

Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data*

Seyed M. Hosseini[†] Guy Lichtinger[‡]

Preliminary

August 2025

Abstract

We study whether generative artificial intelligence (AI) constitutes a form of *seniority-biased technological change*, disproportionately affecting junior relative to senior workers. Using U.S. résumé and job posting data covering nearly 62 million workers in 285,000 firms (2015–2025), we track within-firm employment dynamics by seniority. We identify AI adoption through a text-analysis approach that flags postings for dedicated “AI integrator” roles, signaling active implementation of generative AI. Difference-in-differences and triple-difference estimates show that, beginning in 2023Q1, junior employment in adopting firms declined sharply relative to non-adopters, while senior employment continued to rise. The junior decline is driven primarily by slower hiring rather than increased separations, with the largest effects in wholesale and retail trade. Heterogeneity by education reveals a U-shaped pattern: mid-tier graduates see the largest declines, while elite and low-tier graduates are less affected. Overall, the results provide early evidence of a seniority-biased impact of AI adoption and its mechanisms.

*We thank Larry Katz, Gabriel Chodorow-Reich, Ludwig Straub, David Lagakos, Jesse Shapiro, Jonathon Hazell, Amanda Pallais, Oren Danieli, Aristotle Epanomeritakis, Austin Zhang, Cameron Deal, Santiago Medina, James Stratton, Fiona Chen, and Sarah Gao for helpful comments and discussions. All errors are our own.

[†]Harvard University. Email: shosseinimaasoum@fas.harvard.edu

[‡]Harvard University. Email: guylichtinger@g.harvard.edu

1 Introduction

The impact of artificial intelligence (AI) on juniors, especially in high-skill, white-collar jobs, has attracted growing attention from both researchers and the media. In many such jobs, workers begin at the bottom of the career ladder performing *intellectually mundane* tasks, i.e., routine yet cognitively demanding activities such as debugging code or reviewing legal documents, which are likely to be especially exposed to recent advances in AI. As these workers gain experience, they typically move up the career ladder to more senior roles that involve more complex problem-solving or managerial responsibilities (Becker, 1966; Garicano, 2000). If AI disproportionately targets entry-level tasks, the bottom rungs of these ladders may be eroding.

The stakes go beyond short-term job losses. A large share of college graduates' lifetime wage growth comes from within-firm advancement, starting in low-paid entry roles (Deming, 2023). Early-career earnings also have long-lasting effects on inequality: Guvenen et al. (2022) find that recent increases in U.S. earnings inequality are driven primarily by disparities in starting wages rather than later income growth. If AI disproportionately affects junior positions, it could have lasting consequences for the college wage premium, upward mobility, and income disparities.

The expected impact of generative AI on junior employment is theoretically ambiguous. On one hand, experimental studies suggest that AI tools can significantly enhance the productivity of less-experienced workers (e.g., Noy and Zhang, 2023; Brynjolfsson et al., 2025b), suggesting that AI may complement junior employees and therefore increase demand for their labor. On the other hand, the automation of routine cognitive tasks may directly substitute for the work typically performed by junior workers.

Anecdotal accounts have amplified concerns about such displacement in high-skill industries (e.g., The New York Times, 2025a,b; Independent, 2025; The Atlantic, 2025; The Wall Street Journal, 2025). For instance, a July 2025 *Wall Street Journal* article highlighted a sharp drop in demand for junior workers, citing perspectives from major employers, recruiters, labor market analysts, and recent graduates. One executive at the recruiting firm Hirewell noted that “marketing agency clients have all but stopped requesting entry-level staff—young grads once in high demand but whose work is now a ‘home run’ for AI.” Some observers have also linked AI to recent labor market trends: since late 2022,

unemployment among recent college graduates has risen sharply, even as the unemployment rate for young workers overall has remained steady (Appendix Figure A.1). Others, however, question the importance of AI in these developments, pointing to alternative factors such as economic uncertainty, post-Covid retrenchment, and increased offshoring (e.g., [Financial Times, 2025](#)).

In this paper, we aim to measure the potential *seniority-biased* impact of AI on the labor market. Specifically, we ask whether AI adoption by firms disproportionately affects junior roles relative to more senior positions. This perspective extends the classic literature on skill-biased technological change (SBTC), which emphasizes shifts in labor demand across education or occupation groups (e.g., [Autor et al., 2003](#); [Acemoglu and Autor, 2011](#)), to a related but distinct dimension: *seniority*.

