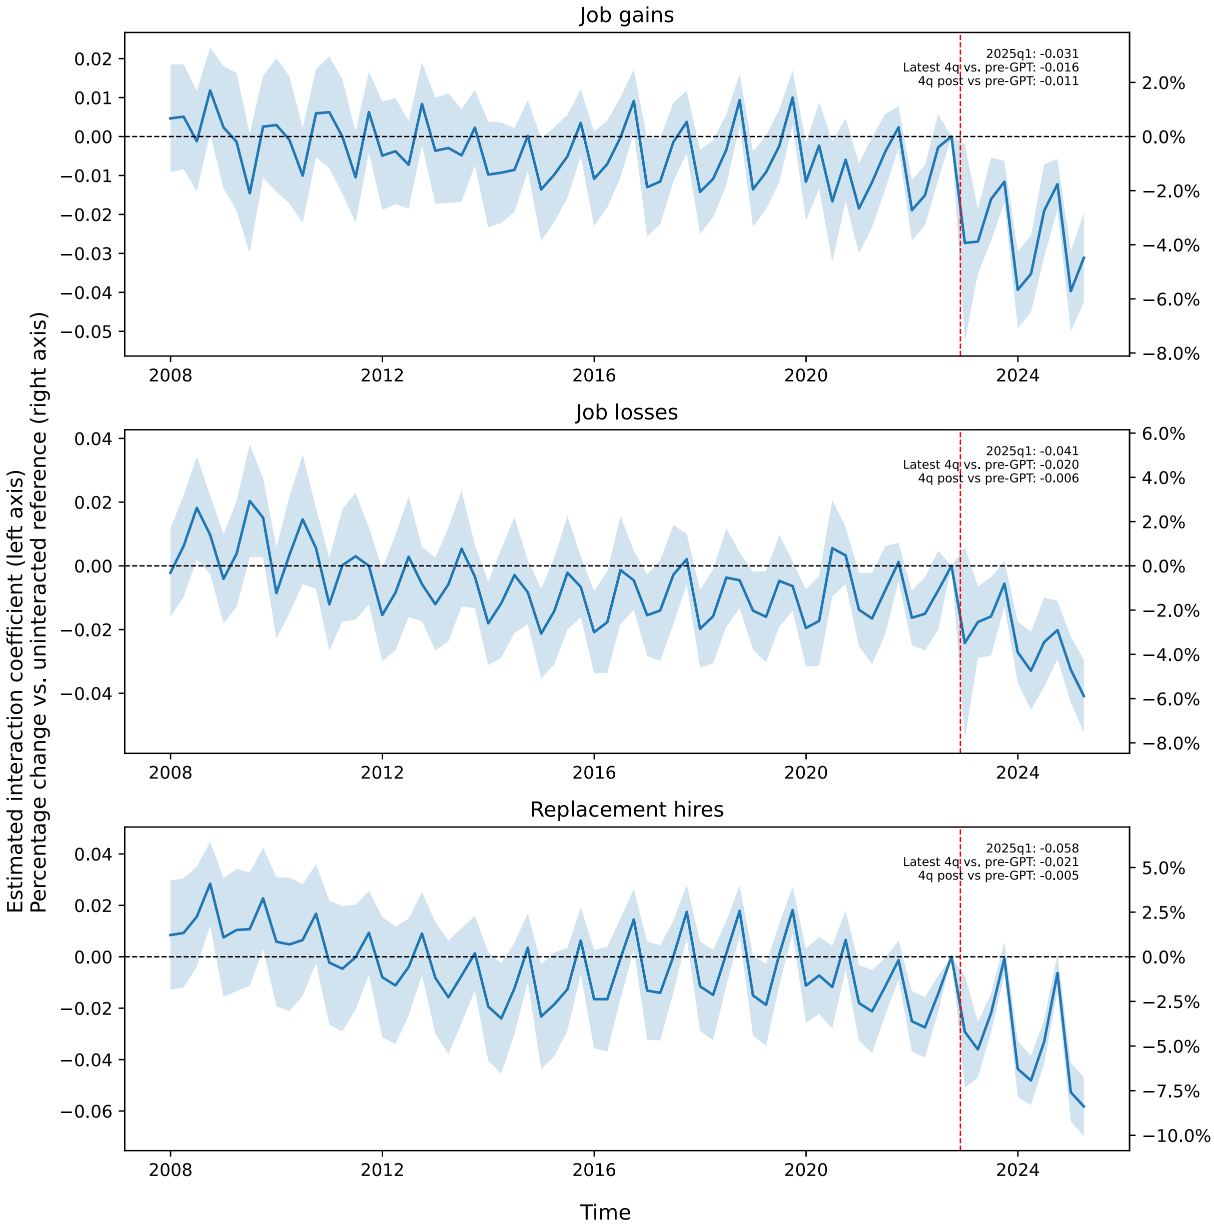

Note: Each plotted series is estimated from a separate regression of the indicated outcome for ages 22-24 on that of ages 25-54, interacted with a period by most-exposed quintile indicator. The reference quarter is 2022q3 for all outcomes. Text at top right shows the coefficient for the latest available period, the last four quarters compared to the four quarters ending in the reference period, and the first four post-reference quarters compared to pre-reference quarters. Standard errors are two-way clustered by detailed industry and state.

Figure 16: Triple difference backfill regressions: Early career replacement hires on older worker separations

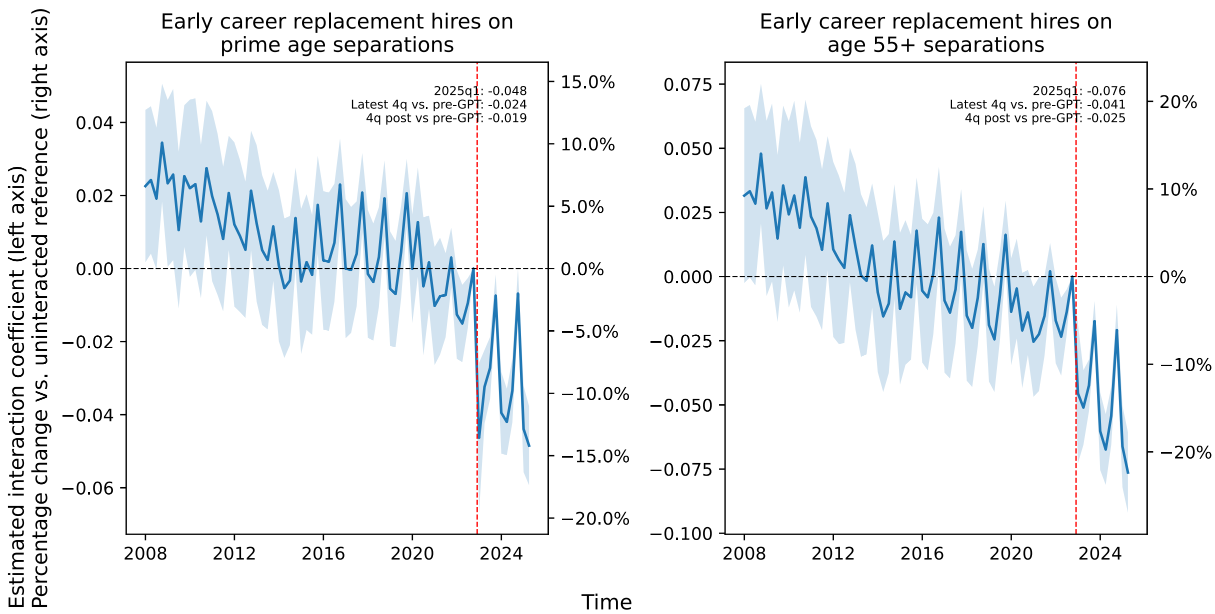

Note: Each plotted series is estimated from a separate regression of hires resulting in end-of-quarter employment for ages 22-24 on separations from beginning-of-quarter employment for older workers in the same quarter, interacted with a period by most-exposed quintile indicator. The reference quarter is 2022q3. Text at top right shows the coefficient for the latest available period, the last four quarters compared to the four quarters ending in the reference period, and the first four post-reference quarters compared to pre-reference quarters. Standard errors are two-way clustered by detailed industry and state.

the unwinding of changes in employment patterns brought about the preceding period of low interest rates. Several researchers and analysts have discussed this competing narrative as a challenge to interpretation of trends in employment (e.g., Hosseini Maasoum and Lichtinger, 2025; Iscenko and Millet, 2026; Frank et al., 2026).

