

Generative AI and Early-Career Hiring: Counts, Entry Rates, and the Employment Base

Yi Wu

wuyi@utexas.edu

The University of Texas at Austin

Research Article

Keywords: generative AI, early-career workers, stable hiring, worker flows, industry exposure, labor demand

Posted Date: October 19th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11342200/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Generative AI and Early-Career Hiring: Counts, Entry Rates, and the Employment Base

Abstract

Stable new hires of workers aged 22–34 fell relative to those of workers aged 35–54 after 2022 in U.S. industries with greater exposure to language-model AI. Using the Quarterly Workforce Indicators, I compare stable with all new hires and separate the stable-hire count into the published entry rate and its stable-employment denominator. With equal industry weights, the early-versus-established post-period coefficient per standard deviation of industry exposure is -0.0165 for stable hires, compared with -0.0081 for all hires; their difference is statistically significant. On a common positive sample, the stable-hire log coefficient of -0.01796 consists of a $+0.00310$ entry-rate component and a -0.02106 employment-base component. The employment-base differential was already trending downward before the pandemic, partly offset in the count by an opposing movement in the measured rate. A separate comparison with the extrapolated 2017–2019 trend finds a more negative count slope after 2023, although the rate contribution varies with industry weights and the chosen preperiod. The average post-2022 stable-hire count differential is therefore concentrated in the stable-employment denominator, whose exposure-related age gradient predates ChatGPT; a distinct decline in the QWI-defined entry rate is less firmly supported. Employment-weighted count and base estimates are less precise.

Keywords: generative AI, early-career workers, stable hiring, worker flows, industry exposure, labor demand

JEL Classification: J23 , J24 , O33 , C23

1 Introduction

A decline in new-hire counts need not imply an equal decline in hires per unit of the stable-employment denominator used to construct the hiring rate. The distinction is relevant for early-career labor demand: changes in workforce size and age composition can alter the count even without a comparable change in the measured entry rate. Counts and employment-scaled rates accordingly describe different features of labor-market adjustment.

Using the Quarterly Workforce Indicators (QWI), Tucker (2026) documents weaker hiring among young workers in industries with greater AI exposure and discusses hiring rates alongside the changing employment base. I examine the more specific contrast

between stable and all new hires, date each aggregate hiring series to its hiring-origin quarter, and decompose stable-hire counts into the published QWI entry rate and employment denominator. The question is not whether counts differ from rates, but how the stable-hiring differential is distributed across these measurements and when its components began to diverge.

The post-2022 exposure differential is more negative for stable new hires than for all new hires among workers aged 22–34 relative to those aged 35–54: -0.0165 versus -0.0081 . The difference is -0.0084 ($p = 0.0068$). This comparison concerns two aggregate flow counts measured for the same hiring-origin quarter, not the probability that an individual worker remains employed.

On a common positive sample, the log stable-hire coefficient of -0.01796 comprises $+0.00310$ for the QWI-defined entry rate and -0.02106 for its stable-employment denominator. Before the pandemic, the negative exposure-related base trend was partially offset by the rate component, leaving the count path comparatively flat. Separately, the count slope after 2023 is more negative than an extrapolation of the 2017–2019 trend, although the corresponding rate component varies with weights and the selected preperiod.

The average post-period hiring shortfall and the departure from an earlier trend answer different questions. Taken together, they show why the employment denominator deserves attention alongside worker flows when studying adjustments associated with technological exposure.

2 Related Literature and Empirical Background

2.1 Technology and Early-Career Hiring

Changes in task costs may appear in hiring and separations before they are readily visible in aggregate employment (Acemoglu & Restrepo, 2020; Autor, 2015). Generative AI affects many language-intensive tasks, yet its consequences for employment flows are not determined by its effect on task productivity alone. Experiments in professional writing and evidence from customer support report productivity improvements, including for relatively inexperienced employees (Brynjolfsson, Li, & Raymond, 2025; Noy & Zhang, 2023). Whether that improvement raises or reduces demand for new workers is a separate labor-market question.

Evidence on early labor-market responses differs across outcomes and settings. Humlum and Vestergaard (2025) find changes in task content alongside modest aggregate wage and hours effects; Kauhanen and Rouvinen (2025) also report limited early aggregate labor-market responses in Finnish data. In contrast, the U.S. QWI analysis of Tucker (2026) documents differential hiring declines among younger workers. That hiring result provides the empirical starting point here.

