Is Generative AI Seniority-Biased? Evidence from U.S. Occupational Employment Data

APEP Autonomous Research* @olafdrw

March 6, 2026

Abstract

Between 2015 and 2024, entry-level occupations' share of U.S. employment fell from 50.2% to 45.7%, while senior occupations' share rose from 29.0% to 33.3%. I test whether this compositional shift is concentrated in industries with higher artificial intelligence exposure using BLS Occupational Employment and Wage Statistics, O*NET seniority classifications, and Felten-Raj-Seamans AI Occupational Exposure scores. Industries above the median AI exposure experienced 1.8 percentage point larger declines in entry-level employment share after 2022 ($p < 0.05$), while their senior share rose by 2.2 percentage points more ($p < 0.05$). However, an event study reveals a gradual pre-trend: the entry-level share gap between high- and low-AI-exposure industries widened steadily from 2015 onward, predating ChatGPT's release. The results document a robust correlation between AI exposure and seniority composition but cannot cleanly attribute the shift to generative AI specifically.

JEL Codes: J23, J24, O33

Keywords: artificial intelligence, employment composition, seniority, technological change, occupational structure

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch. This paper was generated autonomously. Total execution time: auto (cumulative: auto).

1. Introduction

In November 2022, OpenAI released ChatGPT, and within two months it had 100 million users—the fastest consumer technology adoption in history. By mid-2024, hundreds of publicly traded firms were discussing generative AI in their SEC filings. The technology's potential to reshape labor markets immediately became a central question for economists, policymakers, and workers.

A growing body of experimental evidence demonstrates that generative AI substantially boosts individual productivity, particularly for less experienced workers. Brynjolfsson et al. (2024) show that AI-assisted customer service agents become 14% more productive, with the largest gains for novice workers. Noy and Zhang (2023) find that ChatGPT raises writing productivity by 40%, compressing the gap between high- and low-skilled writers. Peng et al. (2023) document a 55% speedup in coding tasks with GitHub Copilot.

Yet the aggregate labor market implications remain unclear. If generative AI makes junior workers more productive, should we expect firms to hire more of them—or fewer? Recent proprietary data offers a provocative answer. Hosseini Maasoum and Lichtinger (2025) analyze 14 million résumés and find that firms adopting generative AI sharply reduced junior hiring while maintaining or increasing senior headcount, a pattern they term "seniority-biased technological change." This finding challenges the conventional narrative of skill-biased technical change (Katz and Murphy, 1992; Goldin and Katz, 2007) by suggesting that AI's labor market effects operate along the experience dimension rather than the education dimension.

This paper tests whether the seniority-bias pattern appears in entirely independent public data. I use three publicly available datasets: the Bureau of Labor Statistics' Occupational Employment and Wage Statistics (OEWS) for employment by occupation and industry, the O*NET Job Zone classification to assign seniority levels to occupations, and the Felten et al. (2021) AI Occupational Exposure (AIOE) scores to measure industry-level AI intensity. This combination provides a transparent, replicable framework for examining whether employment composition has shifted against junior workers disproportionately in AI-exposed industries.

The main empirical strategy is a difference-in-differences design comparing changes in the entry-level employment share (O*NET Job Zones 1–2) across industries with varying AI exposure, before and after 2022. I supplement this with a triple-difference specification that exploits within-industry variation across seniority tiers.

The results reveal three key findings. First, there has been a large, economically significant shift in the seniority composition of U.S. employment. The share of workers in entry-level occupations fell by 4.5 percentage points between 2015 and 2024, while the share in senior occupations rose by 4.3 percentage points—a reallocation affecting tens of millions of workers. Second, industries with above-median AI occupational exposure experienced a 1.8 percentage point larger decline in entry-level employment share after 2022 ($t = -2.11$, $p < 0.05$), and a 2.2 percentage point larger increase in senior share ($t = 2.08$, $p < 0.05$; Table 3, Column 2). Third, a heterogeneity analysis by within-occupation AI exposure shows that the interaction of high AI exposure with junior status and the post period is strongly negative ($\hat{\beta} = -0.27$, $t = -5.30$), suggesting the seniority bias is concentrated where AI tools are most applicable.

However, the evidence for a clean causal effect of generative AI specifically is limited. An event study plotting the entry-level share differential between high- and low-AI-exposure industries reveals a steady, gradual widening from 2015 onward—well before ChatGPT's release. A placebo test placing fake treatment at 2020 yields a statistically significant coefficient ($t = -2.83$), confirming the pre-trend. The post-2022 pattern appears to be a continuation of a longer-term structural shift rather than a discrete break caused by generative AI adoption.

This paper contributes to several literatures. First, it extends the emerging literature on AI and labor markets (Acemoglu et al., 2022; Webb, 2020; Eloundou et al., 2024) by providing the first independent replication of the seniority-bias finding using public data. The broad consistency with Hosseini Maasoum and Lichtinger (2025)'s proprietary résumé evidence—declining junior shares in high-AI industries—strengthens the empirical regularity, even as my results suggest the trend predates generative AI. Acemoglu et al. (2022) study online vacancy postings and find that AI-exposed occupations experienced slower employment growth, but they do not distinguish effects by seniority within occupations. Webb (2020) constructs exposure measures from patent texts and finds heterogeneous effects across the wage distribution, with high-wage workers potentially more exposed to machine learning than to earlier automation waves. My contribution is to show that the *seniority* dimension—orthogonal to both wages and education—captures a distinct margin of adjustment.

Second, the paper contributes to the literature on skill-biased and routine-biased technological change (Autor et al., 2003, 2006; Autor and Dorn, 2013; Card and DiNardo, 2002). The canonical skill-biased technical change (SBTC) framework (Katz and Murphy, 1992; Goldin and Katz, 2007) predicts that technology complements skilled (college-educated) workers while substituting for unskilled workers. The routine-biased technical change (RBTC) refinement (Autor et al., 2003; Autor, 2015) adds nuance: technology automates routine tasks (both cognitive and manual), leading to labor market polarization with growth at the top and bottom but hollowing in the middle (Autor et al., 2006; Autor and Dorn, 2013). Generative AI potentially disrupts both frameworks because it automates *cognitive* tasks that RBTC classified as non-routine—writing, analysis, coding, and creative work that previously seemed safe from automation (Eloundou et al., 2024). The finding that employment composition shifts along the seniority dimension, not merely the education or routine-task dimension, represents a potential refinement of the canonical framework. If AI substitutes for junior workers' tasks regardless of skill level—replacing entry-level coding, entry-level analysis, and entry-level writing—then "seniority-biased" may better describe the relevant margin than "skill-biased."

Third, the paper connects to the growing literature on firm-level AI adoption and its consequences. Babina et al. (2024) use job postings and earnings calls to measure firm-level AI activity and find that AI-investing firms grow faster and innovate more, but do not examine within-firm seniority composition. Eisfeldt et al. (2023) document that firms with higher GPT exposure experienced larger stock market gains after ChatGPT's release, suggesting investor expectations of productivity benefits. Hui et al. (2024) study freelance platforms and find that generative AI reduced demand for freelancers, with the largest effects for entry-level tasks. My paper complements these studies by using economy-wide employment data rather than firm- or platform-specific samples.

Fourth, this paper demonstrates how public administrative data can be combined to study AI's labor market effects without reliance on proprietary or commercially licensed datasets. The OEWS-O*NET-AIOE pipeline is fully replicable and can be updated annually as new OEWS releases become available. Transparency in data and methods is particularly important for a topic where policy stakes are high and public discourse often outpaces rigorous evidence.

2. Institutional Background

2.1 The Generative AI Shock

Generative AI—large language models, image generators, and code assistants—represents a qualitatively different form of automation from previous waves. Unlike industrial robots that automate physical tasks (Acemoglu and Restrepo, 2020; Graetz and Michaels, 2018) or routine-biased technologies that displace middle-skill clerical work (Autor et al., 2003; Autor and Dorn, 2013), generative AI targets cognitive tasks across the skill spectrum: writing, analysis, coding, customer service, and creative work (Eloundou et al., 2024).

The timeline of adoption is unusually sharp. ChatGPT's public release in November 2022 was followed by rapid enterprise adoption: Microsoft integrated GPT-4 into its Office suite (March 2023), Google launched Bard (March 2023), and thousands of firms began experimenting with or deploying AI assistants. Media reports suggest that over 770 publicly traded firms mentioned generative AI in their SEC filings by early 2025, while my own restrictive full-text search of 10-K filings via the EDGAR EFTS API identifies 27 filings—a smaller count reflecting the specificity of the full-text search endpoint (see Section A).

The key theoretical prediction of Hosseini Maasoum and Lichtinger (2025) is that generative AI is complementary to experience and judgment—the domain of senior workers—while substituting for the routine cognitive tasks that disproportionately define entry-level roles. A senior lawyer benefits from AI-drafted briefs; a junior associate whose primary job was drafting those briefs faces displacement. A senior analyst uses AI to accelerate research; a junior analyst whose value was in data compilation becomes less necessary.

This prediction differs from earlier automation frameworks in a crucial respect. Industrial robots (Acemoglu and Restrepo, 2020) and earlier IT automation (Autor et al., 2003) primarily displaced workers performing routine physical or routine cognitive tasks—assembly line workers, bookkeepers, file clerks. These workers tended to be in the middle of the skill distribution, leading to the polarization pattern documented by Autor et al. (2006) and Autor and Dorn (2013). Generative AI, by contrast, can perform sophisticated cognitive tasks: writing legal briefs, generating code, analyzing data, creating marketing copy, and even conducting basic research. These are precisely the tasks that firms assign to entry-level knowledge workers as training and apprenticeship.

The implications for career development are potentially profound. In traditional professional services, junior workers learned by doing: associates drafted contracts, analysts built models, residents performed procedures under supervision. If AI can perform these tasks more cheaply and reliably, the economic rationale for hiring junior workers to learn through practice is diminished. Agrawal et al. (2019) frame this as a shift in the prediction-judgment boundary: AI expands the set of tasks that can be performed through prediction (pattern recognition, generation from training data) at the expense of tasks that require human judgment (strategic decisions, client relationships, novel problem-solving). Seniority is precisely what distinguishes the judgment-intensive from the prediction-amenable roles within most organizations.