Plotting interest rates against the estimated employment trend (Figure 17), it's easy to see why this is an appealing explanation. The rise in interest rates through 2022 immediately precedes the release of ChatGPT. However, since the impacts of monetary policy are well-known to lag the shocks themselves (e.g., Christiano, Eichenbaum and Evans, 1999), one cannot conclude the role of interest rates as an explanation through a simple comparison of the series. More critically, the plausibility of this explanation hinges the untested assumption that AI-exposed industries are also particularly sensitive to monetary policy shocks. Testing that assumption is the goal of this section.

Prior literature has shown that industries vary widely in their sensitivity to monetary policy shocks, and it has identified a range of reasons why some industries might be expected to more exposed to such shocks than others (Dedola and Lippi, 2005; Peersman and Smets, 2005; Durante, Ferrando and Vermeulen, 2022). In this section, I build on this literature to provide an initial test of the hypothesis that monetary policy might explain recent declines in early-career employment in AI-exposed occupations and industries, and in particular the discontinuous decline in hires that coincided with ChatGPT's introduction.

The local projections method (Jordà, 2005) is a widely used, flexible method for estimating impulse responses to shocks in a panel setting. I implement this methodology by estimating a series of regressions of the form:

Figure 17: Regression-adjusted early career employment in exposed firms vs. interest rates

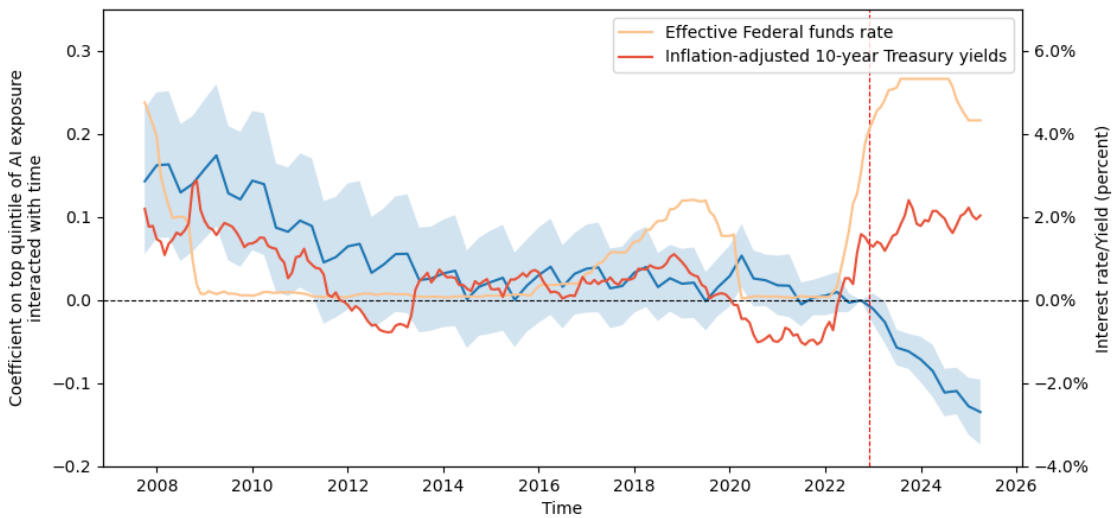

Source: Federal Reserve Bank of St. Louis (FRED). The event study regression result from Figure A.6 is plotted against the selected indicators.

$$y_{i,s,t+h} - y_{i,s,t} = \alpha_{i,s,h} + \beta_{i,h} \cdot \text{MPS}_t + \epsilon_{i,s,t+h} \tag{4}$$

Here, the left-hand side of the equation represents the change in outcome y for 4-digit NAICS industry group i and state s over h periods following an observed shock that takes place in period t. The regression is separated estimated for leads $h \in [1, ...8]$. Industry-state fixed effects $\alpha_{i,s,h}$ control for persistent trends in y over the sample period that may differ by state and industry. The covariates of interest $\beta_{i,h}$ are estimated on a measure of the orthogonal component of monetary policy shocks in time t from Bauer and Swanson (2023). An estimate is obtained for each 4-digit NAICS industry group in each regression, while estimates for different leads h are obtained from separate regressions. A more negative estimated coefficient $\beta_{i,h}$ represents a stronger negative response to an unexpected monetary policy shock, while positive coefficients represent a countercyclical response. Outcomes measured include logs of employment, hires, separations, and hiring and separation rates (i.e., scaled by average employment over the quarter). All estimations using all available time periods of data from QWI through 2016q4, so that all impulse response periods exclude the COVID pandemic.

This approach has distinct advantages and disadvantages relative to alternative methods such as structural VAR. The local projection approach permits estimation across all industries in a pooled panel, allowing for heterogeneous responses at the NAICS industry group level and only aggregating for a summary of net effects. Relative to structural VAR, local projections are more robust to model misspecification, since impulse responses are estimated directly at each horizon rather than iterated forward from a single estimated system. A practical tradeoff is that local projections tend to yield less precise estimates, particularly at longer horizons, where VARs impose additional smoothness through their dynamic structure. Use of the Bauer and Swanson (2023) measure helps to improve the estimates of monetary policy's effect, by controlling for confounding information effects to better isolate the causal role of policy shifts themselves.

The main limitation of this approach is an exact counterpoint to that strength in causal identification. High-frequency measurement of surprise monetary policy shocks, by design, omit non-surprise sources of variation in monetary policy. If changes in monetary policy are largely anticipated, or if the effects of interest rates operate through the anticipated component of policy, then this method may underestimate the magnitude of subsequent impacts.

The goal of this exercise is not to maximally predict changes in employment across industries. Rather, it is to see whether AI-exposed industries are more sensitive to such shocks than others, and if so, to such an extent that it could plausibly explain the declines in employment and hiring following ChatGPT's introduction. In this context, the high-frequency measure of monetary policy shocks is appropriate because it ensures that estimates are unlikely to be influenced by contemporaneous confounding factors.

Figure 18 shows the estimated impulse response functions for five outcomes of interest, separately for all workers (left side) and early career workers (right side). Results are scaled to show the impulse response to a single one standard deviation positive surprise monetary policy shock. As expected, both employment and hires are negative following the shock. Hires bottom out between four and six quarters later with a decline of about 1.5%. Separations are also negative, bottoming out about 8 quarters after the shock with a similar magnitude. QWI data do not distinguish between voluntary and involuntary separations, but this pattern of declining separations aligns with previous findings that reduced quits outweigh increased involuntary separations following monetary policy shocks (Graves, Huckfeldt and Swanson, 2023). Employment following monetary policy shock declines gradually and levels out eight to ten quarters following the shock, but it does not appear to recover. Response patterns for early career workers are broadly similar to those of all workers.

That both AI-exposed and less-exposed industries respond negatively to monetary policy shocks is clear. To assess whether monetary policy shocks are a plausible alternative explanation for the decline in top-quintile early career hires and employment, what matters is the difference between the most exposed industries and other firms. Here, the results shown in Figure 18 are less conclusive. The early career hires series differ relatively little between the most exposed quintile of industries and others; this is also true of separations. However, the employment series do show a gap of about 0.5% in employment between employment in the most exposed quintile and other following a one standard deviation monetary policy shock.[16]

Of course, as Figure 17 demonstrates, the period preceding ChatGPT's introduction saw several large interest rate increases. Having estimated monetary policy sensitivity for each industry on pre-COVID data, I now use these estimates to trace the projected path of each outcome over the period that includes COVID, the introduction of ChatGPT, and subsequent quarters, using data on monetary shocks over that time period. That is, I compute:

$$\hat{y}_{i,t} = \sum_{h=1}^{12} \hat{\beta}_{i,h} \cdot \mathrm{MPS}_{t-h} \tag{5}$$

where each $\hat{\beta}_{i,h}$ is the estimate from the local projection regression on the pre-COVID period and MPS_{t-h} is the monetary policy shock from subsequent periods. I normalize 2022q4 to 0, so that the magnitude of the employment change (measured in logs) can be compared to the indexed employment change shown in Figure 3. In contrast, the event study regression figures are appropriately compared to the gap between the most exposed quintile and other industries series in each figure. Note that since the Bauer and Swanson series ends in December 2023, I assume 0 shocks from that point forward; all changes after this point are the impulse responses to preceding shocks.