2.2 Industry Exposure and Worker-Flow Measurement

I use the Language-Model AI Industry Exposure (LM-AIIE) index of Felten, Raj, and Seamans (2023), which maps occupational exposure to language-model capabilities into industry-level variation, drawing on the AI-exposure framework of Felten, Raj, and

Seamans (2021). The measure represents potential task exposure, not observed firm adoption. A comparison with late-2025 Business Trends and Outlook Survey (BTOS) reports provides an ex-post check on its relationship with reported AI use, but cannot establish the timing of adoption in 2023.

The distinction between a composite flow and its components also bears on pre-trend assessment. Preperiod tests alone do not establish counterfactual parallel trends (Rambachan & Roth, 2023; Roth, 2022); moreover, a relatively flat count can conceal offsetting movements in the published rate and its denominator. The QWI accounting relationship can locate those movements arithmetically, without identifying their separate economic causes.

3 Data and Measurement

3.1 LM-AIIE and the Stable Industry Sample

The sample covers 2017–2025. QCEW and QWI change from NAICS 2017 to NAICS 2022 beginning in 2022, potentially changing the composition of a four-digit industry even if its code is unchanged. To avoid artificial breaks, I retain only industries with identical four-digit titles and constituent six-digit code–title sets in the two classification vintages, using the BLS hierarchy crosswalk. Of 286 four-digit industries mapped to LM-AIIE, 218 meet this strict criterion. Exposure is standardized within the retained set, so a unit increase represents one standard deviation across these 218 industries. Data sources and construction details are summarized in Online Resource 1.

3.2 Quarterly Workforce Indicators and Flow Timing

The QWI are public data from the U.S. Census Bureau’s Longitudinal Employer-Household Dynamics program (U.S. Census Bureau, 2026b). I use private-sector, nonseasonally adjusted observations by state, four-digit industry, age, and quarter, pooled across sex. The main comparison combines ages 22–24 and 25–34 as the early-career group and ages 35–44 and 45–54 as the established-career group. These adjacent age categories reduce contrasts at the extremes of the age distribution during the pandemic; age nevertheless measures neither tenure nor individual labor-market entry. Original-bin estimates appear below.

QWI distinguishes all new hires (HirN) from stable new hires (HirNS), recorded when a new job first qualifies as full-quarter employment. Let EmpS and EmpSpv denote current-quarter and previous-quarter stable employment, respectively. At the aggregation level used here, I recompute the QWI-defined stable-new-hire rate from the summed counts and denominator components. I denote the aggregated rate in percentage units by $\text{HirNSR}_{jgt}^{(\%)}$. For the accounting identity I express that rate as a fraction:

$$R_{jgt} = \frac{\text{HirNSR}_{jgt}^{(\%)}}{100},$$

$$\text{HirNS}_{jgt} = R_{jgt} B_{jgt}, \quad B_{jgt} = \frac{\text{EmpSpv}_{jgt} + \text{EmpS}_{jgt}}{2}. \quad (1)$$

The denominator B_{jgt} scales the flow by average stable employment; it is neither vacancies nor a pool of potential applicants, and R_{jgt} is not an individual hiring hazard. Importantly, B_{jgt} is not entirely predetermined: current-quarter EmpS includes employment newly attaining full-quarter status, so part of the denominator is measured contemporaneously with stable hiring. The identity separates accounting components, not independently identified causal channels.

Because HirNS is recorded when an individual first attains full-quarter status, the corresponding job generally began in the preceding quarter. I therefore associate the values of HirNS, HirNSR, and StableBase reported in $t + 1$ with hiring-origin quarter t ; HirN stays dated to quarter t . For example, stable new hires reported in 2023Q1 are placed alongside all new hires in 2022Q4 for the origin-quarter comparison. This aligns aggregate series by their published timing definitions; it does not link individual records. The three accounting components are shifted together.

Before constructing broad age categories or aggregating over states, I sum hiring numerators and the two denominator components, then recompute the QWI-defined rate as $100 \times \text{HirNS}/B$ in percentage-point units or HirNS/B in fractional units. A simple average of age- or state-cell rates would not preserve Equation 1. The principal analysis aggregates available state cells to an industry-by-career-stage-by-quarter panel. The raw QWI files include all 50 states and the District of Columbia, but the available vintage has no usable 2017-onward Alaska records and Michigan coverage ends in 2021Q4. Michigan’s exit may therefore change the composition of the available-state aggregate across the comparison period. Suppressed and unavailable cells are not imputed, and the resulting aggregation is not a complete, balanced national census. The state-cell comparison is in Section 7.1; coverage and the last usable hiring-origin quarters are discussed in Online Resource 1.