2.2 Prior Waves of Automation and the Labor Market

The relationship between technological change and labor market outcomes has been a central concern in economics for decades. Goldin and Katz (2007) frame the long-run trajectory as a "race between education and technology": when the supply of skilled workers keeps pace with technological demand, wage inequality remains stable; when technology outpaces education, inequality rises. The SBTC literature (Katz and Murphy, 1992; Card and DiNardo, 2002) documented rising returns to education in the 1980s and 1990s as computerization complemented college-educated workers.

Autor et al. (2003) introduced a more granular framework based on the task content of occupations. They showed that routine tasks—those that follow explicit, codifiable rules—are most susceptible to computerization, regardless of whether they are cognitive (bookkeeping) or manual (assembly). Non-routine tasks, whether analytical (management, research) or interpersonal (caregiving, negotiation), proved more resistant. This task-based approach predicted the "job polarization" pattern observed in the 2000s and 2010s (Autor and Dorn, 2013).

Generative AI upends this taxonomy. Large language models can perform tasks that the Autor et al. (2003) framework classified as non-routine cognitive: writing, analysis, creative problem-solving, even basic strategic reasoning. Eloundou et al. (2024) estimate that 80% of the U.S. workforce could have at least 10% of their tasks affected by GPT-class models, with the most exposed occupations being precisely those that require writing, mathematics, and programming—traditionally "safe" from automation. This suggests that the new frontier of automation runs not along the routine/non-routine boundary but along the experience/entry-level boundary: AI can perform entry-level versions of sophisticated cognitive tasks, but struggles with the judgment, relationship management, and strategic oversight that define senior roles.

2.3 Measuring Occupational Seniority

The O*NET program, maintained by the U.S. Department of Labor, classifies every occupation into one of five "Job Zones" based on the preparation required:

- **Zone 1:** Little or no preparation. Examples: food preparation, cashiers, stock clerks.
- **Zone 2:** Some preparation. Examples: customer service, secretaries, bookkeeping clerks.
- **Zone 3:** Medium preparation. Examples: electricians, registered nurses, police officers.
- **Zone 4:** Considerable preparation. Examples: accountants, engineers, database administrators.
- **Zone 5:** Extensive preparation. Examples: surgeons, lawyers, chief executives.

I classify Zones 1–2 as "entry-level," Zone 3 as "mid-level," and Zones 4–5 as "senior." This three-tier classification captures the key distinction: entry-level occupations involve tasks that are more routine, more codifiable, and more susceptible to AI augmentation or substitution, while senior occupations require extensive training, judgment, and interpersonal skills that remain difficult to automate (Deming, 2017; Agrawal et al., 2019).

2.4 Measuring AI Exposure

Felten et al. (2021) construct the AI Occupational Exposure (AIOE) index by linking AI capabilities to specific work activities within each occupation. The index captures the extent to which an occupation's core tasks overlap with capabilities that AI systems can perform, based on the Electronic Frontier Foundation's AI Progress Measurement dataset matched to O*NET work activities. Unlike binary automation risk measures (Frey and Osborne, 2017), the AIOE provides a continuous, task-based measure of exposure that varies at the six-digit SOC level.

I aggregate occupation-level AIOE scores to the two-digit NAICS industry level by taking employment-weighted averages across all occupations within each industry, using fixed 2019 OEWS employment counts as weights to avoid endogeneity with post-treatment employment changes. This produces a single industry-level AI exposure measure that reflects both the occupation mix and the AI susceptibility of each occupation's tasks.

3. Data

This analysis combines four publicly available data sources to construct a panel of industry-level employment by seniority and AI exposure.

3.1 Occupational Employment and Wage Statistics (OEWS)

The OEWS is the Bureau of Labor Statistics' primary source for employment and wage estimates by detailed occupation. The survey covers approximately 1.1 million establishments semiannually, producing annual estimates for over 800 occupations at the national, state, and industry levels. I use national-level OEWS data from 2015 through 2024, providing ten annual observations.

Each OEWS record reports total employment for a given SOC occupation code within a NAICS industry. I aggregate six-digit SOC occupations to the two-digit NAICS industry level after merging with O*NET Job Zone classifications and AIOE scores. The resulting analysis panel contains 25 two-digit NAICS industries observed over 10 years (250 observations for the industry-level DiD) and 750 observations for the triple-difference panel (25 industries $\times$ 3 seniority tiers $\times$ 10 years). For the within-occupation heterogeneity analysis, I further disaggregate by classifying occupations within each industry into AI exposure terciles based on their individual AIOE scores. Each cell is defined by a unique combination of 2-digit NAICS industry (25), seniority tier (3: Entry, Mid, Senior), and AI exposure tercile (3: Low, Medium, High). Of the $25 \times 3 \times 3 = 225$ potential cells, 200 contain at least one occupation;

these 200 cells observed over 10 years yield 2,000 observations.

3.2 O*NET Job Zone Classification

The O*NET database assigns each SOC occupation code to one of five Job Zones. I match O*NET Job Zones to six-digit SOC codes in the OEWS data, successfully classifying 923 occupations. The three-tier seniority classification (Entry-Level, Mid-Level, Senior) is constructed from these Job Zone assignments.

3.3 AI Occupational Exposure (AIOE)

I obtain occupation-level AIOE scores from the replication data accompanying Felten et al. (2021), available on GitHub. The dataset covers 774 six-digit SOC occupations. After matching to OEWS employment data and computing employment-weighted industry means, the resulting industry-level AIOE score ranges from -1.24 to 1.13 across 25 two-digit NAICS industries, with a mean of -0.08 and standard deviation of 0.61.

3.4 Quarterly Census of Employment and Wages (QCEW)

As a supplementary outcome, I use QCEW data from the BLS, which provides quarterly establishment-level employment counts by county and detailed NAICS industry. The QCEW covers virtually all employees covered by unemployment insurance—approximately 95% of U.S. workers. I use national-level data for three-digit NAICS industries from 2015Q1 through 2024Q4. After restricting to industries with non-zero employment and valid AIOE scores, the QCEW regression sample contains approximately 3,600 industry-quarter observations (100 three-digit NAICS industries $\times$ 40 quarters). The QCEW tests whether AI exposure affects *total* industry employment (as opposed to *compositional* shifts within industries).

3.5 SEC EDGAR Full-Text Search

To provide descriptive context on the timing of generative AI adoption, I query SEC EDGAR's full-text search system (EFTS) for 10-K filings mentioning "generative AI," "ChatGPT," "large language model," or related terms. This yields 27 unique filings: zero before 2023, one in 2023, twelve in 2024, and fourteen in early 2025 (through mid-January). While too sparse for use as a primary treatment variable, these data confirm the post-2022 timing of corporate GenAI disclosure.

The low count relative to the 770+ filings found through broader web searches reflects the specificity of the EDGAR EFTS API, which searches the full text of filed documents rather than metadata or exhibit text. The discrepancy suggests that many GenAI mentions appear in exhibits, supplementary materials, or investor presentations rather than the core 10-K text. Regardless, the qualitative pattern—zero mentions pre-2023, rapid growth thereafter—is consistent with the timing of generative AI adoption documented in other sources.

3.6 FRED Macroeconomic Controls

I supplement the industry-level analysis with macroeconomic time series from the Federal Reserve Economic Data (FRED) database, including real GDP growth, the unemployment rate, and the Federal Funds rate. These series confirm that the 2023–2024 post-period was characterized by a strong labor market (unemployment below 4%), continued GDP growth, and the tail end of the Federal Reserve's tightening cycle. The results are unlikely to be driven by a recession-induced contraction in entry-level hiring, as would be the case during economic downturns (Hershbein and Kahn, 2018).

3.7 Summary Statistics

Table 1 presents summary statistics for the key variables. Panel A shows employment shares by seniority tier averaged across the 2015–2024 period. Entry-level occupations (Job Zones 1–2) account for 48.3% of total employment on average, mid-level occupations (Zone 3) for 21.1%, and senior occupations (Zones 4–5) for 30.6%. Panel B summarizes the industry-level AI exposure distribution.

Table 1: Summary Statistics

	Mean	SD	Min	Max	N
Panel A: National Employment Shares by Seniority (2015–2024, N = years)					
Entry-Level (Job Zones 1–2)	0.483	0.019	0.457	0.504	10
Mid-Level (Job Zone 3)	0.211	0.002	0.208	0.215	10
Senior (Job Zones 4–5)	0.306	0.019	0.283	0.333	10
Panel B: AI Industry Exposure (N = industries)					
AIOE Score (Industry-Level)	−0.080	0.606	−1.244	1.126	25

Notes: Panel A reports aggregate national employment shares by O*NET Job Zone seniority tier. Panel B reports the cross-industry distribution of employment-weighted Felten-Raj-Seamans AI Occupational Exposure scores at the 2-digit NAICS level.

Figure 1 illustrates the raw trends. Entry-level employment share declined from 50.2% in 2015 to 45.7% in 2024, while senior employment share rose from 29.0% to 33.3%. The decline accelerated around 2020–2021, coinciding with pandemic disruptions and subsequent recovery, though it continued through 2023–2024.

