

Online Resource 1: Supplementary Data and Robustness Evidence Generative AI and Early-Career Hiring: Counts, Entry Rates, and the Employment Base

1 Source data, exposure mapping, and coverage

Table S1: Source datasets and their roles

Source	Unit, measures, and use
LM-AIIE	Industry-level language-model task exposure; treatment intensity across retained NAICS4 industries.
QWI	State by NAICS4 by age by quarter; new and stable hires, stable-employment components and worker-flow rates.
QCEW	County by NAICS4 by quarter; employment, quarterly payroll, and reconstructed weekly wages.
BTOS	Three-digit subsector by survey period; late-2025 current AI-use shares for exposure relevance.
CPI-U	Monthly national consumer-price index; deflates QCEW wage outcomes to 2022Q4 dollars.

The analysis uses the Language-Model Artificial Intelligence Industry Exposure (LM-AIIE) measure mapped to four-digit North American Industry Classification System (NAICS4) codes. To prevent classification revisions from being mistaken for economic changes, the stable-industry baseline requires the same NAICS4 title and the same constituent six-digit code–title sets in the 2017 and 2022 NAICS editions. Of 286 NAICS4 industries mapped to LM-AIIE, 218 satisfy this strict rule. Exposure is standardized across these 218 industries. The raw QCEW processing retains private-sector ordinary counties with exact four-digit codes and aggregates employment and payroll to state–industry–quarter observations.

The raw QWI directory includes 51 files (the 50 states and Washington, DC). The public QWI vintage used in the project contains no usable observations for Alaska from 2017 onward; Michigan coverage ends in 2021Q4. Because Michigan contributes to the preperiod but not the postperiod, state aggregation may inherit a change in geographic composition. An explicit check excluding Michigan from every quarter is reported below, alongside the state-cell robustness. Thus 50 geographic units contribute to the 2017+ QWI panel, with uneven coverage over time. The data are not imputed. A hiring-origin quarter is observed for stable hires only when the next quarter has a published stable-hire record; therefore the last usable origin quarter may precede the last reporting quarter. Only the all-sex records are retained. The four main age bins are 22–24, 25–34, 35–44, and 45–54. The numerator and the two QWI-defined denominator components are summed before recomputing rates. The rate used for the log identity is the fraction HirNS/B ; its percentage-unit counterpart is $100 \times \text{HirNS}/B$. Current-quarter stable employment enters B , so this denominator is partly contemporaneous with stable hiring, rather than a strictly predetermined stock.

For the exact accounting exercise, the common positive sample retains 14,176 industry–career-stage–quarter observations versus 14,348 in the cohort-aligned IHS count analysis (98.8%). The corresponding estimates are -0.01796 (log count, SE 0.00528) and -0.01650 (IHS count, SE 0.00469). Both samples retain all 218 industries. The original QWI accounting-panel construction requires nonmissing HirNS, EmpS, EmpSpv, and stable separations (SepS) in all four underlying age bins. The SepS requirement was inherited from the common worker-flow analysis, not from the accounting identity itself. An additional analysis removes this requirement and reconstructs the aggregate accounting sample from the underlying state–industry–age files. The primary estimates retain the historical specification for direct comparison; the alternate samples are reported in the table below. These restrictions are not evidence of random missingness.

2 Sample-construction robustness checks

The original decomposition can be replicated directly from the state–industry–age QWI panel: the rebuilt baseline has the same 14,176 observations, 218 industry clusters, and count, rate, and base coefficients as the saved analysis panel. To test changing coverage, Michigan is removed from all years *before* the state records are aggregated. To test the inherited flow-completeness rule, the requirement that stable separations be observed in every original age bin is removed while retaining the three variables required by the accounting identity. The final row applies both changes. These checks leave the estimating equation and clustering unchanged.

Table S2: Robustness to Michigan coverage and the stable-separation completeness rule

Sample	Log count	Log rate	Log base	N
Original	-0.01796 (0.00528)	$+0.00310$ (0.00452)	-0.02106 (0.00500)	14,176
No Michigan	-0.01843 (0.00525)	$+0.00309$ (0.00452)	-0.02152 (0.00486)	14,176
No SepS restriction	-0.01516 (0.00458)	$+0.00582$ (0.00362)	-0.02098 (0.00496)	14,270
No Michigan / no SepS	-0.01541 (0.00459)	$+0.00605$ (0.00365)	-0.02146 (0.00482)	14,266

Notes: Four-digit-industry-clustered standard errors in parentheses; 218 industry clusters in all rows. Negative count and base coefficients remain significant at the 1% level in all four specifications; log-rate estimates are not significant at the 5% level. The No Michigan specification retains the same number of aggregated cells, because eliminating Michigan changes the values of the available-state aggregate without necessarily eliminating an entire industry–career-stage–quarter cell. The no-SepS sample retains 94 more observations than the original. All reported three-part coefficients satisfy the log identity up to machine precision. These are descriptive estimates, not identified causal effects of AI adoption.