[16] These series are estimated separately, so hires, separations, and employment changes are not constrained to align.

Figure 18: Estimated impulse responses to a one standard deviation monetary policy shock

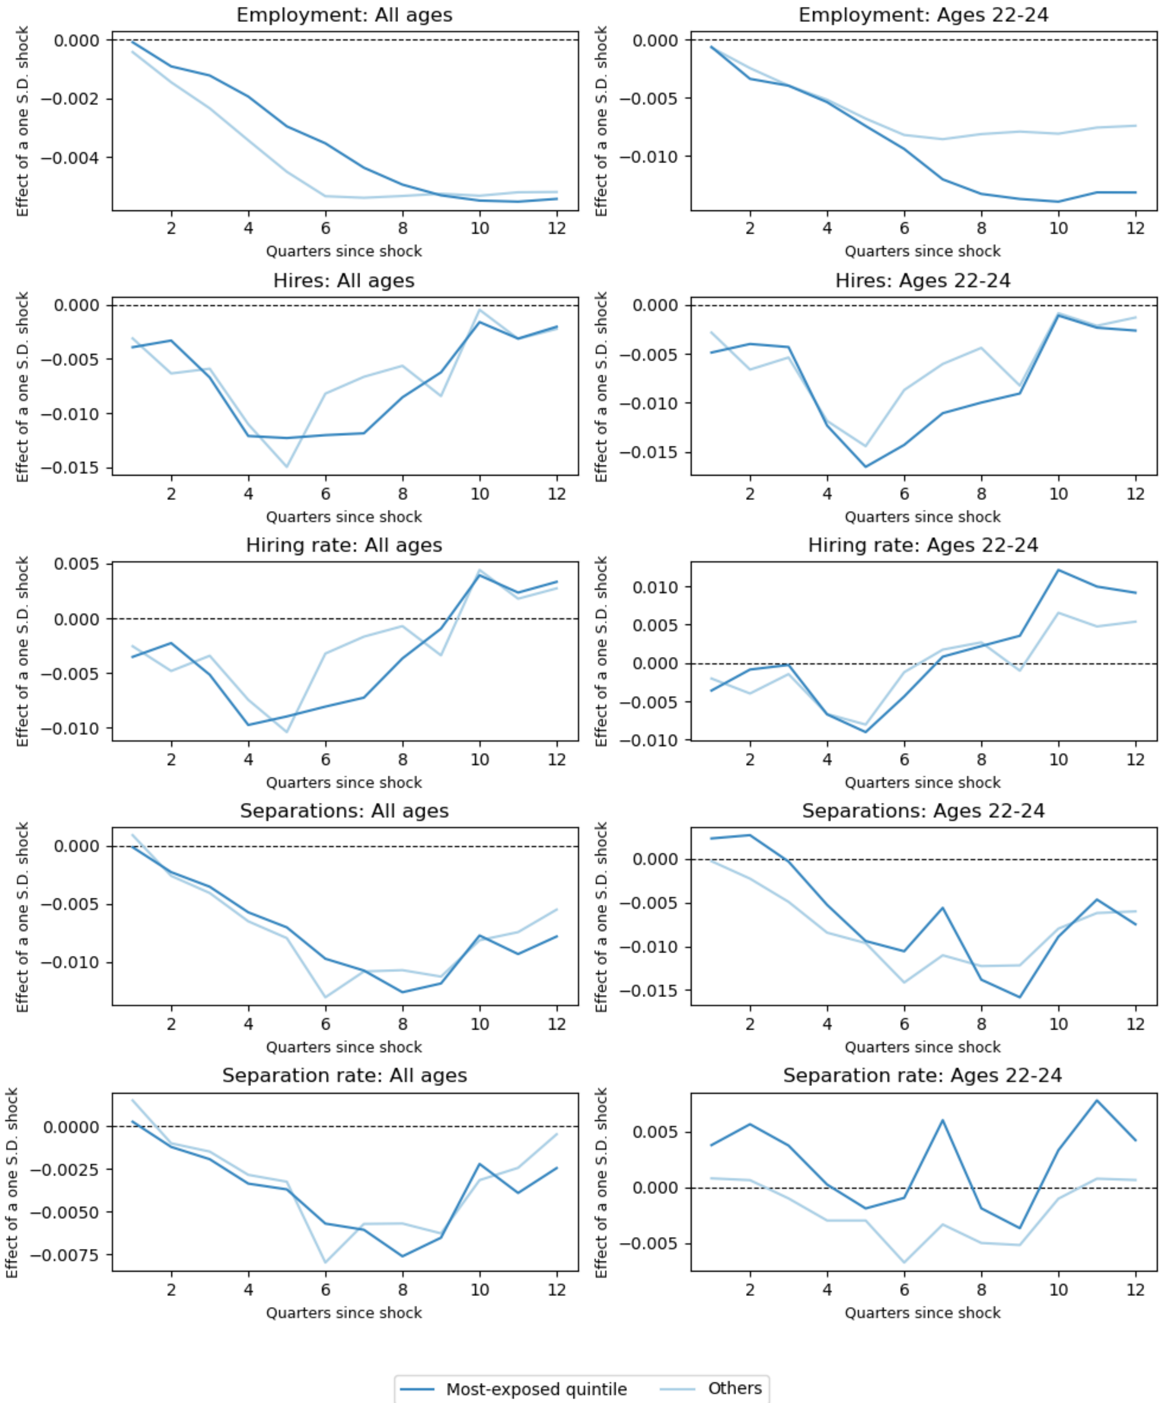

Note: Each plot shows the estimated impulse response for the selected outcome and age group to a one standard-deviation positive monetary policy shock. Local projections are estimated separately for each NAICS industry group using all available QWI data through 2016q4, and the employment-weighted average response for the most AI-exposed quintile of industries states and all other industry-states is plotted. All outcomes are in logs. Monetary policy shock data from Bauer and Swanson (2023).