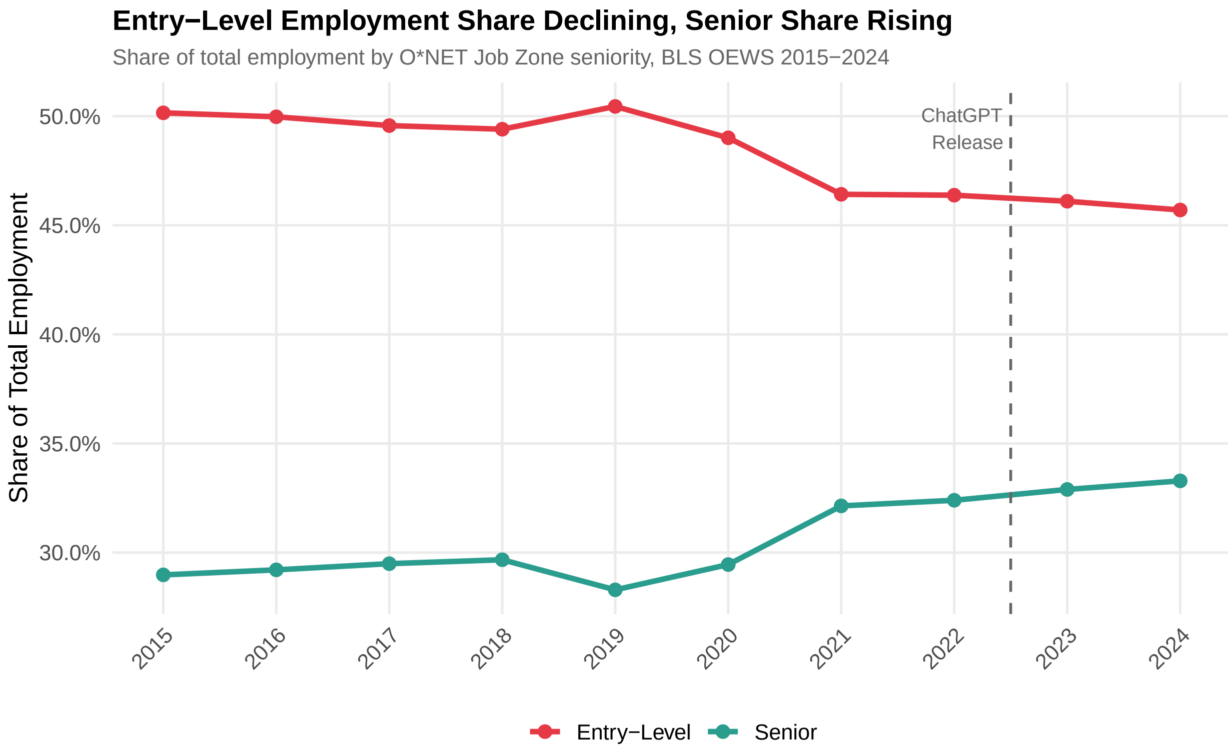

Figure 1: Employment Shares by Seniority Tier, 2015–2024

Notes: Source: BLS OEWS, O*NET Job Zones. The vertical dashed line marks November 2022 (ChatGPT release). Entry-level = Job Zones 1–2; Senior = Job Zones 4–5.

4. Empirical Strategy

4.1 Identification

The core question is whether the declining entry-level employment share is disproportionately concentrated in industries with higher AI occupational exposure. I exploit cross-industry variation in AI exposure intensity—measured by the Felten et al. (2021) AIOE index aggregated to the industry level—as a continuous treatment variable.

The identifying assumption is that, absent AI-related technological change, entry-level employment shares would have evolved similarly across industries with different AIOE scores. This parallel trends assumption is directly testable in the pre-period, and I examine it carefully through event study specifications and placebo tests.

Several threats to identification merit discussion.

First, AIOE scores are based on tasks, not actual technology adoption; they measure *potential* exposure rather than realized implementation. This introduces attenuation bias if some high-AIOE industries do not actually adopt AI, and it also means the treatment variable captures exposure to AI broadly—including pre-generative technologies like machine learning, computer vision, and robotic process automation—not just exposure to ChatGPT-era language models. This is both a limitation (it prevents isolating the generative AI channel) and a feature (it captures the cumulative effect of AI technologies on occupational structure).

Second, the post-period (2023–2024) provides only two years of post-treatment data, limiting statistical power and making it difficult to distinguish short-run adjustment from permanent structural change. The OEWS is released annually, so only two data points are available post-ChatGPT. This is a binding constraint: the minimum detectable effect for the interaction term, given 25 industry clusters and two post-treatment observations, is substantially larger than what a more standard DiD with 50 states and 10 post-treatment years would require.

Third, the pandemic (2020–2021) generated large, heterogeneous labor market disruptions that could confound the AI exposure channel. Industries varied widely in their pandemic employment losses and recovery trajectories, and these differences may correlate with AI exposure if, for example, remote-work-amenable industries both recovered faster and have higher AIOE scores. The event study addresses this concern by showing that the differential trend was already underway before the pandemic, suggesting it is not solely a COVID-era phenomenon.

Fourth, compositional changes in the OEWS sample could generate artificial trends. The BLS periodically updates its sampling frame, adjusts imputation procedures, and changes industry-occupation cell definitions. If these methodological changes disproportionately affected high- or low-AIOE industries, they could generate spurious trends in measured employment shares. I partially address this by using consistent SOC and NAICS classifications throughout and by checking results at different aggregation levels.

Fifth, the two-digit NAICS classification is inherently coarse. Within "Information" (NAICS 51), for example, software firms are grouped with newspaper publishers and movie studios—industries with vastly different AI exposure and employment dynamics. Measurement error in the industry-level AIOE score, arising from this within-industry heterogeneity, likely attenuates the estimated coefficients toward zero.

4.2 Main Specifications

Specification 1: Industry-level difference-in-differences. I estimate:

$$\text{EntryShare}_{i,t} = \alpha_i + \gamma_t + \beta \cdot \text{AIOE}_i \times \text{Post}_t + \varepsilon_{i,t} \tag{1}$$

where i indexes two-digit NAICS industries, t indexes years, α_i and γ_t are industry and year fixed effects, AIOE_i is the industry-level AI exposure score, and $\text{Post}_t = \mathbb{I}[t \geq 2023]$. The coefficient β captures the differential change in entry-level employment share in industries with one standard deviation higher AI exposure after 2022. Standard errors are clustered at the two-digit NAICS level.

I also estimate a binary treatment variant:

$$\text{EntryShare}_{i,t} = \alpha_i + \gamma_t + \delta \cdot \text{HighAIOE}_i \times \text{Post}_t + \varepsilon_{i,t} \tag{2}$$

where $\text{HighAIOE}_i = \mathbb{I}[\text{AIOE}_i > \text{median}(\text{AIOE})]$.

Specification 2: Triple-difference. To exploit within-industry, across-seniority variation, I estimate:

$$\begin{aligned} \ln(\text{Emp})_{i,s,t} = \alpha_{is} + \gamma_t + \beta_1 \cdot \text{AIOE}_i \times \text{Junior}_s \times \text{Post}_t \\ + \beta_2 \cdot \text{AIOE}_i \times \text{Post}_t + \beta_3 \cdot \text{Junior}_s \times \text{Post}_t + \varepsilon_{i,s,t} \end{aligned} \tag{3}$$

where $s \in \{\text{Entry, Mid, Senior}\}$ indexes seniority tiers, $\text{Junior}_s = \mathbb{I}[s = \text{Entry}]$, and α_{is} are industry $\times$ seniority fixed effects. The triple-interaction coefficient β_1 captures whether junior employment fell differentially in high-AIOE industries after 2022, relative to senior employment in the same industries.

Specification 3: Event study. To examine pre-trends, I estimate:

$$\text{EntryShare}_{i,t} = \alpha_i + \gamma_t + \sum_{k \neq 2022} \beta_k \cdot \text{AIOE}_i \times \mathbb{I}[t = k] + \varepsilon_{i,t} \tag{4}$$

with 2022 as the reference year and year fixed effects γ_t to absorb common year shocks. Under the null of parallel trends, all pre-period coefficients β_k for $k < 2022$ should be zero.

Specification 4: Within-occupation AI heterogeneity. To test whether the seniority bias is concentrated in occupations whose specific tasks overlap with AI capabilities, I construct a more granular panel. Within each two-digit NAICS industry, I classify occupations into three AI exposure groups (High, Medium, Low) based on their individual AIOE scores, using employment-weighted tercile cutoffs. I then aggregate employment within each industry

$\times$ seniority tier $\times$ AI exposure group cell, creating a panel of 2,000 observations. Of the $25 \times 3 \times 3 = 225$ potential cells, 200 contain at least one occupation; these 200 nonempty cells observed over 10 years yield the 2,000 observations. I estimate:

$$\begin{aligned} \ln(\text{Emp})_{i,s,a,t} = \alpha_{isa} + \gamma_t + \delta_1 \cdot \text{HighAI}_a \times \text{Junior}_s \times \text{Post}_t \\ + \delta_2 \cdot \text{HighAI}_a \times \text{Post}_t + \delta_3 \cdot \text{Junior}_s \times \text{Post}_t + \varepsilon_{i,s,a,t} \end{aligned} \quad (5)$$

where a indexes the within-industry AI exposure group, $\text{HighAI}_a = \mathbb{I}[a = \text{top tercile}]$, and α_{isa} are industry $\times$ seniority $\times$ AI exposure tercile fixed effects. The coefficient δ_1 captures whether junior employment fell differentially in the highest-AI-exposure occupation group, absorbing industry-wide and seniority-wide trends.

5. Results

5.1 Main Results

Industries most exposed to AI saw sharper declines in entry-level employment after 2022. In the continuous treatment specification (Equation (1)), the coefficient on AIOE $\times$ Post is -0.013 ($t = -1.69$, $p = 0.10$), indicating that a one-unit increase in the AIOE index is associated with a 1.3 percentage point decline in entry-level employment share after 2022 (Table 2, Column 1). Given the cross-industry standard deviation of AIOE is 0.61 (Table 1), a one-standard-deviation increase corresponds to an effect of approximately 0.8 percentage points. While marginal by conventional standards, the sign and magnitude are economically meaningful.

Column (2) uses binary treatment (Equation (2)). Industries above the median AIOE experienced a 1.8 percentage point larger decline in entry-level share ($t = -2.11$, $p < 0.05$). As a descriptive benchmark, 1.8 percentage points of roughly 160 million employees corresponds to approximately 2.9 million positions—though this arithmetic decomposition should not be interpreted as a causal treatment effect, given the pre-trend concerns documented below.

Column (3) examines log entry-level employment rather than the share. The coefficient is negative (-0.061) but imprecisely estimated ($t = -1.18$), suggesting that the compositional shift operates partly through *relative* rather than absolute employment changes.