3.3 QCEW, BTOS, and Prices

From private-sector county records in the Quarterly Census of Employment and Wages (QCEW), I construct state–industry–quarter employment and real average weekly wages (U.S. Bureau of Labor Statistics, 2026b). Quarterly employment averages the three monthly levels; weekly wages equal quarterly payroll divided by 13 times average employment and are deflated using CPI-U to 2022Q4 dollars (U.S. Bureau of Labor Statistics, 2026a). These outcomes pool ages and serve as an aggregate benchmark rather than a direct measure of early-career hiring.

For BTOS (U.S. Census Bureau, 2026a), I aggregate four-digit LM-AIIE to three-digit subsectors using 2022Q1–Q3 employment weights and compare it with reported current AI use. In the available extraction, that item begins only in late 2025, so the check concerns the relevance of potential exposure to subsequent reported use, not the date when employers adopted AI. Additional diagnostics are in Online Resource 1.

4 Empirical Design

4.1 Early-Career Difference-in-Differences by Exposure

I compare changes in the early-versus-established outcome gap across industries with different LM-AIIE values. For industry j , career group g , and quarter t , the main regression is

$$Y_{jgt} = \beta A_j \times Early_g \times Post_t + \lambda_{jt} + \mu_{gt} + \phi_{jgq(t)} + \varepsilon_{jgt}, \quad (2)$$

where A_j is standardized exposure, $Early_g$ equals one for ages 22–34, and $Post_t$ equals one from 2023Q1 onward. The transition quarter 2022Q4 is excluded from static regressions, leaving 2017Q1–2022Q3 as the comparison period, including the pandemic and subsequent recovery. Hiring counts enter as inverse-hyperbolic-sine (IHS) transformations; the exact decomposition uses natural logs on a common strictly positive sample.

Industry-by-quarter effects (λ_{jt}) absorb shocks shared by the two career groups within an industry. Career-by-quarter effects (μ_{gt}) absorb national age-group movements, while industry-by-career-by-quarter-of-year effects ($\phi_{jgq(t)}$) account for group-specific seasonality within industries. Standard errors are clustered by four-digit NAICS industry, where exposure varies. The coefficient compares the change in the early-minus-established gap across industries one standard deviation apart in exposure. It is not the absolute change in young-worker hiring; a causal interpretation would also require ruling out other exposure-related changes in age-specific industry paths.

The baseline weights four-digit industries equally. The alternative uses 2022Q1–Q3 employment weights, placing greater emphasis on larger industries and changing the weighted comparison substantially. The two estimands need not coincide.

4.2 Same-Season Event Study

The early-minus-established hiring gap exhibits industry-specific seasonal patterns. A single omitted quarter can obscure the movement of interest, so I compare each quarter with the same calendar quarter in 2019:

$$Y_{jgt} = \sum_{\tau \in T \setminus \mathcal{R}} \beta_{\tau} A_j \times Early_g \times \mathbf{1}\{t = \tau\} + \lambda_{jt} + \mu_{gt} + \phi_{jgq(t)} + \varepsilon_{jgt}, \quad (3)$$

where $\mathcal{R} = \{2019Q1, 2019Q2, 2019Q3, 2019Q4\}$. Each quarter of 2019 is a normalization point for its own season; the plotted zeros for 2019 are imposed references, not estimated effects. The 2017–2018 quarters provide pre-COVID comparisons relative to 2019. Online Resource 1 reports the alternative age-group and single-reference diagnostics: the adjacent group comparison attenuates pandemic-era contrasts, and same-season normalization reduces recurring oscillations.

Figure 1 gathers the main specification comparisons. The stable-hire differential remains negative as the age comparison, seasonality controls, and timing alignment change. The last three estimates move to the common positive log sample needed for the accounting analysis.