Our analysis draws on a new dataset that combines LinkedIn résumé and job-posting data from Revelio Labs. The dataset covers nearly 285,000 U.S. firms, more than 150 million employment spells from roughly 62 million unique workers between 2015 and 2025, and over 245 million job postings. A key advantage of these data is the standardized seniority classification assigned to each position by Revelio’s algorithm, which enables us to track junior (Entry/Junior) versus senior (Associate and above) employment within firms over time.¹

We develop a novel approach to identifying firm-level AI adoption by detecting job postings that explicitly recruit “AI integrator” roles. The method proceeds in two steps: first flagging postings with AI-related keywords, then using a large language model to determine whether they represent genuine AI-integrator vacancies—roles dedicated to implementing or operating AI systems. A firm is classified as an adopter if it posted at least one such vacancy. This approach allows us to capture firms that have actively begun integrating generative AI into their operations. By this measure, 10,599 firms, about 3.7 percent of our sample, adopted generative AI during the study period.

Our analysis proceeds in several steps. First, we examine aggregate patterns: from 2015 to mid-2022, junior and senior employment grew at similar rates. Beginning in mid-2022, junior employment flattened and, by early 2023, declined, while senior employment continued to rise. These results are consistent with the findings of [Brynjolfsson et al. \(2025a\)](#), who documented a similar employment divergence in U.S. payroll data by

¹See Appendix A.3 for further details and validation of the seniority classification.

worker age.

Next, we compare adopting and non-adopting firms using a difference-in-differences (DiD) framework that tracks quarterly employment outcomes separately for juniors and seniors. From 2015 through 2022, adopters and non-adopters followed parallel trends in junior employment. Beginning in 2023Q1, however, junior headcount at adopting firms fell sharply relative to controls—declining by 7.7 percent after six quarters. Senior employment, by contrast, had been rising more quickly in adopting firms since 2015, and we find no evidence of a post-2022 break in that trend.

To strengthen identification, we employ a triple-difference design that incorporates firm-by-time fixed effects, which absorb any shocks or trajectories unique to each firm at a given moment in time, ensuring that identification comes only from differences between juniors and seniors within the same firm and time. The results confirm the DiD findings. Through 2022, coefficients show only a mild downward trend, consistent with stronger senior growth at adopting firms, while junior relative growth remains stable. Starting in 2023Q1, however, the coefficients decline much more steeply, aligning with the decline in relative junior employment in adopting firms observed in the DiD results.

We then investigate mechanisms behind these shifts. Leveraging our linked employer-employee data, we decompose workforce changes into inflows (hires), outflows (separations), and internal promotions. The relative decline in junior employment at adopting firms is driven primarily by a sharp slowdown in hiring after 2023Q1. Interestingly, separation rates for juniors also fell relative to non-adopters, but the effect is far smaller than the hiring contraction. Moreover, we find that promotions of juniors into more senior roles increased in adopting firms after early 2023.

Industry-level estimates confirm that the decline in junior hiring is widespread but uneven across sectors. The largest reduction occurred in wholesale and retail trade, where adopting firms hired roughly 40% fewer juniors per quarter than non-adopters. This pattern may reflect the greater substitutability of junior tasks in these sectors with generative AI tools, which can automate routine communication, customer service, and documentation. In contrast, across all sectors, changes in senior hiring at adopting firms are either positive or statistically indistinguishable from zero, underscoring that the impact of AI adoption is concentrated at the entry level.

Finally, we explore heterogeneity by educational background as a proxy for human

capital. Using a large language model to classify institutions into five tiers, we uncover a U-shaped pattern. The steepest relative declines occur among graduates of strong (Tier 2) and solid (Tier 3) schools, while the declines are smaller for juniors from elite (Tier 1) and less selective (Tier 4) institutions. Interestingly, the smallest, and statistically insignificant, decline occurs for graduates of the lowest tier (Tier 5). This pattern could suggest that AI adoption is not uniformly depressing demand for juniors but is reshaping it in ways that disproportionately penalize the broad upper-middle segment of the human capital distribution.