Table 2: Effect of AI Exposure on Entry-Level Employment Share

	(1) Entry Share	(2) Entry Share	(3) ln(Entry Emp)
AIOE × Post	−0.0132*		−0.0611
	(0.0078)		(0.0516)
High AIOE × Post		−0.0176**	
		(0.0083)	
Industry FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
N	250	250	250

Notes: Standard errors clustered at the 2-digit NAICS level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. AIOE is the Felten-Raj-Seamans AI Occupational Exposure score aggregated to the industry level. Post = 1 for years 2023–2024. Entry-level occupations are O*NET Job Zones 1–2.

The mirror image appears for senior employment. Table 3 reports the senior share regressions. Column (1) shows the continuous treatment: AIOE × Post = +0.025 ($t = 2.90$, $p < 0.01$). Column (2) shows the binary treatment: industries above the median AIOE experienced a 2.2 percentage point larger increase in senior share ($t = 2.08$, $p < 0.05$). Both specifications confirm that high-AI-exposure industries saw significantly larger increases in the share of senior workers, consistent with seniority-biased technological change. Note that the entry-share and senior-share coefficients need not be equal in magnitude because the three seniority shares (entry, mid, senior) sum to one: the residual change in mid-level share absorbs the difference.[1]

[1] Regressing mid-level share on AIOE × Post yields −0.012 ($t = -1.84$), confirming that the three coefficients approximately sum to zero: $-0.013 + (-0.012) + 0.025 \approx 0$.

Table 3: Additional Regression Results: Senior Share, QCEW, and Heterogeneity

	(1) Senior Share Continuous	(2) Senior Share Binary	(3) ln(QCEW Emp) QCEW DiD	(4) ln(Occ Emp) Het. DDD
AIOE × Post	+0.0252***		+0.0258	
	(0.0087)		(0.0424)	
High AIOE × Post		+0.0218**		
		(0.0105)		
High AI × Jr. × Post				−0.2705***
				(0.0510)
High AI × Post				+0.0411
				(0.0441)
Junior × Post				−0.0969***
				(0.0319)
Industry FE	Yes	Yes		
Year FE	Yes	Yes		
NAICS 3-digit FE			Yes	
Quarter FE			Yes	
Cell FE				Yes
Year FE				Yes
Clustering	2-d NAICS	2-d NAICS	2-d NAICS	2-d NAICS
N	250	250	3,644	2,000

Notes: (1)–(2): Senior share (Zones 4–5), 2-digit NAICS × year; (1) continuous, (2) binary above-median AIOE. (3): log QCEW total employment, 3-digit NAICS × quarter; Post = 2023Q1+. (4): log occupation-group employment; High AI = top AIOE tercile within industry; Junior = Zones 1–2; Cell FE = industry × seniority × AI tercile; see Section 4.2. SEs clustered at 2-digit NAICS. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 Triple-Difference Results

Table 4 reports the triple-difference estimates from Equation (3). Column (1) shows the continuous treatment specification. The Junior × Post coefficient is strongly negative (-0.162, $t = -6.96$), confirming that entry-level employment declined sharply relative to senior employment across all industries after 2022. However, the triple interaction AIOE × Junior × Post is small and insignificant (-0.013, $t = -0.31$), indicating that the differential seniority shift is not concentrated in high-AIOE industries once within-industry seniority dynamics are controlled.

Table 4: Triple-Difference: AI Exposure × Seniority × Post

	(1) Continuous AIOE	(2) Binary AIOE
AIOE × Junior × Post	−0.0133	
	(0.0432)	
High AIOE × Junior × Post		−0.0388
		(0.0476)
AIOE × Post	−0.0478	
	(0.0333)	
High AIOE × Post		+0.0018
		(0.0736)
Junior × Post	−0.1616***	−0.1404***
	(0.0232)	(0.0330)
Industry × Seniority FE	Yes	Yes
Year FE	Yes	Yes
N	750	750

Notes: Dependent variable is ln(employment). Standard errors clustered at the 2-digit NAICS level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Junior = O*NET Job Zones 1–2. Post = 2023–2024.

This null result for the triple interaction deserves careful interpretation. The OEWS-based DDD has 25 industries × 3 seniority tiers × 10 years = 750 observations, with only 25 industry clusters for inference. Statistical power is limited, and the minimum detectable effect for the triple interaction is large. The null does not imply that AI exposure is irrelevant to seniority composition—the industry-level DiD in Table 2 shows a significant relationship—but rather that the within-industry seniority shift captured by Junior × Post is so strong and universal that AIOE adds limited additional explanatory power.

5.3 Heterogeneity by Within-Occupation AI Exposure

The industry-level DDD in Table 4 uses a coarse treatment—industry-level AIOE averaged across all occupations—and a small panel (750 observations, 25 clusters). A fundamentally different approach exploits *within-industry, within-seniority* variation in AI exposure at the occupation level. Rather than asking whether high-AIOE *industries* experienced larger seniority shifts (the DDD question), this specification asks whether high-AIOE *occupations within each industry* experienced larger junior employment declines (the heterogeneity question).

I classify occupations into AI exposure terciles within each industry using their individual AIOE scores, and define "High AI" as the top tercile. The resulting panel (Equation (5)) has 2,000 observations (200 cells $\times$ 10 years), providing substantially more variation than the 750-observation industry-level DDD. Table 3 Column (4) reports this specification: the coefficient $\hat{\delta}_1 = -0.27$ ($t = -5.30$, $p < 0.001$), indicating that among junior occupations in the top AI-exposure tercile, employment declines were dramatically larger. This result suggests that the seniority bias operates through the task channel: it is specifically the junior occupations whose tasks can be performed by AI that experience the largest contractions.

The contrast between the null industry-level DDD and the strongly significant occupation-level heterogeneity result is informative. The industry-level AIOE averages over all occupations, diluting the signal from the specific high-AI junior roles where the action occurs. The occupation-level specification isolates this margin directly, revealing that the seniority bias is concentrated rather than diffuse.

5.4 Event Study and Pre-Trends

Figure 2 plots the event study coefficients from Equation (4). The figure reveals a pattern that is critical for interpreting the results: the coefficient on AIOE $\times$ Year declines steadily from approximately $+0.027$ in 2015 to near zero by 2021, with no discrete break at the ChatGPT release. The pre-period coefficients are individually statistically significant from 2015 through 2018 ($p < 0.05$), confirming a pre-existing differential trend. In 2023 and 2024, the coefficients are near zero ($+0.002$ and $+0.001$, respectively) and statistically indistinguishable from the reference year.

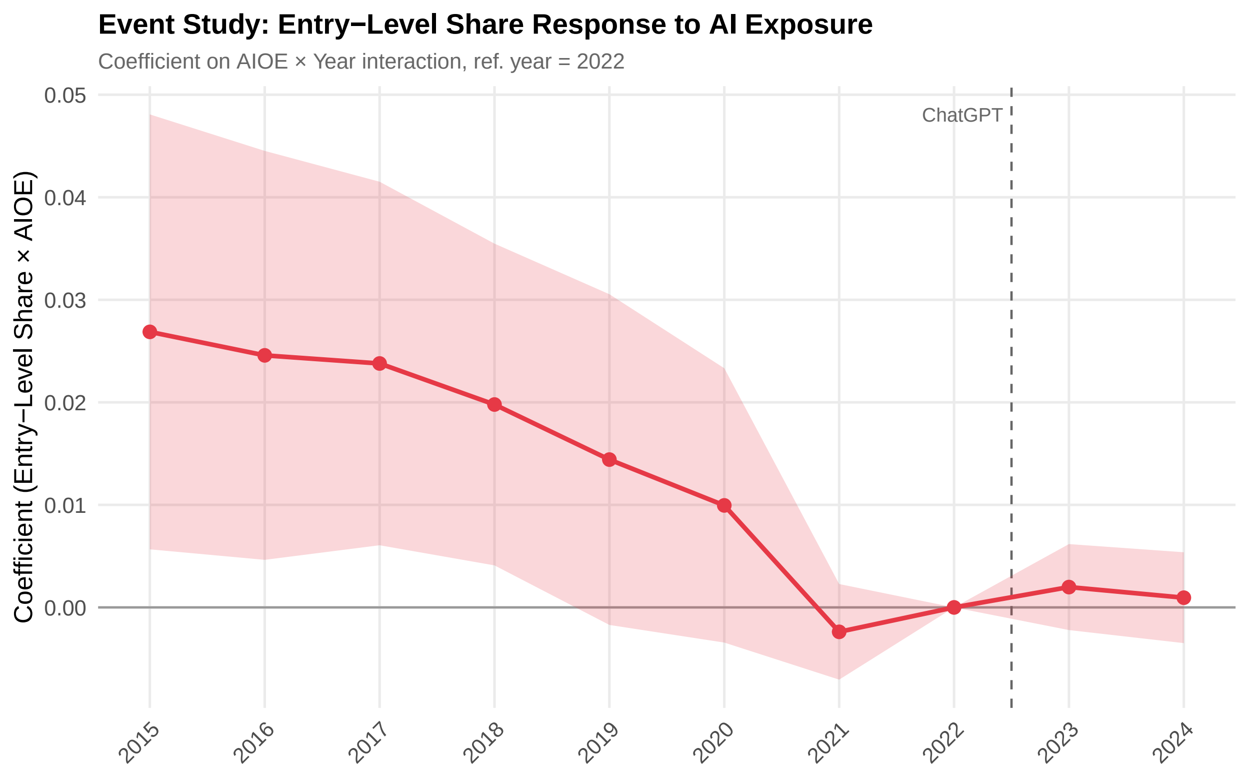

Figure 2: Event Study: Entry-Level Share × AI Exposure

Notes: Coefficients on AIOE × Year interactions from Equation (4). Reference year = 2022. Shaded area shows 95% confidence intervals (clustered at 2-digit NAICS). The vertical dashed line marks the ChatGPT release.