Fig. 1: Specification Comparisons Leading to the Stable-Hiring Decomposition

Notes: The first four points compare IHS stable-hire counts under alternative age groups, seasonality controls, and origin-quarter timing. The final three points use logs on a common positive sample, with identical regressors and weights. Estimates spanning the two transformations should not be read as a single fixed-outcome specification. All standard errors are clustered by four-digit industry.

4.3 Post-Period Averages and Trend Decompositions

For any identical positive sample and linear-regression design, Equation 1 implies $\log H = \log R + \log B$. The corresponding coefficients add exactly when the observations, weights, regressors, and aggregation level are identical:

$$\beta_H^{post} = \beta_R^{post} + \beta_B^{post}, \quad (4)$$

$$\gamma_H^{pre} = \gamma_R^{pre} + \gamma_B^{pre}, \quad (5)$$

$$\Delta\gamma_H^{post} = \Delta\gamma_R^{post} + \Delta\gamma_B^{post}. \quad (6)$$

These identities hold for the static post-period coefficient and, in separate regressions, for the 2017–2019 pre-COVID slope and the post-2023 deviation from that slope. The static comparison includes the pandemic years in its reference period. The segmented-trend exercise instead fits 2017–2019 and 2023–2025, excluding 2020–2022; it measures deviations from an extrapolated path, not a causally identified break. Its full estimating equation is given in Online Resource 1. As a result, the accounting identity disciplines interpretation but cannot identify the forces changing the components.

4.4 Age-Pooled QCEW Benchmark

For the age-pooled QCEW comparison, I estimate

$$Y_{sijt} = \theta A_j \times Post_t + \lambda_{st} + \psi_{sjq(t)} + u_{sijt}, \quad (7)$$

where λ_{st} are state-by-quarter and $\psi_{sjq(t)}$ are state-by-industry-by-season effects. Level estimates use 2022Q1–Q3 state-industry employment weights; four-quarter changes also describe year-on-year movements. This benchmark is not the career-stage triple difference in Equation 2.

5 Main Results

5.1 Stable New Hires versus All New Hires

Stable new hires exhibit a more negative early-career exposure differential than all new hires. After aligning the series to their hiring-origin quarter, the IHS coefficient is -0.0165 (SE 0.0047, $p < 0.001$) for stable hires, compared with -0.0081 (SE 0.0057, $p = 0.157$) for all hires. The formal stable-minus-all contrast is -0.0084 (SE 0.0031, $p = 0.0068$), estimated using the same sample and regressors. Figure 1 shows that the negative stable-hire coefficient is not generated by the principal timing or age-comparison choices. Figure 2 displays the two event-study paths. These are contrasts between aggregate flows, not individual job-retention probabilities.




Fig. 2: Event Paths for Stable and All New Hires

Notes: Each plotted coefficient is an IHS early-versus-established contrast associated with one standard deviation of industry exposure, relative to the corresponding 2019 season. Stable hires are dated to the quarter of initial hiring. The four 2019 seasonal references are normalizations, not estimated zero effects. Industries are equally weighted and standard errors are clustered by NAICS4 industry.

5.2 Stable-Hire Counts, Entry Rates, and Employment Base

The stable-employment-denominator event path closely follows the stable-hire-count path in Figure 3, and its exposure-related differential was already declining before 2023. In most post-2022 quarters, the official entry-rate estimates are closer to zero and comparatively imprecise.

The common positive sample retains 14,176 observations, 98.8% of the corresponding aligned IHS sample, and all 218 industries. Its log-count coefficient (-0.01796 , SE 0.00528) is close to the IHS count estimate (-0.01650). The common-sample coefficient satisfies

$$\underbrace{-0.01796}_{\text{log count}} = \underbrace{+0.00310}_{\text{log official rate}} + \underbrace{(-0.02106)}_{\text{log employment base}} . \quad (8)$$

The count and base estimates have confidence intervals below zero; the entry-rate interval contains zero. The post-2022 stable-hire count differential is therefore concentrated in the stable-employment denominator, whose exposure-related age gradient predates ChatGPT. Because current-quarter stable employment enters this denominator, the accounting allocation does not identify a separate stock shock that caused hiring to fall. All coefficients compare age groups across differentially exposed industries; none represents an unconditional 2.1% decline in young-worker employment.