Figure 19: Cumulative response to monetary policy shocks

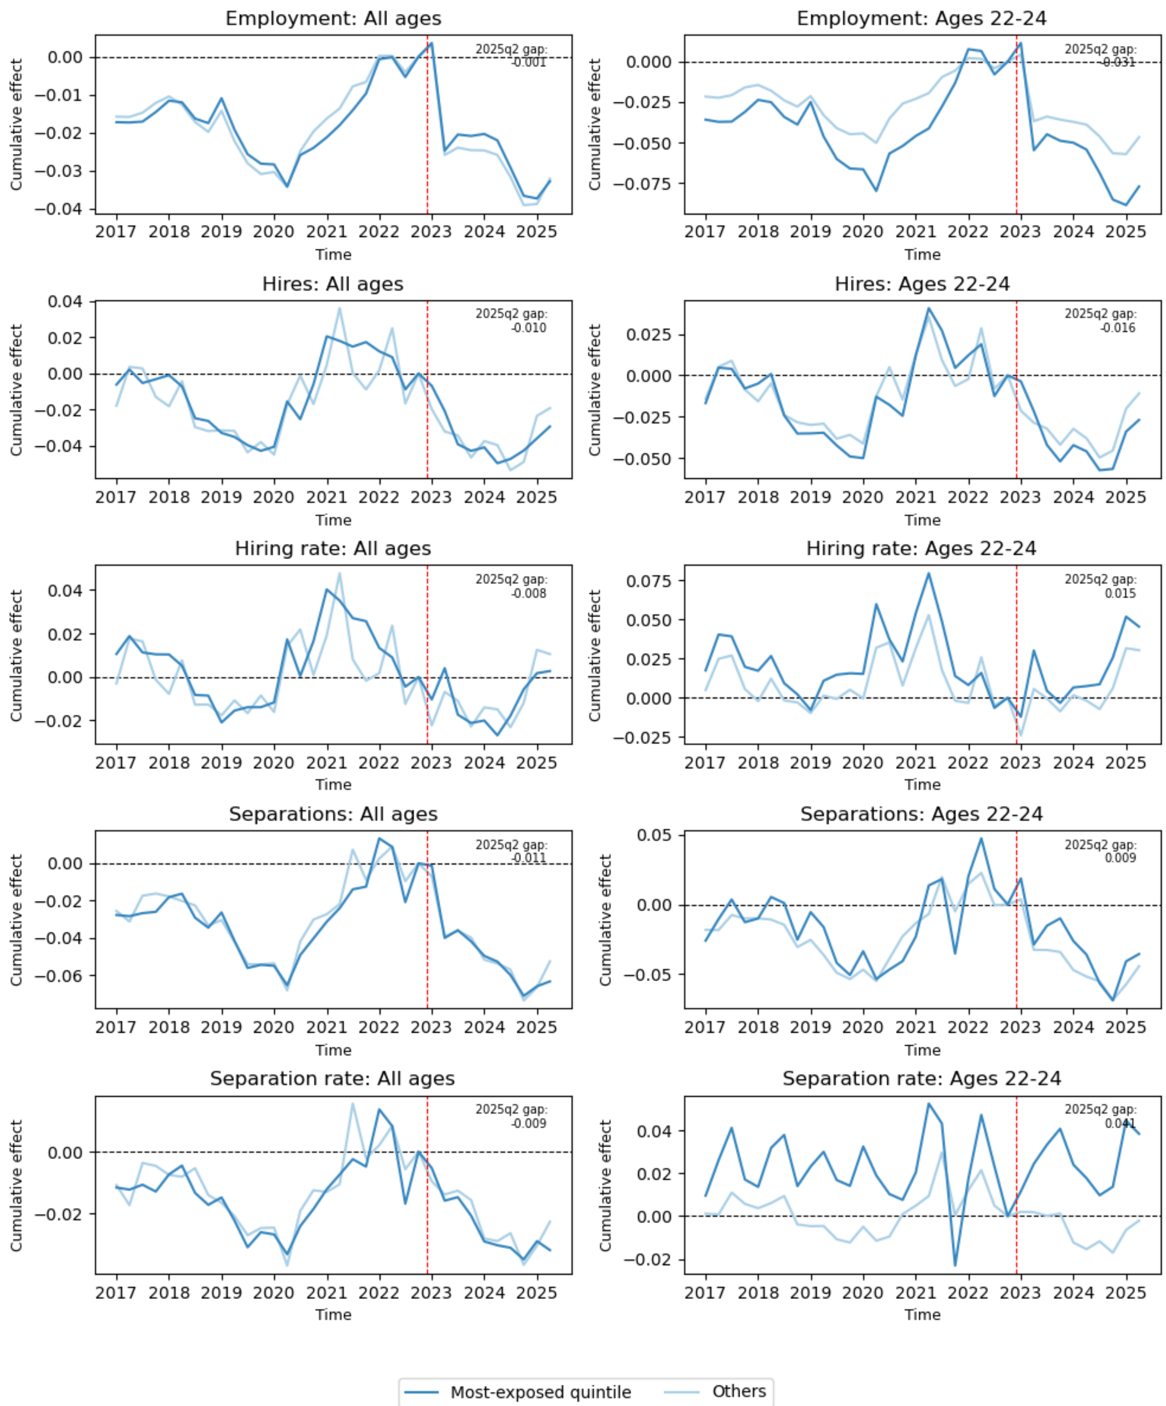

Note: Each plot shows the cumulative estimated impact of 12 lags of monetary policy shocks from Bauer and Swanson (2023) on the selected outcome, age group, and aggregation of NAICS industry groups. All outcomes are in logs, normalized so that 2022q4 is 0. The shocks series ends in 2023q4, so 0 shocks are incorporated from 2024q1 to 2025q2.

Figure 19 shows the result, for the same outcomes and age groups as the preceding figure. Of these panels, the most supportive of the role of monetary policy is the one showing predicted changes in early career employment (top right), as anticipated by the difference in sensitivity coefficients. Predicted employment of both most AI-exposed quintile industries and others declines sharply in early 2023, with the most AI-exposed industries a bit more affected. By 2025q2, the predicted gap between most exposed quintile and other industry early career employment reaches 3.1%, one quarter of the 12.4% regression-adjusted point estimate. The predicted gap in employment is confined to early career workers; the two all ages employment series differ little over the entire time frame, though they show a similar shape of employment decline in early 2023.

In several other ways, the employment series do not match reality well. Most notably, they predict a large employment decline across all industries and age groups in early 2023, consistent with predictions of a recession driven by monetary tightening that were made at that time. but no substantial employment decline occurred among early career workers in less AI exposed industries, nor among older workers.

The series for hires and separations also show evidence that the local projection model predicts economic weakness in 2023. Early career hires decline by just under 5% by mid-2023, about half of the observed decline in the most AI-exposed industries. But, the model also predicts a contemporaneous decline in other industries that did not occur. The predicted gap in hires as of 2025q2 of 1.6% is only one eighth the size of the event study estimate. And, the predicted gap in separations is positive, despite an event study estimate of 11.5% relative decline.

Taken collectively, these figures provide fairly thin evidence for the role of monetary policy in explaining differences between early career workers in the most exposed industries and in others. Most of the negative impacts of monetary policy on employment, hires, and separations are shared across AI-exposed industries and others, because historical analysis suggests that these industries' sensitivity to such shocks is about average. And, there is little evidence to suggest that early career workers might bear an unusually large impact, especially through a rapid decline in hires.

8 Conclusion

In the ten quarters after generative AI became widely available, employers in the most AI-exposed industries shed over 150,000 early career jobs, even as employment of older workers in the same industries was comparatively unaffected. The results of this analysis encourage some caution in attributing all of this to AI. Triple-difference estimators suggest that relative employment of early careers within the most AI-exposed firms was declining for years beforehand, perhaps as a side effect of remote work or delayed career entry.

Yet, the immediate and persistent decline in early career hiring in AI-exposed industries upon ChatGPT's release is hard to reconcile with the direct effect of AI on production activity, and also hard to reconcile with monetary policy shocks or labor supply explanations. How, for example, can an immediate 9% decline in hires be consistent with the BTOS data shown in Figure 2, indicating that only a few subsectors reported 9% adoption of AI within the first year of ChatGPT's release?

One plausible mechanism is that public availability of generative AI caused firms to delay hiring because it increased their uncertainty about the technology's future capabilities, and perhaps also about future economic conditions. Firms may see their early career workforce as an investment, and uncertainty shocks can cause firms to defer investments (Bernanke, 1983; Dixit and Pindyck, 1994). Specifically, hiring may decline rapidly if firms seek to avoid costly misallocation of resources and preserve the option value to reallocate once the uncertainty is resolved (Bloom, 2009).This logic may be particularly relevant to the case of AI, not merely because of the potential for AI to automate or augment certain types of human labor, but because changes to firms' production methods might necessitate an overhaul of their broader hiring practices and organizational structures (Brynjolfsson, Rock and Syverson, 2017). There may be not only increased risk of misallocating resources, but also greater uncertainty about when that risk will resolve. And, higher interest rates may also have played a contributing role, by raising the hurdle for uncertain investments in speculative early career hires.

If this is the mechanism behind the rapid decline in hiring, then given the continued widespread disagreement about what the future economic effects will be (Karger et al., 2026), it's conceivable that firms will remain cautious about early career hiring for the foreseeable future. The hiring rate has roughly returned to its previous level as firms' early career workforce has shrunk, but there is no evidence of catch-up, nor any evidence in the QWI data of early career AI-exposed employment leveling off.