We close by acknowledging the limitations of our empirical design. AI adoption is not random: adopting firms are systematically different from non-adopting ones. They tend to be larger, more technologically oriented, and more concentrated in highly educated labor. Although the triple-difference design helps address these concerns by controlling for firm-specific shocks and trajectories, we cannot fully rule out other explanations for the diverging junior-senior patterns by adoption. In the absence of a natural experiment, however, we believe our approach provides the most credible available evidence that the diffusion of generative AI constitutes a form of seniority-biased technological change, with adverse consequences for junior relative to senior employment within the firm.

It is also important to note that only a small fraction of firms in our sample adopted generative AI (3.7 percent). Thus, while the estimated effects for these firms are economically and statistically significant, the implications for aggregate labor market dynamics may be more modest. More broadly, as with any empirical analysis that exploits cross-sectional variation, our results do not necessarily extend to the aggregate economy without additional assumptions, due to the well-known “missing intercept problem.”

The rest of the paper proceeds as follows. Section 2 reviews related literature. Section 3 describes the data sources, variable construction, and descriptive patterns. Section 4 presents the empirical strategy and main results. Section 5 concludes.

2 Related Literature

This section situates our contribution within three strands of literature: (i) skill- and task-biased technological change, (ii) recent evidence on the labor-market effects of generative

AI, and (iii) firm-level studies using postings and résumé data. We highlight how our focus on realized firm adoption and within-firm seniority composition extends exposure-based approaches.

The classic literature on skill-biased technological change shows that computers and automation have historically displaced workers in routine, codifiable tasks while complementing more complex ones. Autor et al. (2003) documented how computerization reduced demand for routine cognitive and manual work, leading to job polarization. Acemoglu and Autor (2011) emphasized that technology replaced mid-skill tasks while raising demand for high-skill labor, and Autor and Dorn (2013) showed that this was accompanied by growth in low-skill service jobs. More recently, Acemoglu and Restrepo (2022) estimated that automation explains a large share of rising U.S. wage inequality since 1980. While this literature focuses on differences across education or occupations, our paper extends the analysis to seniority within firms. We ask whether generative AI is a “*seniority-biased*” technological change, disproportionately affecting juniors who typically perform simpler, more routinized tasks even in high-skill fields.

Since 2023, a rapidly growing empirical literature has examined AI’s labor-market effects. Experimental studies generally find that AI complements less-experienced workers by boosting their productivity. For example, Noy and Zhang (2023) show that access to ChatGPT substantially reduces completion time and improves output quality, with especially large benefits for lower-ability workers. Brynjolfsson et al. (2025b) similarly find that AI assistance in customer support raised productivity by roughly 14 percent on average, with the largest gains for novices. Dell’Acqua et al. (2023) report comparable improvements in consulting workflows. Related field-experimental evidence in Dell’Acqua et al. (2025) shows that adding a generative-AI copilot reshapes teamwork and the division of expertise, shifting routine cognitive work to the tool and reorienting human effort toward higher-level tasks. These findings are consistent with the view that AI can act as a “leveler,” narrowing productivity gaps between less and more experienced workers (Autor, 2024).

A second strand, closer to our study, uses broad-economy data to trace employment trajectories of occupations and industries by AI exposure, yielding mixed results. In a very recent paper, Brynjolfsson et al. (2025a) show that since the late-2022 debut of generative AI, employment of young entry-level workers (ages 22-25) in the most AI-exposed

occupations fell by about 13 percent relative to trend, while more experienced workers in those occupations saw stable or rising employment. [Simon \(2025\)](#) document that entry-level job postings have declined more than 35 percent since January 2023, with the steepest drops in highly exposed roles: a 10-point increase in exposure predicts an 11 percent decline in entry-level demand, while senior roles in those same occupations rise by 7 percent. [Dominski and Lee \(2025\)](#) link occupational exposure scores to CPS data and, using a first-difference design, show that higher AI exposure is associated with reduced employment. By contrast, [Chandar \(2025\)](#) and [Murray et al. \(2025\)](#) do not find systematic differences in employment patterns between more- and less-exposed occupations in CPS data. [Eckhardt and Goldschlag \(2025\)](#) compare unemployment patterns using five exposure measures, finding statistically significant differences for only two, with even those effects relatively small. Our contribution to this literature is to move beyond occupation-level exposure indices and provide broad-based evidence using firm-level adoption, identified via postings for explicit “AI integrator” roles. This design lets us study realized adoption decisions and their within-firm seniority consequences.