On the *same* 14,176-observation positive sample, the coefficient is $+0.00310$ (SE 0.00452 , $p = 0.494$) when the QWI-defined rate is logged and -0.03579 percentage points (SE 0.03908 , $p = 0.361$) when the rate is measured in levels. Their signs therefore differ even with identical observations, controls, and weights. This difference reflects the distinct nonlinear outcome transformations, not a change from fractional to percentage units: multiplying a level rate by 100 changes the coefficient’s scale but not its sign. Neither transformation provides evidence of a statistically reliable decline in the average stable-entry rate. Online Resource 1 reports the same-sample comparison and the original percentage-point benchmark.

5.3 Age Patterns in the Original QWI Categories

Taking ages 45–54 as the comparison group, the log stable-hire coefficients are approximately -0.032 for ages 22–24, -0.018 for ages 25–34, and -0.011 for ages 35–44. The corresponding employment-base coefficients show an even steeper age gradient. The youngest group’s count differential coexists with a positive point estimate for the QWI-defined entry rate, as shown in Figure 4. Thus the age gradient in counts is not merely a gradient in stable hires per unit of incumbent employment. The age categories remain proxies for career stage, not direct measures of experience.

6 Pre-COVID Trends and Later Dynamics

6.1 Offsetting Trends before the Pandemic

In 2017–2019, the exposure-related early-minus-established count slope is small, while its rate and employment-base components move in opposite directions. In log-coefficient




Fig. 3: Stable-Hire Count, Stable-Entry Rate, and Stable-Employment Base

Notes: All three paths are natural-log estimates on the same positive industry-career-quarter sample, with the same weights and fixed effects. The values sum by coefficient under the exact accounting identity. The 2019 quarterly points are imposed same-season references; the other points have 95% confidence intervals clustered by NAICS4 industry.




Fig. 4: Count, Rate, and Base across Original QWI Age Groups

Notes: Coefficients use ages 45–54 as the omitted group and the common positive log-accounting sample. The three panels have different vertical-axis scales; magnitudes should be compared numerically rather than by their visual heights. Standard errors are clustered by NAICS4 industry.

units per quarter, for a one-standard-deviation difference in exposure,

$$-0.000824 = +0.000791 - 0.001615. \quad (9)$$

The base slope is negative and precisely estimated; the count slope is smaller and not significantly different from zero. In the corresponding 2017–2018 lead tests, $p = 0.301$ for the count but $p = 0.0035$ for the base. A comparatively flat count pretrend can therefore coexist with economically distinct movements in its components. It does not, on its own, establish the counterfactual trend required for a causal exposure interpretation.

6.2 Post-2023 Deviations from the Pre-COVID Trend

The segmented comparison asks whether the post-2023 trajectory differs from a linear extrapolation of 2017–2019; it is separate from the average post-period regression. Under equal industry weights, the estimated change in the count slope is -0.00195 log-coefficient units per quarter, with components

$$-0.00195 = -0.00119 - 0.00076. \quad (10)$$

The count slope change is negative and relatively precise. The entry-rate slope change is marginally different from zero under this specification, whereas the base slope-change interval is wider. These deviations depend on extrapolating through the unmodeled pandemic period, so they do not establish a new AI-driven break.




Fig. 5: Pre-COVID Slope Decomposition, 2017–2019

Notes: Quarterly slope coefficients for the early-minus-established difference per one-standard-deviation increase in exposure during 2017–2019. The same positive sample and design are used for all three log components, so the slopes add exactly. Whiskers show 95% confidence intervals clustered by NAICS4 industry.




Fig. 6: Exact Accounting Decomposition of the Post-2023 Slope Change

Notes: Change in the exposure-by-early quarterly slope from 2023 onward relative to a fitted 2017–2019 linear trend (not the static post-period level difference). 2020–2022 are excluded from this regression. Estimates use equal industry weights and the same log-accounting sample; 95% confidence intervals are clustered by NAICS4 industry.

6.3 Trend Windows, Weighting, and Placebos

Across the three pre-COVID windows—2017–2018, 2017–2019, and 2018–2019—the equal-weighted post-2023 count-slope deviation remains negative. The division between rate and base is less stable when the preperiod or employment weights change. This matters because the pre-2022 employment-weight Herfindahl index is approximately 0.026, equivalent in a concentration calculation to about 38 equally weighted industries, although the regression still includes 218 industry clusters. The 20 largest industries represent more than half the baseline weight.