Another possibility is that firms made the choice to reduce their hiring based on technologies that they had not yet adopted, but expected to adopt during the job spells of the workers they would otherwise have hired. Further work using non-public tabulations of LEHD data may shed further light on this distinction, by observing changes in the job spells of these workers. Limited linked employer-employee evidence from Europe (Humlum and Vestergaard, 2025) suggests that large changes in tasks following adoption have had little effect on employment and earnings there, but evidence from U.S. job posting data supports the notion that adopting firms may be using AI adoption to reduce their early-career hiring (Hosseini Maasoum and Lichtinger, 2025). As survey sources on AI adoption become linkable to matched employer-employee data, more direct evidence of the role of AI adoption may emerge.

Either way, the implications of a persistent decline in hires could be substantial. Young workers often move up a "job ladder" from less productive to more productive firms, and in previous periods of weak demand, the job ladder has collapsed, with negative consequences for these workers (Haltiwanger, Hyatt and McEntarfer, 2018). Further research will needed to be done to assess whether the decline in hires has similar effects on workers who might otherwise have been employed in AI-exposed industries.

Undoubtedly, much more work will need to be done before a cohesive picture of AI's early impacts is in focus. The finding of a clear discontinuity in hires among the most AI-exposed firms is only one piece of that picture. Interpreting it as clear evidence of how AI can or will affect early career employment in the long run is premature. However, this evidence bolsters the case that AI is playing at least some role in labor market weakness among early career workers in these fields, whether directly through its use, or indirectly through the uncertainty it has created. An appropriately measured interpretation of this finding may help us to better understand the underlying forces at play, and more quickly identify longer run impacts as they occur.

References

Abowd, John M, Bryce E Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L McKinney, Marc Roemer, and Simon Woodcock. 2009. "The LEHD infrastructure files and the creation of the Quarterly Workforce Indicators." In *Producer dynamics: New evidence from micro data.* 149–230. University of Chicago Press.

Bauer, Michael D, and Eric T Swanson. 2023. "A reassessment of monetary policy surprises and high-frequency identification." *NBER Macroeconomics Annual*, 37(1): 87–155.

Bernanke, Ben S. 1983. "Irreversibility, Uncertainty, and Cyclical Investment*." *The Quarterly Journal of Economics*, 98(1): 85–106.

Bick, Alexander, Adam Blandin, and David J Deming. 2026. "The rapid adoption of generative AI." *Management Science*.

Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica*, 77(3): 623–685.

Bonney, Kathryn, Cory Breaux, Cathy Buffington, Emin Dinlersoz, Lucia S Foster, Nathan Goldschlag, John C Haltiwanger, Zachary Kroff, and Keith Savage. 2024. "Tracking firm use of AI in real time: A snapshot from the Business Trends and Outlook Survey." National Bureau of Economic Research.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine?: Six Facts about the Recent Employment Effects of Artificial Intelligence." Stanford Institute for Economic Policy Research (SIEPR).

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2017. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." National Bureau of Economic Research Working Paper 24001.

Cajner, Tomaz, Leland Dod Crane, Ryan Decker, Adrian Hamins-Puertolas, Christopher Johann Kurz, and Tyler Radler. 2018. "Using payroll processor microdata to measure aggregate labor market activity." *FEDS Working paper*.

Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans. 1999. "Monetary policy shocks: What have we learned and to what end?" *Handbook of macroeconomics*, 1: 65–148.

Council of Economic Advisers. 2024. "Potential labor market impacts of artificial intelligence: An empirical analysis." Council of Economic Advisers.

Dedola, Luca, and Francesco Lippi. 2005. "The monetary transmission mechanism: evidence from the industries of five OECD countries." *European Economic Review*, 49(6): 1543–1569.

Dixit, Avinash K, and Robert S Pindyck. 1994. *Investment under uncertainty.* Princeton university press.

Durante, Elena, Annalisa Ferrando, and Philip Vermeulen. 2022. "Monetary policy, investment and firm heterogeneity." *European Economic Review*, 148: 104251.

Eisfeldt, Andrea L, Gregor Schubert, and Miao Ben Zhang. 2023. "Generative AI and Firm Values." National Bureau of Economic Research Working Paper 31222.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs are GPTs: Labor market impact potential of LLMs." *Science*, 384(6702): 1306–1308.

Emanuel, Natalia, Emma Harrington, and Emma Pallais. 2023. "The Power of Proximity to Coworkers: Training for Tomorrow or Productivity Today?" ¡p¿Revise and resubmit, Quarterly Journal of Economics¡/p¿.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses." *Strategic management journal*, 42(12): 2195–2217.

Frank, Morgan R, Alireza Javadian Sabet, Lisa Simon, Sarah H Bana, and Renzhe Yu. 2026. "AI-exposed jobs deteriorated before ChatGPT." *arXiv preprint arXiv:2601.02554*.

Goldfarb, Avi, and Catherine Tucker. 2019. "Digital economics." *Journal of economic literature*, 57(1): 3–43.

Graves, Sebastian, Christopher Huckfeldt, and Eric T Swanson. 2023. "The labor demand and labor supply channels of monetary policy." *NBER Working Paper*, , (w31770).

Haltiwanger, John, Henry Hyatt, and Erika McEntarfer. 2018. "Who Moves Up the Job Ladder?" *Journal of Labor Economics*, 36(S1): S301–S336.

Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, Kevin K. Troy, Dario Amodei, Jared Kaplan, Jack Clark, and Deep Ganguli. 2025. "Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations."

Hosseini Maasoum, Seyed Mahdi, and Guy Lichtinger. 2025. "Generative AI as seniority-biased technological change: Evidence from US résumé and job posting data." *Available at SSRN 5425555*.

Hoynes, Hilary, Douglas L Miller, and Jessamyn Schaller. 2012. "Who suffers during recessions?" *Journal of Economic perspectives*, 26(3): 27–48.

Humlum, Anders, and Emilie Vestergaard. 2025. "Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI." National Bureau of Economic Research Working Paper 33777.

Iscenko, Zanna, and Fabien Curto Millet. 2026. "Looking for the ladder: Is AI impacting entry-level jobs." Economic Innovation Group.

Johnston, Andrew C, and Christos Makridis. 2025. "The labor market effects of generative AI: A difference-in-differences analysis of AI exposure." *Available at SSRN 5375017*.

Jordà, Òscar. 2005. "Estimation and inference of impulse responses by local projections." *American economic review*, 95(1): 161–182.

Karger, Ezra, Otto Kuusela, Jason Abaluck, Kevin A Bryan, Basil Halperin, Todd R Jones, Connacher Murphy, Philip Trammell, Matt Reynolds, Dan Mayland, Ria Viswanathan, Ananaya Mittal, Rebecca Ceppas de Castro, Josh Rosenberg, and Philip Tetlock. 2026. "Forecasting the Economic Effects of AI." National Bureau of Economic Research Working Paper 35046.

Massenkoff, Maxim, and Peter McCrory. 2026. "Labor market impacts of AI: A new measure and early evidence." *Anthropic Research*, 5.

Olden, Andreas, and Jarle Møen. 2022. "The triple difference estimator." *The Econometrics Journal*, 25(3): 531–553.

Peersman, Gert, and Frank Smets. 2005. "The industry effects of monetary policy in the euro area." *The Economic Journal*, 115(503): 319–342.

Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. 2025. "Working with AI: Measuring the Applicability of Generative AI to Occupations."