3 Exposure validation and aggregate labor-market outcomes

Business Trends and Outlook Survey (BTOS) subsector measures become available for the current-AI-use item only in late 2025 in the supplied extraction. These data therefore validate exposure *ex post* and are not used to construct a 2023 adoption treatment. The pooled observation-level coefficient is 0.0803 (SE 0.0120, $p < 0.001$; $N = 921$); averaging the subsector observations across survey periods yields 0.0757 (SE 0.0117; $N = 58$). Figure S1 displays the exposure–reported-use relationship.

Table S3: QCEW static estimates

Outcome	Estimate	SE	<i>p</i> value	<i>N</i>
Real weekly wage level	-0.00697	0.00519	0.181	339,398
Employment level	0.00679	0.01071	0.527	339,398
Year-over-year real wage growth	0.00268	0.00244	0.275	297,879
Year-over-year employment growth	0.00209	0.00514	0.684	297,879

Figure S4 presents the corresponding QCEW wage and employment event-study paths. The joint 2017–2018 preperiod test rejects for wages ($p < 0.001$) and yields $p = 0.027$ for employment. These aggregate age-pooled outcomes do not identify an early-career hiring margin.

4 Age comparisons and seasonal reference periods

The primary study compares ages 22–34 with ages 35–54, rather than comparing prime-age workers with both very young and older workers whose pandemic-era trajectories differ substantially. The event study uses the same quarter of 2019 as the reference for each calendar quarter, with industry–career-stage–quarter-of-year fixed effects. The alternative single-2022Q3-reference specification exhibits greater apparent seasonal patterning. Figure S2 compares the age-group definitions, and Figure S3 contrasts the two reference-period choices.

5 Other hiring and worker flows

Table S4: Worker-flow benchmarks and the formal stable-minus-all comparison

Outcome	Estimate	SE	<i>p</i> value	<i>N</i>
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 minus all new hires	-0.00843	0.00309	0.0068	14,348
Stable hires, employment-weighted	-0.01293	0.00876	0.141	14,348
Official stable-entry rate (p.p.)	-0.0351	0.0391	0.370	14,178
Official stable-separation rate (p.p.)	0.0572	0.0315	0.071	14,628
Official stable hiring-flow rate (p.p.)	-0.0367	0.0427	0.392	13,964

Table S5: Official stable-entry rate: same-sample functional-form comparison

Outcome	Coefficient	Clustered SE	<i>p</i> value	<i>N</i>
Log QWI-defined stable-entry rate	+0.00310	0.00452	0.494	14,176
Stable-entry rate (percentage points)	−0.03579	0.03908	0.361	14,176

Notes: Both estimates use the identical positive accounting sample, fixed effects, and equal-industry weights; uncertainty is clustered by four-digit industry (218 clusters). The difference in sign is attributable to the nonlinear outcome transformation, not the factor of 100 between percentage and fractional units. Neither coefficient rejects zero at conventional levels.

The stable-minus-all estimate is a contrast of two aligned QWI aggregate flows, not the probability that a new hire survives to full-quarter status. For example, HirNS_{t+1} and HirN_t are aligned

to origin quarter t without matching persons. The QWI-defined stable-hiring rates use their own denominators and must not be compared directly with log coefficients. For the stable-entry rate, the original percentage-point regression has 14,178 observations versus 14,176 in the common positive log sample. Re-estimating the percentage-point specification on precisely the 14,176-observation log sample gives -0.03579 (SE 0.03908, $p = 0.361$), while the log-rate regression gives $+0.00310$ (SE 0.00452, $p = 0.494$). The sign difference therefore persists when sample support is identical; it reflects the different functional forms, not the two-observation sample difference or the simple percent-versus-fraction conversion. Neither coefficient is statistically distinguishable from zero. The entry- and separation-rate event studies appear in Figures S12 and S13, while Figure S5 places the rate estimates alongside their preperiod benchmarks.

6 Accounting samples and the original age categories

On the same positive sample, with the same regression matrix and weights, $\log H = \log R + \log B$ implies exact additivity of every linear-regression coefficient. The baseline postperiod differential satisfies

$$-0.01796 = +0.00310 - 0.02106.$$

An alternative “clean years” sample gives $-0.02407 = +0.00425 - 0.02832$ ($N = 9,466$). The reference group for the detailed age comparison is ages 45–54. For ages 22–24 the log stable-hiring count coefficient is about -0.032 , while the base differential is larger in magnitude and the entry-rate point estimate is positive. Figure S6 overlays the three accounting paths, and Figure S7 reports the corresponding paths by original age group.