None of the three reported pre-COVID placebo dates (2018Q1, 2018Q3, and 2019Q1) yields a count-slope deviation as negative as the post-2023 exercise. These checks are not an exhaustive placebo distribution. The 2023–2025 residuals in Online Resource 1 differ from the 2019-normalized event-study paths: they describe deviations from the extrapolated pre-COVID trend. On that benchmark, base residuals are mostly positive while rate residuals become more negative later, with the count residual equal to their sum.

7 Robustness and Additional Evidence

7.1 Available-State Aggregation versus State Cells

To assess whether aggregation conceals state heterogeneity or public-data suppression, I re-estimate the count, rate, and base regressions on state-by-industry-by-career-by-quarter cells, with state-by-industry-by-quarter, state-by-career-by-quarter, and state-by-industry-by-career-by-season effects. Under equal weights, the count and base coefficients are -0.0182 and -0.0222 , compared with -0.0180 and -0.0211 in the aggregated analysis. The count and base remain negative under employment weighting; the entry-rate point estimate is close to zero and changes sign. Table 1 records the four comparisons and their different industry support.

Table 1: Average Post-2022 Decomposition across Estimands

Estimand	log count	log entry rate	log stable base	N	clusters
Aggregated, equal wt.	-0.01796 (0.00528)	$+0.00310$ (0.00452)	-0.02106 (0.00500)	14,176	218
Aggregated, emp. wt.	-0.01291 (0.00876)	-0.00314 (0.00335)	-0.00977 (0.01091)	14,176	218
State cells, equal wt.	-0.01820 (0.00479)	$+0.00404$ (0.00261)	-0.02223 (0.00524)	444,090	217
State cells, emp. wt.	-0.01127 (0.00857)	-0.00244 (0.00316)	-0.00883 (0.01061)	442,552	217

Notes: Industry-clustered standard errors appear in parentheses. Every row uses identical samples, regressors, and weights for the three log components. The aggregated panel sums available state records but is not a balanced national total. Employment weighting changes the estimand; none of the three employment-weighted coefficients is statistically distinguishable from zero at the 5% level. The log-accounting identity holds to numerical precision within each row; displayed values may differ slightly because of rounding. The state-cell sample supports fewer industry clusters than the aggregated specifications.

As a check on changing geographic support, excluding Michigan before constructing the aggregated panel gives a log-count coefficient of -0.01843 (SE 0.00525) and a log-base coefficient of -0.02152 (SE 0.00486). Dropping the inherited requirement that stable separations be observed for every original age bin yields -0.01516 (SE 0.00458) for counts and -0.02098 (SE 0.00496) for the base. Both modifications together preserve

the negative count and base differentials. The rate remains imprecisely estimated in all four versions; complete results are in Online Resource 1.

7.2 Other Hiring and Separation Flows

Other QWI flow outcomes do not show a uniform change in the same direction among younger workers. The stable-separation-rate coefficient is positive but does not attain the conventional 5% significance threshold ($p = 0.071$); the stable hiring-flow-rate estimate is small and imprecise. These are percentage-point rates with QWI-specific denominators, not natural-log count coefficients. Table 2 also reports the formal stable-minus-all comparison underlying the main result.

Table 2: Other Hiring and Stable-Employment Flow Outcomes

Outcome	Estimate	SE	p value	N
Stable new hires HirNS_{t+1} (IHS)	-0.01650	0.00469	< 0.001	14,348
All new hires HirN_t (IHS)	-0.00807	0.00568	0.157	14,348
Stable – all hires (IHS contrast)	-0.00843	0.00309	0.0068	14,348
Stable-entry rate (p.p.)	-0.0351	0.0391	0.370	14,178
Stable-separation rate (p.p.)	+0.0572	0.0315	0.071	14,628
Stable hiring-flow rate (p.p.)	-0.0367	0.0427	0.392	13,964

7.3 Age-Pooled Wages, Employment, and Exposure Relevance

In the QCEW benchmark, the post-period estimates are about -0.7 log points for real weekly wages and $+0.7$ log points for employment, both imprecise. The year-over-year growth specifications are also near zero, while the pre-COVID QCEW event paths contain exposure-related differences. As these outcomes pool workers of all ages, they do not test the early-versus-established QWI hiring differential. Detailed coefficients and event-study plots appear in Online Resource 1.