A Appendix Materials

Figure A.1: Event study regressions: Early career employment by firm size

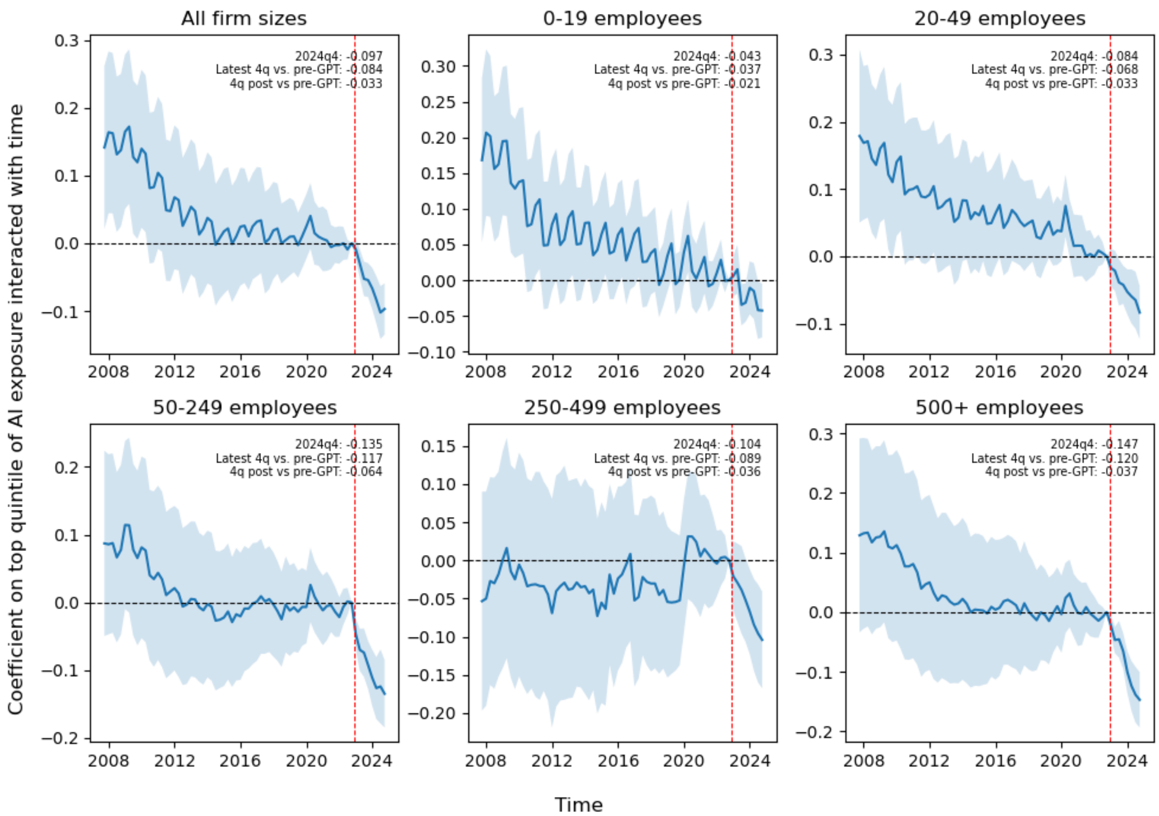

Note: Each series is estimated from an event study regression of log beginning-of-quarter employment of ages 22-24 on a period by most-exposed quintile indicator, separately by binned firm size. The reference period is 2022q4. AI exposure, quintile definitions, and fixed effects specifications are at the NAICS industry group (4-digit) level. Text at top right shows the coefficient for the latest available period, and the last four quarters compared to the four quarters ending in the reference period. Standard errors are two-way clustered by detailed industry and state.

Figure A.2: Event study regressions: Early career hires by firm size

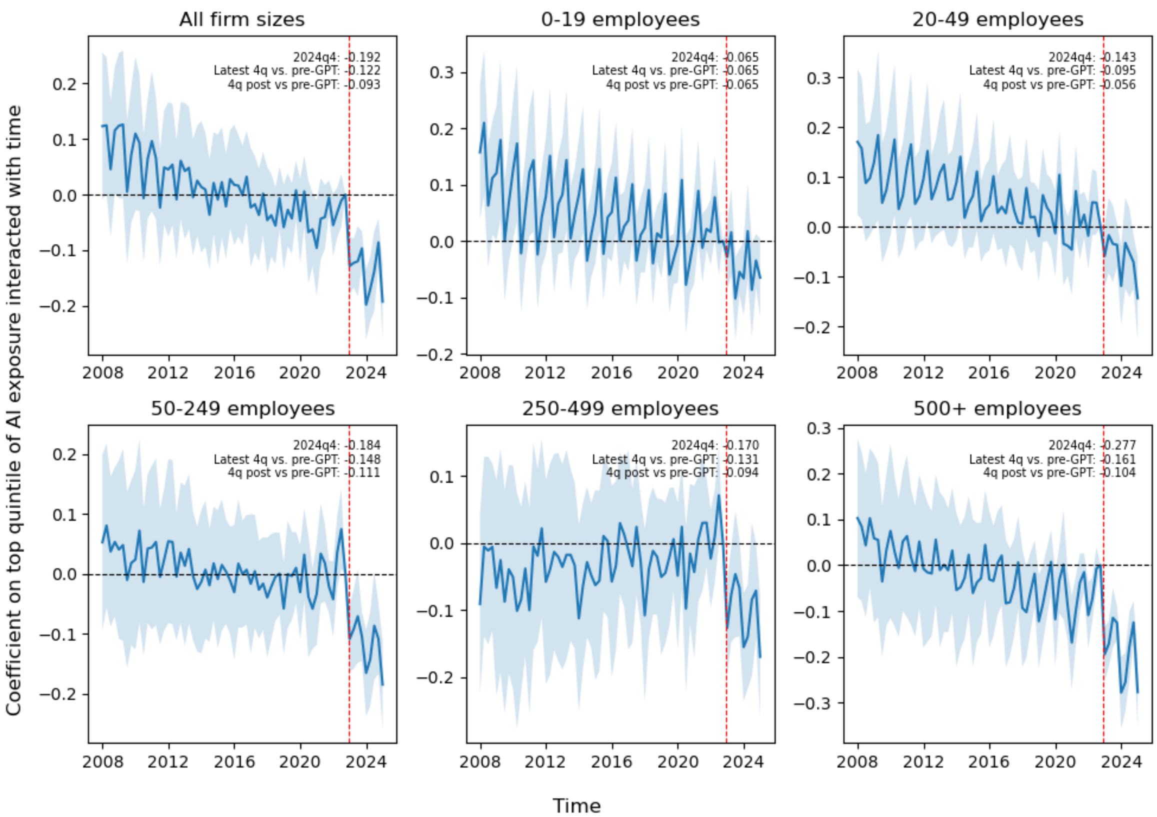

Note: Each series is estimated from an event study regression of log hires resulting in end-of-quarter employment of ages 22-24 on a period by most-exposed quintile indicator, separately by binned firm size. The reference period is 2022q3. AI exposure, quintile definitions, and fixed effects specifications are at the NAICS industry group (4-digit) level. Text at top right shows the coefficient for the latest available period, and the last four quarters compared to the four quarters ending in the reference period. Standard errors are two-way clustered by detailed industry and state.