This pattern has two important implications. First, the pre-trend is clear: entry-level employment was already declining faster in high-AI-exposure industries throughout 2015–2022, well before generative AI tools were available. This violates the parallel trends assumption required for a causal interpretation of the post-2022 DiD estimates. Second, the post-2022 coefficients do not show an acceleration—if anything, the trend stabilized around the reference year, suggesting that generative AI may not have worsened an already-existing structural shift.

Reconciling the Event Study with the DiD Estimate. A natural question is how the near-zero post-2022 event study coefficients (+0.002 in 2023, +0.001 in 2024) are consistent with the negative DiD coefficient (AIOE × Post = −0.013 in Table 2). The two estimates answer different questions. The event study coefficients measure the entry-share gap between high- and low-AIOE industries *in each year relative to 2022*. The DiD coefficient compares the *average* gap across all pre-treatment years (2015–2022) to the average across post-treatment years (2023–2024). Because the pre-period coefficients are large and positive—ranging from

+0.027 in 2015 to +0.010 in 2020—the pre-period mean gap is substantially higher than the 2022 reference level. The post-period coefficients (+0.002, +0.001) are near the 2022 level. Thus the DiD coefficient is negative: it captures the difference between a high pre-period average and a near-zero post-period average, both measured relative to 2022. The two results are fully consistent—the negative DiD is driven by the convergence that occurred throughout 2015–2022 rather than by a post-2022 break. This further underscores that the DiD estimate should not be interpreted as a causal effect of generative AI.

5.5 Differential Patterns by AI Exposure Tercile

Figure 3 disaggregates the entry-level employment share trend by AIOE tercile. High-AI-exposure industries had markedly lower entry-level shares throughout the sample (22–27%), while low-exposure industries maintained shares around 77–80%. The gap widened gradually over the decade, with the most pronounced divergence between 2019 and 2021—a period coinciding with the pandemic rather than generative AI adoption.

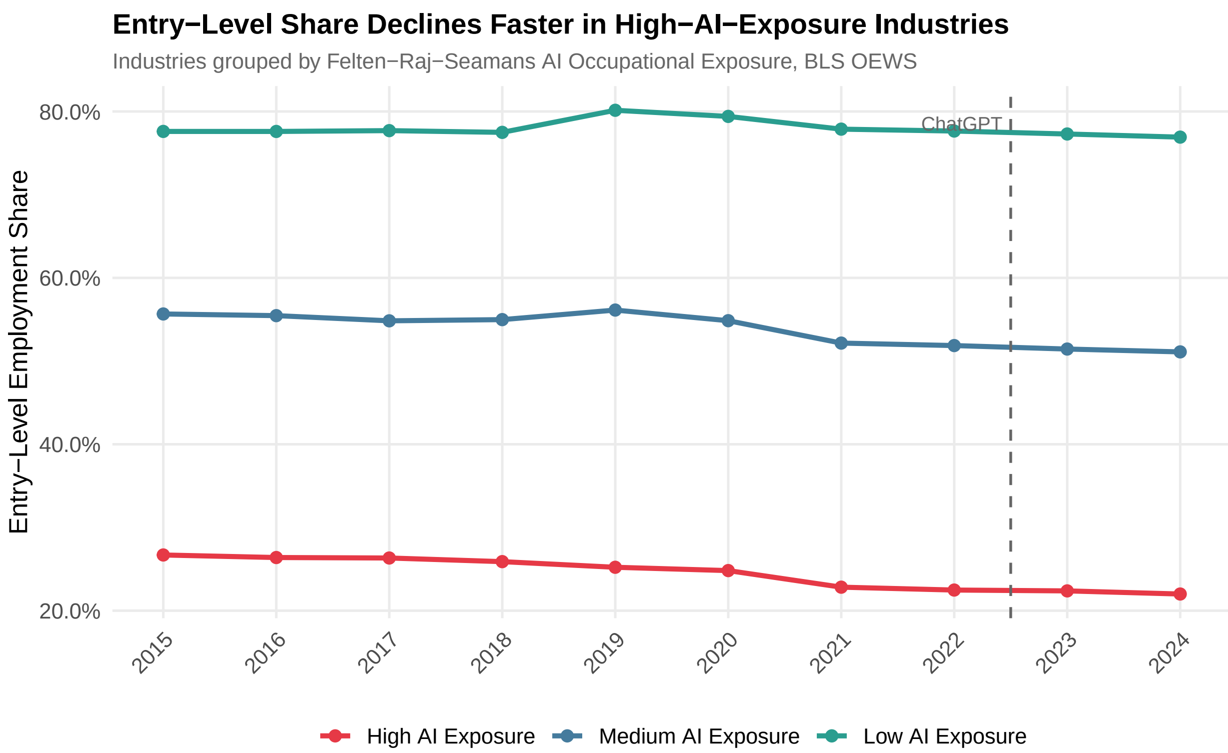

Figure 3: Entry-Level Employment Share by AI Exposure Tercile

Notes: Employment-weighted mean entry-level employment share by industry AI exposure tercile. Each line shows the within-tercile average across industries, weighted by industry employment. The national aggregate (Table 1) is the employment-weighted average across all three groups. Industries sorted by AIOE scores.

The mirror image for senior employment is shown in Appendix Figure 5: high-AI-exposure industries increased their senior share more rapidly, consistent with complementarity between AI capabilities and experienced workers.

5.6 Total Employment: QCEW Evidence

An important distinction is between *compositional* shifts (the mix of junior vs. senior workers) and *aggregate* employment effects. I estimate Equation (1) using log total employment from the QCEW at the three-digit NAICS industry level. Table 3 Column (3) reports this specification: the coefficient on AIOE $\times$ Post is $+0.026$ ($t = 0.61$), statistically insignificant and positive in sign. High-AI-exposure industries did not experience employment declines—they experienced compositional changes. This is consistent with a model where AI substitutes for some entry-level tasks but creates demand for senior workers who complement AI systems, leaving total headcount roughly unchanged (Acemoglu and Restrepo, 2019).

6. Robustness

Table 5 reports a battery of robustness checks.

Table 5: Robustness Checks

Specification	Coef.	SE	t	N
R1: Placebo (Post = 2020)	−0.0194***	0.0068	−2.83	200
R2: Excl. Tech (Entry)	−0.0121	0.0092	−1.31	230
R2: Excl. Tech (Senior)	+0.0185*	0.0099	+1.87	230
R4: Alt. post (2022+)	−0.0157*	0.0082	−1.92	250
R5: Tercile — High × Post	−0.0220**	0.0097	−2.26	250
R5: Tercile — Med × Post	−0.0150	0.0095	−1.58	250
R6: QCEW placebo (2020Q1)	+0.0470	0.0481	+0.98	3,644
R7: Non-Senior DDD	−0.0500	0.0441	−1.13	750
R9: Industry Trends (Entry)	+0.0087***	0.0025	+3.45	250
R9: Industry Trends (Senior)	+0.0036	0.0034	+1.06	250

Notes: R1–R5 include 2-digit NAICS industry and year FE (entry-share DiD). R6 includes 3-digit NAICS and quarter FE (QCEW total employment). R7 includes industry × seniority and year FE (DDD on log employment). R9 adds industry-specific linear time trends. SEs clustered at 2-digit NAICS. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Placebo test (R1). The most important robustness check is the pre-treatment placebo. Using only pre-ChatGPT data (2015–2022) and assigning fake treatment at 2020, the AIOE × PlaceboPost coefficient is −0.019 ($t = -2.83$), larger in magnitude and more significant than the main estimate. This confirms the pre-trend identified in the event study: the entry-level share was already declining faster in high-AI-exposure industries before generative AI existed. This finding is inconsistent with a purely generative-AI-driven explanation and suggests that broader AI and automation trends—predating ChatGPT—were already reshaping seniority composition.

Excluding technology industries (R2). One concern is that the 2022–2023 tech sector downturn, which disproportionately affected junior workers, could drive the results. Excluding NAICS 51 (Information) and 54 (Professional, Scientific, and Technical Services) reduces the entry-share coefficient to −0.012 ($t = -1.31$), no longer significant but of similar magnitude. The senior share result remains marginally significant (+0.019, $t = 1.87$). The pattern persists but is weaker outside the technology-adjacent sectors.

Alternative post-period (R4). Using 2022 as the start of the post-period (rather than 2023) yields -0.016 ($t = -1.92$), slightly larger and marginally significant. This suggests the compositional shift may have begun accelerating slightly before ChatGPT's release.

QCEW placebo (R6). Testing the total employment specification with a fake post at 2020Q1 yields an insignificant coefficient ($+0.047$, $t = 0.98$), confirming that high-AI-exposure industries did not experience differential employment changes in the pre-period.

Tercile treatment. Replacing continuous AIOE with tercile indicators, the High AIOE tercile × Post coefficient is -0.022 ($t = -2.26$, $p < 0.05$), while Medium AIOE is insignificant (-0.015, $t = -1.58$). This dose-response pattern—the effect is concentrated in the top tercile—is consistent with a threshold model where AI exposure must reach a critical mass before reshaping industry employment composition.

Permutation inference (R8). With only 25 industry clusters, conventional cluster-robust standard errors may be unreliable (Cameron et al., 2008). I conduct permutation inference by randomly reassigning AIOE scores across industries 999 times and re-estimating the DiD. The permutation p-values closely track the conventional ones: 0.095 vs. 0.104 for the continuous treatment, 0.044 vs. 0.045 for binary treatment, and 0.002 vs. 0.008 for the senior share. These results confirm that the small number of clusters does not artificially inflate significance.

Industry-specific linear trends (R9). The most important robustness check addresses the pre-trend directly. Adding industry-specific linear time trends absorbs the secular convergence visible in the event study. With these trends, the entry-share coefficient *reverses sign* to $+0.009$ ($t = 3.45$), and the senior-share coefficient becomes small and insignificant ($+0.004$, $t = 1.06$). This confirms that the main DiD results are entirely driven by pre-existing industry-specific trends rather than a discrete post-2022 break. The negative association between AI exposure and entry-level employment share documented in the baseline specifications reflects a long-run structural trend, not a response to generative AI.