At the three-digit subsector level, a one-standard-deviation difference in LM-AIIE corresponds to roughly eight percentage points more reported current AI use in the late-2025 BTOS check. This supports the relevance of the exposure index to differences in subsequent use. It does not show which firms adopted AI at the start of the post-2022 period or separate adoption from other exposure-related industry changes.

8 Discussion

A smaller stable-employment denominator need not signal a contemporaneous reduction in hiring per unit of that denominator. It can reflect accumulated changes in employment composition, including earlier inflows, separations, and aging across QWI categories. But the denominator also contains current-quarter stable employment and is partly measured alongside the stable-hiring flow. The components are consequently informative descriptions of the hiring shortfall, not independent structural mechanisms.

The longer-running exposure-related age gradient in the denominator places the post-2022 count differential in a broader period of workforce adjustment. Pandemic recovery and other industry characteristics correlated with potential AI exposure remain alternative explanations. Separately, the extrapolated trend comparison allows for a further post-2023 rate deviation under some specifications, but that evidence is more sensitive to the benchmark and weights. The published QWI rate should be understood as a flow scaled by stable employment, not as an individual hiring probability.

9 Conclusion

In more language-model-exposed industries, stable new hires of early-career workers fell relative to those of established workers after 2022. Most of the average differential is located in the stable-employment denominator, whose age-related exposure trend was already negative before 2023; the QWI-defined entry rate exhibits no comparably robust average decline.

Changes after 2023 may also depart from the pre-COVID trajectory, but the inferred rate component depends on the extrapolation and the weighting of industries. The hiring count alone cannot tell us when the employment composition underlying that count began to change, or whether hiring per unit of stable employment fell at the same time.

AI Use Declaration

GPT-5.6sol and GPT-6sol were used to assist with editing empirical analysis code and with translation and language polishing during manuscript preparation. Responsibility for the empirical results, their interpretation, and the final manuscript remains with the author.

References

- Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188–2244, <https://doi.org/10.1086/705716>
- Autor, D.H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30, <https://doi.org/10.1257/jep.29.3.3>
- Brynjolfsson, E., Li, D., Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942, <https://doi.org/10.1093/qje/qjae044>

- Felten, E.W., Raj, M., Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217, <https://doi.org/10.1002/smj.3286>
- Felten, E.W., Raj, M., Seamans, R. (2023). *How will language modelers like ChatGPT affect occupations and industries?* (Working Paper No. 4375268). SSRN. Retrieved from <https://doi.org/10.2139/ssrn.4375268>
- Humlum, A., & Vestergaard, E. (2025). *Still waters, rapid currents: Early labor market transformation under generative AI* (NBER Working Paper No. 33777). National Bureau of Economic Research. Retrieved from <https://doi.org/10.3386/w33777> (Revised March 2026)
- Kauhanen, A., & Rouvinen, P. (2025). Assessing early labour market effects of generative AI: Evidence from population data. *Applied Economics Letters*, 1–4, <https://doi.org/10.1080/13504851.2025.2513973>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192, <https://doi.org/10.1126/science.adh2586>
- Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. *The Review of Economic Studies*, 90(5), 2555–2591, <https://doi.org/10.1093/restud/rdad018>
- Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3), 305–322, <https://doi.org/10.1257/aeri.20210236>
- Tucker, L.C. (2026). *You're (not) hired: Artificial intelligence and early career hiring in the Quarterly Workforce Indicators* (CES Working Paper Nos. 26–27). U.S. Census Bureau, Center for Economic Studies. Retrieved from <https://www2.census.gov/library/working-papers/2026/adrm/ces/CES-WP-26-27.pdf>
- U.S. Bureau of Labor Statistics (2026a). *Consumer Price Index for All Urban Consumers: All Items in U.S. City Average (CUUR0000SA0)*. Retrieved from <https://www.bls.gov/cpi/>
- U.S. Bureau of Labor Statistics (2026b). *Quarterly Census of Employment and Wages*. Retrieved from <https://www.bls.gov/cew/>

- U.S. Census Bureau (2026a). *Business Trends and Outlook Survey*. Retrieved from <https://www.census.gov/hfp/btos/data>
- U.S. Census Bureau (2026b). *Quarterly Workforce Indicators, Longitudinal Employer-Household Dynamics Program*. Retrieved from <https://www.census.gov/programs-surveys/ces/data/public-use-data/quarterly-workforce-indicators.html>

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [Supplementarymaterial.pdf](#)