7 Pre-COVID slopes and post-2023 trend sensitivity

The exact quarterly pre-COVID slope coefficients satisfy $-0.000824 = +0.000791 - 0.001615$ for count, rate, and base, respectively. The segmented-trend design uses the 2017–2019 and 2023–2025

Table S6: 2017–2019 exposure-by-early-career trend slopes

Component	Slope/quarter	SE	95% CI	p value
Stable-hire count	-0.000824	0.000601	$[-0.002002, 0.000355]$	0.172
Official stable-entry rate	+0.000791	0.000581	$[-0.000348, 0.001930]$	0.175
Stable-employment base	-0.001615	0.000386	$[-0.002372, -0.000857]$	< 0.001

quarters, excluding the entire 2020–2022 interval. Define d_t as quarters since 2023Q1 (zero in 2023Q1), $P_t = \mathbf{1}\{t \geq 2023Q1\}$, and $Z_{jg} = A_j \times \text{Early}_g$. For each log-accounting outcome Y_{jgt} , the Stata specification is

$$Y_{jgt} = \gamma Z_{jg} d_t + \delta Z_{jg} P_t + \kappa Z_{jg} P_t d_t + \lambda_{jt} + \mu_{gt} + \phi_{jgq(t)} + \varepsilon_{jgt}. \quad (1)$$

Here γ is the fitted pre-COVID differential slope, δ is the intercept deviation at 2023Q1 relative to the extrapolation, and κ is the post-2023 slope deviation (not the total post slope). With shared samples, regressors and weights, all three coefficients obey the count–rate–base accounting identity. This descriptive regression does not model the excluded pandemic years. Under equal industry weighting the post-2023 slope change is $-0.0019529 = -0.0011915 - 0.0007613$. The respective confidence intervals for the count, rate, and base slope deviations are $[-0.003130, -0.000776]$,

$[-0.002379, -0.000005]$, and $[-0.001860, 0.000337]$. These are deviations from an extrapolated trend rather than causally identified shocks. Figure S8 compares the alternative pre-COVID trend windows and weighting schemes. Figure S9 shows the three reported pre-COVID placebo dates (2018Q1, 2018Q3, and 2019Q1), not an exhaustive search over possible breaks. Figure S10 plots quarter-specific deviations from the extrapolated trajectory.

8 Geographic aggregation and weighting

Table S7: Accounting coefficients across aggregation levels and weights

Specification	Log count	Log entry rate	Log base	<i>N</i>
Aggregated, equal industry	−0.01796 (0.00528)	+0.00310 (0.00452)	−0.02106 (0.00500)	14,176
Aggregated, employment weighted	−0.01291 (0.00876)	−0.00314 (0.00335)	−0.00977 (0.01091)	14,176
State cells, equally weighted	−0.01820 (0.00479)	+0.00404 (0.00261)	−0.02223 (0.00524)	444,090
State cells, pre-employment weighted	−0.01127 (0.00857)	−0.00244 (0.00316)	−0.00883 (0.01061)	442,552

Notes: Standard errors clustered at the NAICS4 industry level (218 clusters in aggregated rows; 217 in state-cell rows). Each row preserves log-coefficient additivity. The figure labels the national-style sum as an aggregate of available state records; it is not a balanced national census total.

Figure S11 compares the four aggregation and weighting specifications graphically. The baseline aggregated log-count SE is 0.00528; the coefficient additivity restriction does not imply an additivity rule for standard errors. The aggregated specifications include 218 industry clusters while the state-cell versions have 217 clusters because the state-cell and aggregated observation supports differ. The available-state aggregation does not restore missing public QWI records. Preperiod employment weights have Herfindahl index 0.02635 among the 218 industries, implying approximately 37.96 effective treatment-level units. The 5, 10, and 20 most heavily weighted industries account for 25.9%, 38.2%, and 52.4% of total baseline weight, respectively. This concentration explains why reweighting changes both the estimand and its precision.

9 Reproducibility and limitations

The original processing pipeline contains (i) a Python raw-data builder, (ii) Stata estimation code, and (iii) an R plotting script. The source-data construction and estimation steps are specified in the accompanying research scripts; the panels use publicly released aggregates rather than individual records. Limitations include observational exposure differences, potentially confounded postpandemic structural changes, public-use QWI suppression, and age as an imperfect proxy for career experience.

Supplementary Figures

The figure numbers correspond to the filenames supplied with Online Resource 1. Figures were prepared in R (ggplot2 and patchwork) from Stata diagnostic exports. Each figure is supplied as an embedded-font vector PDF and EPS, with a high-resolution TIFF alternative.




Figure S1: LM-AIIE and subsequently reported current AI use




Figure S2: Pandemic recovery patterns under alternative age comparisons




Figure S3: Single-quarter and same-season event-study normalization




Figure S4: Same-season event studies for QCEW real weekly wage and employment




Figure S5: Confidence intervals and preperiod benchmarks for official QWI rates




Figure S6: Combined count, entry-rate, and stable-base event-study paths




Figure S7: Original QWI age-group event-study paths for the three accounting components




Figure S8: Trend-window and employment-weighting sensitivity of post-2023 slope changes




Figure S9: Pre-COVID placebo dates compared with the observed post-2023 break




Figure S10: 2023–2025 quarterly residuals relative to the extrapolated pre-COVID trajectory




Figure S11: Available-state aggregation versus state-cell and employment-weighted estimates




Figure S12: Official stable-entry rate in percentage points




Figure S13: Official stable-separation rate in percentage points