Figure A.3: Event study regressions: Early career separations by firm size

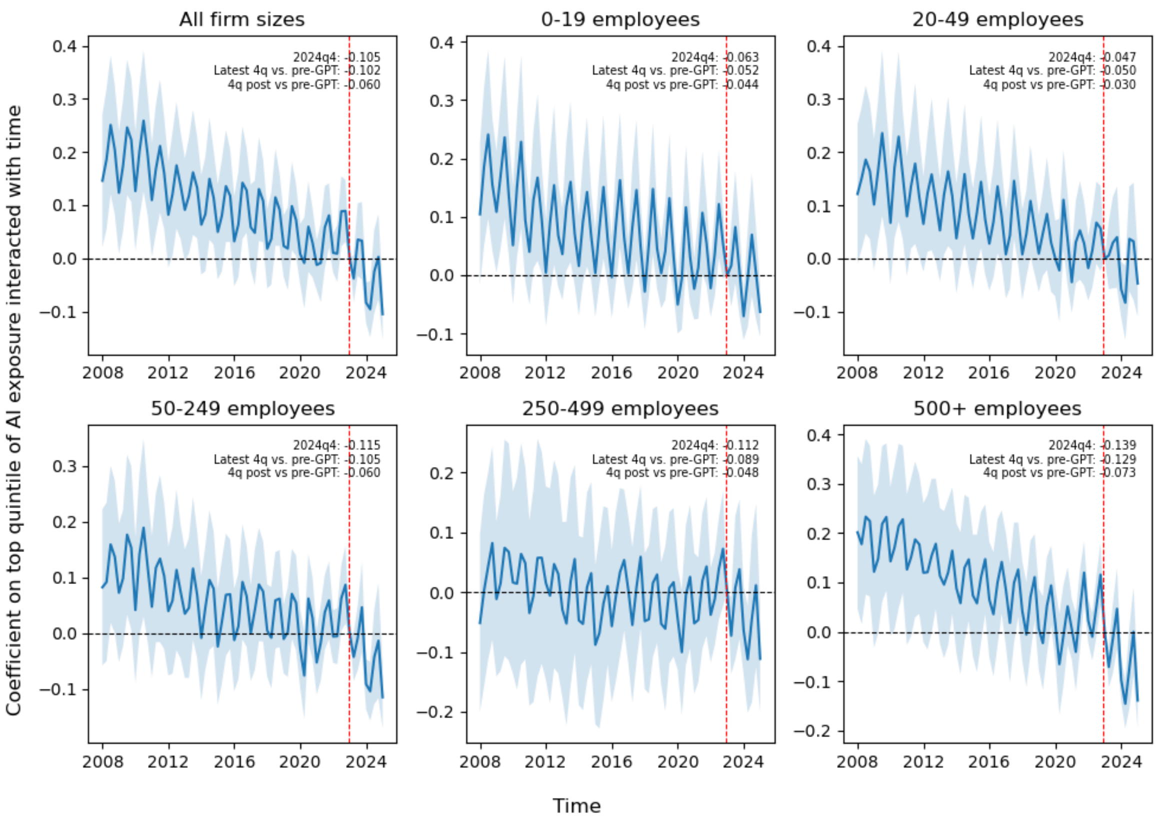

Note: Each series is estimated from an event study regression of log separations from beginning-of-quarter employment of ages 22-24 on a period by most-exposed quintile indicator, separately by binned firm size. The reference period is 2022q4. AI exposure, quintile definitions, and fixed effects specifications are at the NAICS industry group (4-digit) level. Text at top right shows the coefficient for the latest available period, and the last four quarters compared to the four quarters ending in the reference period. Standard errors are two-way clustered by detailed industry and state.

Figure A.4: Event study regressions: Early career employment by firm age

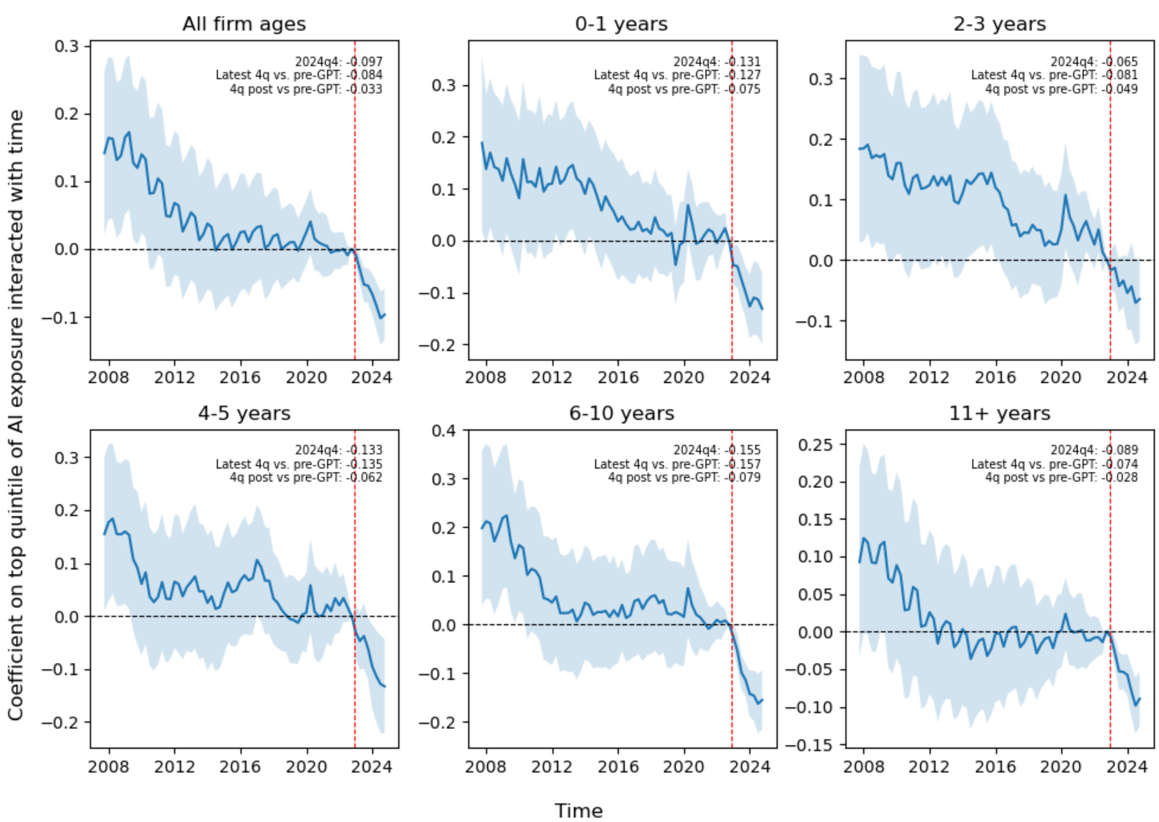

Note: Each series is estimated from an event study regression of log beginning-of-quarter employment of ages 22-24 on a period by most-exposed quintile indicator, separately by binned firm size. The reference period is 2022q4. AI exposure, quintile definitions, and fixed effects specifications are at the NAICS industry group (4-digit) level. Text at top right shows the coefficient for the latest available period, and the last four quarters compared to the four quarters ending in the reference period. Standard errors are two-way clustered by detailed industry and state.