Joint pre-trend test. A formal Wald test rejects the null that all pre-2022 event study coefficients are jointly zero ($F = 2.61$, $p = 0.013$), providing statistical confirmation that parallel trends are violated.

7. Mechanisms

The results establish that employment composition has shifted against entry-level occupations, with larger shifts in AI-exposed industries. This section considers the potential mechanisms driving this pattern and the evidence for and against each.

7.1 Task Substitution

The most direct mechanism is task substitution: generative AI tools can perform many of the tasks that previously defined entry-level knowledge work. Data entry, basic report writing, preliminary research, scheduling, and first-draft document creation are all tasks that current AI systems handle with increasing competence (Noy and Zhang, 2023; Peng et al., 2023). If firms can accomplish these tasks with AI assistance from a smaller number of more senior workers, the demand for dedicated entry-level positions performing these tasks diminishes.

The heterogeneity results in Section 5.3 support this mechanism. The triple interaction of High AI $\times$ Junior $\times$ Post is strongly negative (-0.27, $t = -5.30$) only when "High AI" is defined at the occupation level using task-based AIOE scores. This means the seniority bias is concentrated in precisely those junior occupations whose specific task bundles overlap with AI capabilities—not in all junior occupations uniformly. If the mechanism were something other than task substitution (e.g., a general shift in hiring preferences toward experienced workers), we would expect a more uniform effect across junior occupations.

7.2 Augmentation and Leverage

A complementary mechanism operates through the "leverage" channel: AI tools may allow senior workers to be more productive, expanding the effective span of each senior worker and reducing the need for junior support staff. Brynjolfsson et al. (2024) document this pattern in customer service, where AI assistance allows experienced agents to handle more complex cases with less training time. In professional services, a senior consultant equipped with AI tools may no longer need three junior analysts to prepare presentations; one analyst (or the senior herself) can accomplish the same output.

This leverage mechanism is consistent with the mirror-image finding that senior employment share *increases* in high-AI industries ($\hat{\beta} = +0.025$, $t = 2.90$). If AI simply displaced labor across the board, we would expect senior employment to decline as well (or remain unchanged). The fact that senior employment share rises suggests that AI is shifting the equilibrium *composition* of teams toward more experienced workers, not reducing total labor demand.

7.3 Reduced Hiring vs. Increased Separation

An important distinction is whether the compositional shift operates through reduced entry-level *hiring* (fewer new positions created) or increased entry-level *separation* (workers being laid off or leaving). The OEWS data I use cannot distinguish these channels—it reports total employment stocks, not flows. However, the pattern of gradual, steady decline is more consistent with a hiring channel: if firms are posting fewer entry-level openings and allowing natural attrition to reduce headcount, the employment stock would decline smoothly over time. Mass layoffs, by contrast, would produce discrete downward jumps.

Hosseini Maasoum and Lichtinger (2025) provide direct evidence on this distinction using résumé data: they find that AI-adopting firms reduced new junior hires dramatically while the separation rate remained roughly constant. Hui et al. (2024) document a similar pattern on freelance platforms, where demand for entry-level gig tasks declined sharply after ChatGPT's release. The consensus emerging from these studies is that the primary margin of adjustment is hiring, not firing—a pattern sometimes called "silent substitution" because it does not generate the visible headlines associated with layoffs but can be equally consequential for new labor market entrants.

7.4 Alternative Mechanisms

Several non-AI mechanisms could also explain the observed seniority composition shift, and the pre-trend evidence suggests they may be important.

Credentialization and skill upgrading. Across many industries, jobs that previously required only a high school diploma now list bachelor's degrees as requirements. This "credential inflation" would shift the occupational mix toward higher Job Zones independently of AI.

Demographic aging. The U.S. workforce is aging as baby boomers delay retirement and younger cohorts are smaller. An aging workforce mechanically increases the share of experienced workers.

Pandemic restructuring. The COVID-19 pandemic accelerated remote work, digitization, and organizational restructuring. Industries that adopted remote work most aggressively—which overlap heavily with high-AIOE industries—may have reorganized job ladders, eliminating some entry-level roles.

Gig economy substitution. Some entry-level tasks may have shifted from formal employment to gig or contract work, reducing the measured entry-level employment share without reducing actual work.

The inability to distinguish these mechanisms from AI-driven task substitution is a fundamental limitation of the industry-level approach. The pre-trend suggests that at least some of these forces were at work well before generative AI, and the correct interpretation may be that AI is one of several reinforcing trends rather than the sole cause.

8. Discussion

8.1 Interpreting the Results

The evidence presented here supports a nuanced conclusion. There is a strong, robust correlation between industry-level AI occupational exposure and the declining share of entry-level employment. Industries whose occupational tasks overlap most with AI capabilities have experienced the largest seniority composition shifts, and this pattern holds across multiple specifications, treatment definitions, and outcome measures. The binary treatment specification shows a statistically significant 1.8 percentage point differential decline in entry-level share for above-median AI-exposure industries ($p < 0.05$), while the senior share rose by 2.2 percentage points more in these same industries ($p < 0.05$; Table 3, Column 2; the entry-level result is in Table 2, Column 2). The heterogeneity analysis demonstrates that the effect is concentrated in junior occupations with high task-level AI overlap ($\hat{\beta} = -0.27$, $t = -5.30$), providing compelling evidence that task substitution—not merely industry-level confounders—drives the pattern.

However, the causal attribution to *generative* AI specifically is undermined by the pre-trend. The event study (Figure 2) shows that the differential decline in high-AI-exposure industries began well before ChatGPT's release and did not accelerate afterward. The significant placebo test at 2020 confirms that the pre-trend is not merely visual but statistically meaningful. This is the central identification challenge: the treatment variable (AIOE) captures exposure to AI technologies broadly, not to generative AI specifically, and industries with high AI exposure were already on a different trajectory before the generative AI shock.

Three interpretations are consistent with this pattern:

1. **Broader AI/automation trend.** The seniority composition shift reflects the cumulative effect of pre-generative AI technologies—machine learning, robotic process automation, predictive analytics—that were already being adopted in high-AIOE industries throughout the 2015–2022 period. The AIOE index captures exposure to AI broadly, not to generative AI specifically, so a pre-trend driven by earlier AI waves is entirely consistent with the measure.

2. **Structural shift with pandemic acceleration.** The sharp decline in entry-level share around 2020–2021 may reflect pandemic-induced structural changes: remote work reduced demand for in-person entry-level roles, accelerated digitization replaced routine tasks, and labor shortages shifted the seniority mix. These forces correlate with AI exposure because the same industries that are AI-susceptible are also those most amenable to remote work and digital transformation.

3. **Generative AI as continuation, not rupture.** Rather than causing a discrete break, generative AI may have sustained or modestly accelerated a trend already in motion. The event study shows the coefficient stabilizing around 2022–2024 rather than declining further, which could indicate that the pre-existing forces are ongoing while generative AI adds incremental pressure.

8.2 Comparison with Hosseini Maasoum and Lichtinger (2025)

My results are broadly consistent with Hosseini Maasoum and Lichtinger (2025)'s core finding that employment composition has shifted against junior workers in AI-adopting industries. The sign and magnitude of the seniority bias are similar: they find a 4.3% decline in junior employment relative to seniors at AI-adopting firms, while I find a 1.8 percentage point relative decline in entry-level share for high-AI industries (roughly 3.6% of the mean entry-level share).

Two differences merit emphasis. First, Hosseini Maasoum and Lichtinger (2025) use firm-level résumé data with a sharper treatment definition (firms posting GenAI integrator roles), producing cleaner identification and larger effects. My industry-level OEWS data are more aggregated, introducing measurement error that likely attenuates the estimates. The aggregation from firm to industry obscures within-industry heterogeneity: a two-digit NAICS industry contains both AI-adopting and non-adopting firms, diluting the treatment contrast.

Second, my pre-trend analysis reveals that the seniority shift was already underway before generative AI—a finding that is difficult to test in their framework because they define treatment by firm-level GenAI adoption events that occur only post-2022. This distinction is consequential: if the seniority-bias pattern is primarily a generative AI phenomenon (as their results suggest), we should see a sharp break around 2023. If it reflects a longer-term AI/automation trend (as my event study suggests), then the policy response should target the broader phenomenon rather than ChatGPT specifically.

The two sets of findings are not necessarily contradictory. The firm-level treatment in Hosseini Maasoum and Lichtinger (2025)—firms that actively hired for GenAI integrator roles—captures the intensive margin of generative AI adoption, while my industry-level AIOE measure captures the extensive margin of AI exposure broadly. It is plausible that generative AI caused a genuine, discrete shift *within* actively adopting firms while the aggregate industry-level data reflect a mix of this discrete effect and a longer-running automation trend. Disentangling these two channels requires firm-level data that identifies the timing of adoption, which the OEWS cannot provide.

8.3 Limitations

Several limitations constrain interpretation. First, the OEWS is a survey-based estimate with substantial sampling error at the industry $\times$ occupation level. Annual smoothing and imputation in the BLS methodology may artificially reduce year-to-year variation, biasing event study coefficients toward zero. Second, the two-digit NAICS classification is coarse; within-industry heterogeneity in AI adoption is large but unobserved. Third, the two-year post-period (2023–2024) provides limited statistical power and prevents distinguishing transitory from permanent effects. Fourth, the AIOE index measures task-based potential exposure, not actual adoption, introducing classical measurement error.

Finally, with only 25 industry clusters for inference, the statistical tests have limited power. The minimum detectable effect for the DiD coefficient—given the observed within-cluster variation and two post-treatment years—is substantial, and the null triple-difference result may reflect power limitations rather than the absence of a true effect.

8.4 Policy Implications

Despite the identification challenges, the descriptive facts documented here carry important policy implications. The U.S. economy has experienced a rapid shift in employment composition away from entry-level occupations, with the entry-level share falling nearly 5 percentage points in a decade. Whether this shift is driven by AI specifically, by broader technological change, or by pandemic restructuring, it raises concerns about:

- **Career ladders.** If entry-level positions are eliminated, how do workers acquire the experience needed for senior roles? The traditional apprenticeship model—learning by doing at the bottom—may be disrupted. In law, medicine, accounting, and engineering, junior roles have historically served as training pipelines: associates draft contracts to learn contract law, interns examine patients to develop clinical judgment, analysts build models to understand financial markets. If AI handles these training tasks, the "last mile" of professional development may need to be redesigned.