Figure A.5: Event study regressions: Early career hires by firm age

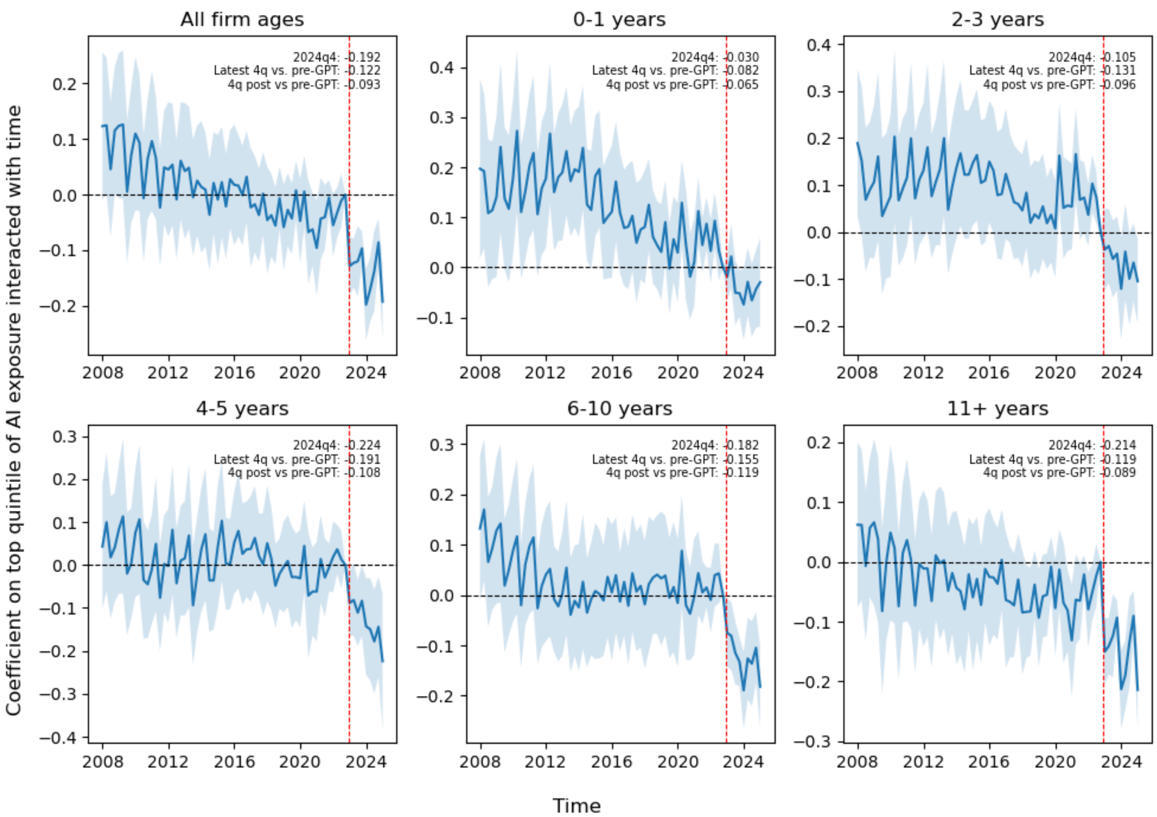

Note: Each series is estimated from an event study regression of log hires resulting in end-of-quarter employment of ages 22-24 on a period by most-exposed quintile indicator, separately by binned firm age. The reference period is 2022q3. AI exposure, quintile definitions, and fixed effects specifications are at the NAICS industry group (4-digit) level. Text at top right shows the coefficient for the latest available period, and the last four quarters compared to the four quarters ending in the reference period. Standard errors are two-way clustered by detailed industry and state.

Figure A.6: Event study regressions: Early career separations by firm age

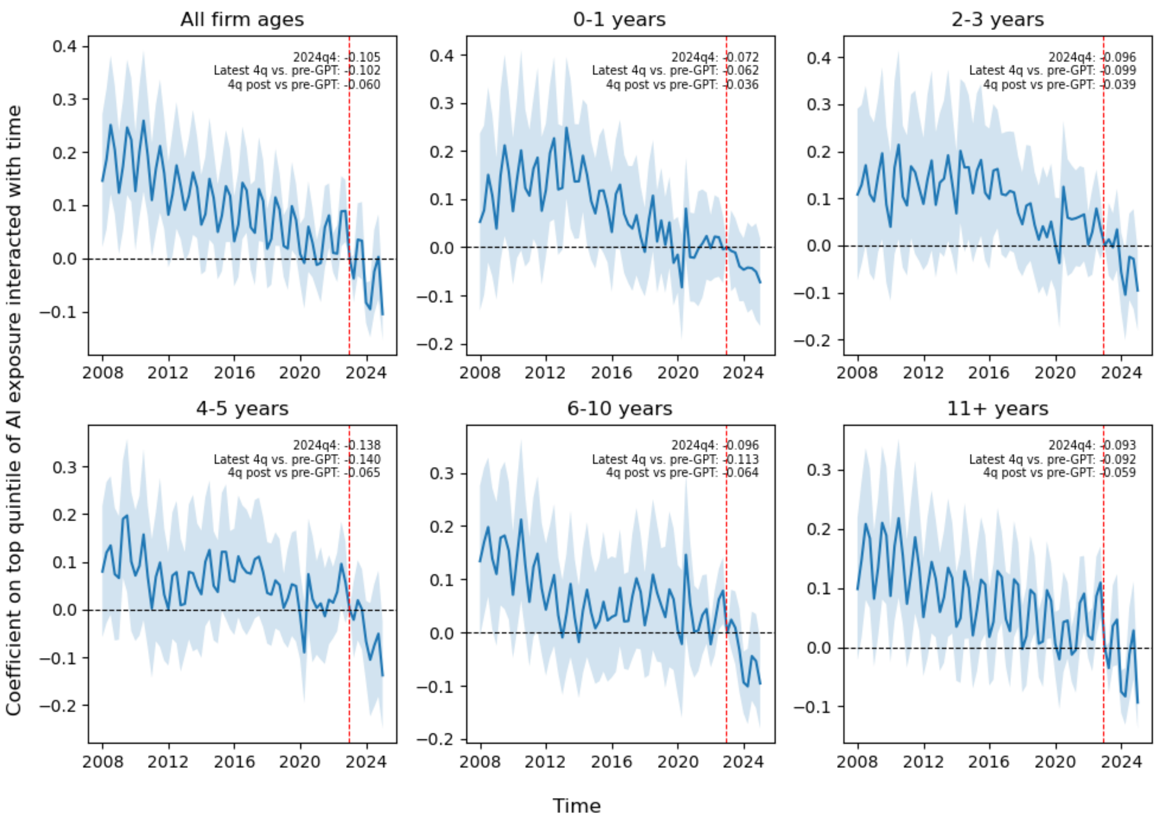

Note: Each series is estimated from an event study regression of log separations from beginning-of-quarter employment of ages 22-24 on a period by most-exposed quintile indicator, separately by binned firm age. The reference period is 2022q4. AI exposure, quintile definitions, and fixed effects specifications are at the NAICS industry group (4-digit) level. Text at top right shows the coefficient for the latest available period, and the last four quarters compared to the four quarters ending in the reference period. Standard errors are two-way clustered by detailed industry and state.

Table A.1: Comparison of measures in predicting industry AI Adoption

Measure		Subsector regression		Sector regression	
Source	Occ. weighting	R-squared	Std. Coefficient	R-squared	Std. Coefficient
Eisfelt, Schubert and Zhang, core task exposure	Agg. hours of full-time, full-year workers	0.582	7.866	0.744	8.306
Eisfelt, Schubert and Zhang, core task exposure	Total employment	0.554	7.577	0.705	8.087
Eisfelt, Schubert and Zhang, total task exposure	Agg. hours of full-time, full-year workers	0.512	7.333	0.689	7.947
Eisfelt, Schubert and Zhang, total task exposure	Total employment	0.495	7.105	0.656	7.752
Eloundou et al., GPT4	Total employment	0.465	6.677	0.696	7.927
Eloundou et al., GPT4	Agg. hours of full-time, full-year workers	0.463	6.741	0.714	8.066
Eloundou et al., Human	Total employment	0.418	6.390	0.711	8.129
Eloundou et al., Human	Agg. hours of full-time, full-year workers	0.416	6.420	0.723	8.215
Felten, Raj, and Seamans, AIOE	Agg. hours of full-time, full-year workers	0.381	6.045	0.734	8.334
Felten, Raj, and Seamans, AIOE	Total employment	0.375	5.947	0.713	8.210
Handa et al., observed exposure	Total employment	0.373	6.088	0.558	7.210
Handa et al., observed exposure	Agg. hours of full-time, full-year workers	0.349	5.937	0.540	7.105
CEA, performance requirements	Agg. hours of full-time, full-year workers	0.303	5.894	0.386	6.178
CEA, performance requirements	Total employment	0.265	5.472	0.322	5.640
CEA, AI exposure	Agg. hours of full-time, full-year workers	0.254	4.965	0.298	5.210
CEA, AI exposure	Total employment	0.239	4.776	0.284	5.086
Eisfelt, Schubert and Zhang, supplemental task exposure	Total employment	0.194	4.358	0.301	5.070
Eisfelt, Schubert and Zhang, supplemental task exposure	Agg. hours of full-time, full-year workers	0.173	4.137	0.290	4.960

Source: Business Trends and Outlook Survey (BTOS), American Community Survey, Quarterly Workforce Indicators, Felten, Raj, and Seamans (2021),Eloundou et al. (2023), Eisfelt, Schubert, and Zhang (2023),Council of Economic Advisers (2024), Handa et al. (2025). Each row shows the result of an unweighted regression of subsector and sector-level AI adoption on the indicated measure of employment-weighted average potential AI exposure or AI-related performance requirements for the subsector/sector. Employment weights are based on total 2022 QWI employment reported in detailed industry by state cells. Latest available data averages the last four rounds of BTOS as of April 9, 2026.