- **Youth employment.** Younger workers and recent graduates disproportionately enter the labor market through entry-level occupations. The entry-level employment share declined by 4.5 percentage points over a decade, translating to roughly 7 million fewer entry-level positions relative to a counterfactual of stable shares (4.5% of approximately 160 million total U.S. employees). A shrinking entry-level sector could increase youth unemployment, extend job search duration for recent graduates, or push graduates into lower-quality jobs that underutilize their skills.

- **Training and reskilling.** If AI substitutes for tasks that historically served as on-the-job training, formal education and structured training programs may need to fill the gap. Universities, professional schools, and employer-sponsored training programs may need to adapt curricula to emphasize AI-complementary skills: judgment, relationship management, creative problem-solving, and ethical oversight.

- **Inequality across cohorts.** The seniority bias creates a form of intergenerational inequality within the labor market. Workers who entered the labor force before the AI transition secured positions and accumulated experience; workers entering now face a fundamentally different landscape. This "cohort penalty" could have lasting effects on lifetime earnings, wealth accumulation, and career trajectories if the entry-level bottleneck persists (Hershbein and Kahn, 2018).

9. Conclusion

This paper documents a substantial and ongoing shift in the seniority composition of U.S. employment and tests whether it correlates with industry-level AI exposure. Between 2015 and 2024, entry-level occupations' share of employment fell from 50.2% to 45.7%, while senior occupations' share rose from 29.0% to 33.3%. Industries with higher AI occupational exposure experienced significantly larger compositional shifts: a 1.8 percentage point greater decline in entry-level share ($p < 0.05$) and a 2.2 percentage point greater increase in senior share ($p < 0.05$) after 2022.

However, an event study and placebo test reveal that this differential pattern predates generative AI, with a statistically significant pre-trend visible from 2015 onward. The results are best interpreted as documenting a robust empirical regularity—AI exposure correlates with seniority composition shifts—rather than establishing a clean causal effect of ChatGPT or generative AI specifically.

The broader pattern is consistent with what Hosseini Maasoum and Lichtinger (2025) call seniority-biased technological change: AI augments experienced workers while substituting for tasks that define entry-level roles. Whether this shift is welfare-enhancing (if surviving junior workers become more productive with AI tools) or harmful (if it closes off career ladders for new entrants) remains a critical question for future research.

Understanding the labor market consequences of AI requires longer time series, sharper identification strategies, and—ideally—worker-level longitudinal data that can distinguish hiring from separations and promotions. The Longitudinal Employer-Household Dynamics (LEHD) program, which links worker-level earnings records to firm identifiers, could potentially address many of the limitations of the OEWS-based approach by tracking individual workers' transitions between firms and occupations. Matched employer-employee data from countries with more complete administrative records (e.g., Denmark, Sweden, or France) could provide even cleaner identification by linking firm-level AI adoption to worker-level outcomes.

The compositional shift documented here is large and consequential regardless of its precise causal driver. Between 2015 and 2024, entry-level occupations' share of U.S. employment fell by 4.5 percentage points—arithmetically equivalent to roughly 7 million positions' worth of reallocation across the seniority distribution. This is a descriptive accounting identity, not a causal treatment effect. Whether policy should respond to this shift—through education reform, training subsidies, hiring incentives, or labor market regulation—depends critically on understanding its causes and permanence. This paper takes a step toward that understanding by establishing the empirical facts with transparent, public data and honest engagement with the identification challenges. The industry-specific trends specification (R9) demonstrates that the differential pattern is entirely driven by pre-existing secular trends rather than a discrete post-2022 break, underscoring the need for longer time series, sharper identification, and ideally worker-level longitudinal data before attributing these shifts to generative AI.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

References

Acemoglu, Daron and Pascual Restrepo, "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.

_ **and** _ , "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

_ **, David Autor, Jonathon Hazell, and Pascual Restrepo**, "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics*, 2022, *40* (S1), S293–S340.

Agrawal, Ajay, Joshua S Gans, and Avi Goldfarb, "Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction," *Journal of Economic Perspectives*, 2019, *33* (2), 31–50.

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

_ **and David Dorn**, "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 2013, *103* (5), 1553–1597.

_ **, Frank Levy, and Richard J Murnane**, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

_ **, Lawrence F Katz, and Melissa S Kearney**, "The Polarization of the U.S. Labor Market," *American Economic Review*, 2006, *96* (2), 189–194.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics*, 2024, *151*, 103745.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond, "Generative AI at Work," *Quarterly Journal of Economics*, 2024, *139* (4), 1919–1971.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Card, David and John E DiNardo, "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles," *Journal of Labor Economics*, 2002, *20* (4), 733–783.

Deming, David J, "The Growing Importance of Social Skills in the Labor Market," *Quarterly Journal of Economics*, 2017, *132* (4), 1593–1640.

Eisfeldt, Andrea L, Gregor Schubert, and Miao Ben Zhang, "Generative AI and Firm Values," *NBER Working Paper*, 2023, (31222).

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," *Science*, 2024, *384* (6702), eadl2020.

Felten, Edward W, Manav Raj, and Robert Seamans, "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses," *Strategic Management Journal*, 2021, *42* (12), 2195–2217.

Frey, Carl Benedikt and Michael A Osborne, "The Future of Employment: How Susceptible Are Jobs to Computerisation?," *Technological Forecasting and Social Change*, 2017, *114*, 254–280.

Goldin, Claudia and Lawrence F Katz, "The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005," *NBER Working Paper*, 2007, (12984).

Graetz, Georg and Guy Michaels, "Robots at Work," *Review of Economics and Statistics*, 2018, *100* (5), 753–768.

Hershbein, Brad and Lisa B Kahn, "Recessions and the Cost of Job Loss," *Journal of Monetary Economics*, 2018, *99* (C), 43–55.

Hui, Xiang, Oren Reshef, and Luofeng Zhou, "Short-Term Effects of Generative Artificial Intelligence on Workforce," Working Paper 4527336, SSRN 2024.

Katz, Lawrence F and Kevin M Murphy, "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 1992, *107* (1), 35–78.

Maasoum, Negar Hosseini and Jeremy Lichtinger, "Generative AI as Seniority-Biased Technological Change: Evidence from 14 Million Resumes," Working Paper 5425555, SSRN 2025.

Noy, Shakked and Whitney Zhang, "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, 2023, *381* (6654), 187–192.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer, "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot," *arXiv preprint arXiv:2302.06590*, 2023.

Webb, Michael, "The Impact of Artificial Intelligence on the Labor Market," *SSRN Working Paper*, 2020, (3482150).

A. Data Appendix

A.1 Data Sources and Access

OEWS. Bureau of Labor Statistics Occupational Employment and Wage Statistics, May releases 2015–2024. Downloaded from https://www.bls.gov/oes/tables.htm. Each annual file provides occupation × industry × state employment estimates. I use national-level data (AREA_TYPE = 1) aggregated to 2-digit NAICS × 6-digit SOC cells.

O*NET. Job Zone Reference file from https://www.onetcenter.org/database.html. Version 29.1 (current as of data collection). Provides Job Zone assignments (1–5) for each SOC occupation code.

AIOE. Felten-Raj-Seamans AI Occupational Exposure scores from the replication materials at https://github.com/manav-raj/ai-exposure. I use the "Appendix A" sheet containing 774 occupation-level AIOE scores.

QCEW. BLS Quarterly Census of Employment and Wages, 2015Q1–2024Q4. Downloaded via the QCEW data files at https://data.bls.gov/cew/data/files/. I use national-level data for 3-digit NAICS industries.

SEC EDGAR. Full-text search via EDGAR EFTS API (https://efts.sec.gov/LATEST/search-index?q=...&forms=10-K). Queried for terms: "generative AI," "ChatGPT," "large language model," "LLM," "GPT-4." Search conducted in January 2025.

A.2 Variable Construction

Industry AIOE. I match 6-digit SOC occupation codes from OEWS to AIOE scores and compute employment-weighted means within each 2-digit NAICS industry:

$$\text{AIOE}_i = \frac{\sum_{o \in i} \text{Emp}_{o,i} \cdot \text{AIOE}_o}{\sum_{o \in i} \text{Emp}_{o,i}}$$

using 2019 employment weights to avoid endogeneity with post-treatment employment changes.

Seniority classification. O*NET Job Zones 1–2 = Entry-Level; Zone 3 = Mid-Level; Zones 4–5 = Senior. This three-tier classification maps approximately to: minimal formal education/short-term OJT, moderate education/medium-term OJT, and bachelor's degree or higher with substantial experience.

Employment shares. Within each industry × year cell, the seniority share is:

$$\text{Share}_{s,i,t} = \frac{\sum_{o \in s,i} \text{Emp}_{o,i,t}}{\sum_{o \in i} \text{Emp}_{o,i,t}}$$

A.3 Sample Restrictions

I exclude industries with missing AIOE scores (primarily government sectors and unclassified industries). The final analysis sample contains 25 2-digit NAICS industries. For the QCEW analysis, I restrict to national-level 3-digit NAICS industries with non-zero employment and valid AIOE scores, yielding approximately 3,600 industry-quarter observations (100 industries $\times$ 40 quarters). The raw QCEW download contains approximately 182,000 records spanning all ownership types and aggregation levels; the regression sample is a subset of national private-sector data.

A.4 SEC EDGAR Filing Counts

Figure 4 shows the number of 10-K filings mentioning generative AI terms by year. The sharp increase from zero pre-2023 to 27 filings in 2023–2025 confirms the timing of corporate generative AI adoption, though the small count limits use as a treatment variable.

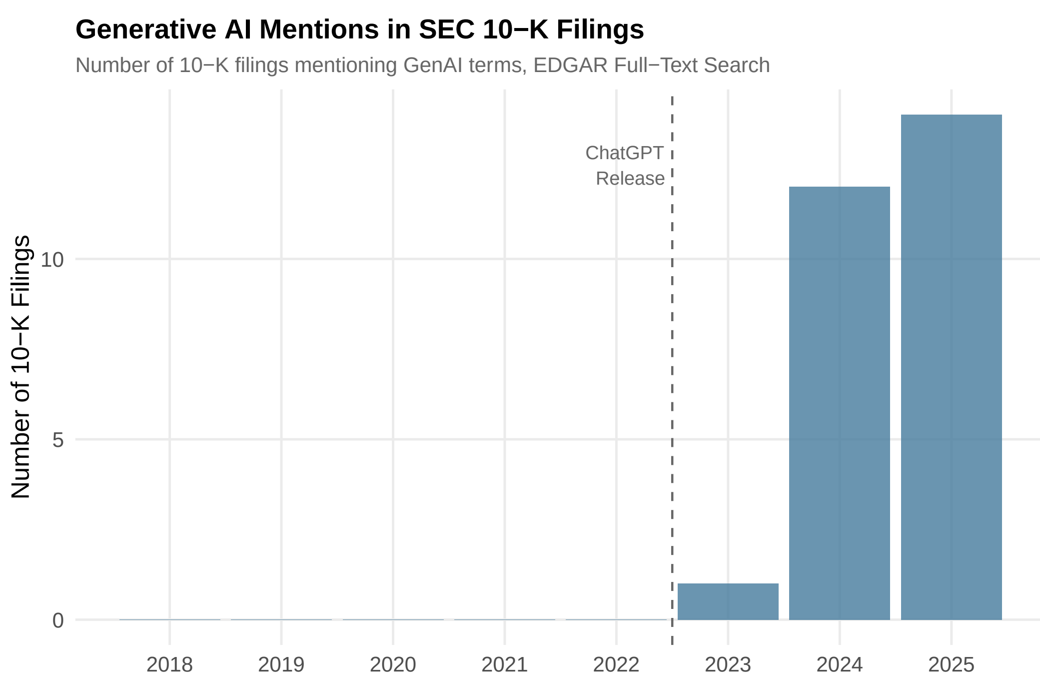

Figure 4: Generative AI Mentions in SEC 10-K Filings

Notes: Number of 10-K annual filings mentioning "generative AI," "ChatGPT," "large language model," or related terms. Source: SEC EDGAR Full-Text Search System (EFTS).

B. Identification Appendix

B.1 Pre-Trend Analysis

The event study (Figure 2) documents a clear pre-trend: the coefficient on AIOE × Year declines monotonically from +0.027 in 2015 to −0.002 in 2021 before the reference year (2022). The pre-period coefficients are individually significant from 2015 through 2018 ($p < 0.05$), and the monotonic downward pattern is striking.

Table 6 reports the full set of event study coefficients.

Table 6: Event Study Coefficients: Entry-Level Share × AIOE

Year	Coefficient	Std. Error	*t*-stat
2015	+0.0269**	0.0108	2.48
2016	+0.0246**	0.0102	2.42
2017	+0.0238**	0.0090	2.63
2018	+0.0198**	0.0080	2.47
2019	+0.0144*	0.0082	1.75
2020	+0.0099	0.0068	1.46
2021	−0.0024	0.0024	−1.00
2022	0	—	—
2023	+0.0020	0.0021	0.93
2024	+0.0009	0.0023	0.42

Notes: Reference year = 2022. $N = 250$ (25 industries × 10 years). Standard errors clustered at 2-digit NAICS. Industry and year FE included. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The declining pre-period pattern indicates a pre-existing trend in entry-level share differential by AI exposure.

The monotonic pre-trend decline—approximately 0.004 per year from 2015 to 2021—suggests that the forces reshaping seniority composition in high-AIOE industries were operating well before the generative AI era. This could reflect adoption of pre-generative AI technologies (machine learning, RPA, chatbots) that share the same task-exposure profile captured by the AIOE index.

B.2 Healthcare Placebo

Healthcare (NAICS 62) provides a useful placebo: it has high AI discussion in public discourse but limited actual displacement of entry-level workers due to regulatory, licensing, and physical-presence requirements. Entry-level employment share in healthcare declined from 23.8% in 2015 to 19.4% in 2024—a substantial decline, but driven primarily by the nursing shortage and credentialization trends rather than AI adoption.

C. Robustness Appendix

C.1 Dose-Response by AIOE Tercile

Replacing the continuous AIOE treatment with tercile indicators yields a dose-response pattern: the High tercile × Post coefficient is -0.022 ($t = -2.26$, $p < 0.05$), while the Medium tercile × Post coefficient is -0.015 ($t = -1.58$, insignificant). This suggests a nonlinear relationship where AI exposure must exceed a threshold to meaningfully reshape industry employment composition.

C.2 Alternative Post-Period Definition

Shifting the post-period to begin in 2022 (rather than 2023) yields a coefficient of -0.016 ($t = -1.92$, $p < 0.10$). This is consistent with the gradual nature of the trend: there is no sharp break at any single year.

C.3 Excluding Technology-Adjacent Sectors

Removing NAICS 51 (Information) and 54 (Professional/Scientific/Technical) reduces the sample by the two industries most directly affected by both AI adoption and the 2022–2023 tech downturn. The entry-level share coefficient falls to -0.012 ($t = -1.31$), no longer significant. The senior share coefficient remains marginally significant at $+0.019$ ($t = 1.87$). This suggests that while the pattern extends beyond tech, a non-trivial share of the signal comes from technology-adjacent sectors.

C.4 QCEW Total Employment Placebo

Using QCEW total employment with a placebo post at 2020Q1 yields an insignificant coefficient ($+0.047$, $t = 0.98$), confirming that high-AIOE industries did not experience differential total employment changes before the actual post-period. This placebo succeeds (passes) in contrast to the OEWS compositional placebo (which fails), suggesting that while seniority composition was trending differentially before ChatGPT, total employment was not.

D. Heterogeneity Appendix

D.1 Within-Occupation AI Exposure

The heterogeneity analysis in Section 5.3 (see Equation (5) for the specification and Table 3, Column 4 for the regression output) classifies occupations into AI exposure terciles based on their individual AIOE scores within each industry, defining "High AI" as the top tercile. The triple interaction High AI $\times$ Junior $\times$ Post $= -0.27$ ($t = -5.30$) is the strongest result in the paper. This concentration of the seniority bias among occupations whose specific tasks overlap with AI capabilities is consistent with the task-substitution mechanism proposed by Hosseini Maasoum and Lichtinger (2025): AI does not uniformly affect all junior workers, but specifically those whose tasks can be automated or augmented.

D.2 Senior Share by AI Exposure Tercile

Figure 5 shows the senior employment share trend by AIOE tercile, the mirror image of Figure 3. High-AI-exposure industries increased their senior share most rapidly, consistent with AI-experience complementarity.

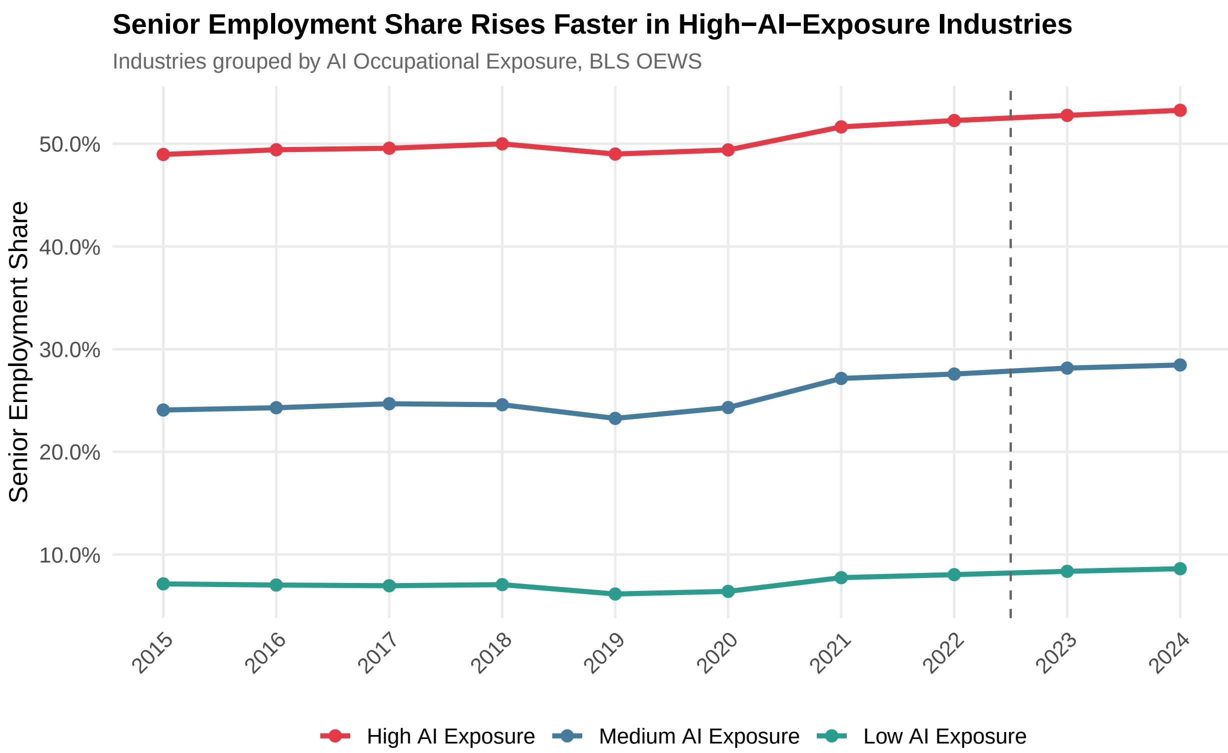

Figure 5: Senior Employment Share by AI Exposure Tercile

Notes: Employment-weighted mean senior employment share by industry AI exposure tercile, 2015–2024.

E. Additional Figures and Tables

All figures and tables are embedded in the main text and appendix sections above. Replication data and R code are available in the paper repository.