A third, related strand—closest to our methodology—examines the implications of firm-level adoption using job postings and résumés. [Babina et al. \(2024\)](#) construct a measure of firm-level AI investment by combining online résumé data from Cognism with Burning Glass postings. Their results suggest that for U.S. firms in the 2010s, AI-adopting firms grow faster in sales, employment, and innovation, with workforces becoming more educated and technologically oriented. [Acemoglu et al. \(2022\)](#) similarly use Burning Glass postings from 2010–2018 to identify AI-exposed establishments based on tasks and skills in vacancies. They find that exposure is associated with lower hiring at the establishment level, but aggregate occupation/industry effects are too small to detect over that period. More recently, [Hampole et al. \(2025\)](#) use data similar to ours, combined with NLP techniques, to construct task-level AI exposure indices. They find that between 2010 and 2023, higher exposure corresponds to lower labor demand, but that firms’ productivity gains offset job losses by expanding employment elsewhere, resulting in muted net changes in total headcount. Together, these studies highlight that pre-2023 AI adoption often entailed internal reallocation rather than aggregate job loss. In contrast, our paper provides evidence on firm-level adoption during the first years of widespread generative AI diffusion (2023–2025). By focusing not only on overall labor demand but on within-firm seniority composition, we provide evidence that AI adoption reduces junior hiring

while leaving senior employment unaffected.

3 Data and Descriptive Patterns

3.1 Data Source and Sample

Our primary data source is a detailed LinkedIn-based résumé dataset provided by Revelio Labs. This dataset contains matched employer-employee information derived from individuals' online profiles. For each worker, we observe all listed employment positions, including job titles, start and end dates, and the employing firm.

Crucially, the dataset also includes a standardized *seniority level* variable for each position, constructed by Revelio through an ensemble modeling approach based on multiple sources of information. This measure combines information from (i) the worker's current job (title, firm, and industry), (ii) their work history (tenure and previous seniority), and (iii) their age. These three inputs produce separate scores, which are averaged into a continuous seniority index and then categorized into seven standardized seniority levels: Entry Level, Junior Level, Associate Level, Manager Level, Director Level, Executive Level, and Senior Executive Level.² In most of the analyses presented below, we differentiate between junior positions (Entry and Junior levels) and senior positions (the remaining five categories).

We complement the worker résumé data with Revelio's database of job postings, which tracks recruitment activity by the firms since 2021. Each posting contains the firm identifier, posting date, job title, standardized occupation codes (mapped to O*NET-SOC), and the raw text description.

The final sample consists of 284,974 U.S. firms that were successfully matched to both employee position data and job postings and that were actively hiring between January 2021 and March 2025.³ For these firms, we observe 156,765,776 positions dating back to 2015 and 245,838,118 job postings since 2021, of which 198,773,384 successfully matched with their raw text description.

²More details on Revelio's seniority classification methodology are available at <https://www.data-dictionary.reveliolabs.com/methodology.html#seniority>.

³We define a firm as active if it recorded at least 20 new hires over this period.

3.2 Key Variables

Workforce Dynamics by Seniority: We construct a monthly panel at the firm level. For each firm-by-month observation, we calculate the number of employees who held a position at the firm that started before and ended after that month, capturing the firm’s workforce size in that period. We repeat this calculation separately for each seniority category, allowing us to track the composition of the workforce across time. In addition, we identify monthly inflows and outflows by seniority. For each firm-by-month, we define new hires as workers who began a new position at the firm that month, having most recently worked at another firm or for whom this is their first observed job. Separations are defined as workers whose position at the firm ended in that month and who either moved to a different firm or exited the labor force (i.e., had no subsequent position listed). Finally, we define promotions as workers who start a new position at the firm after previously holding a lower-seniority role within the same firm.

AI-Adopting Firms: We develop a novel approach to identify firms that adopted AI by detecting job postings explicitly seeking workers to integrate AI technologies into the organization. Ideally, we would apply a large language model (LLM) to analyze the full set of 198,773,384 job postings with raw descriptions. Because such an approach is computationally prohibitive, we proceed in two steps. First, we compile a list of AI-related keywords and flag all postings containing at least one of them.⁴ Out of the 198.8 million postings with raw descriptions, 603,152 contain at least one keyword. Second, we apply an LLM classifier to this subset in order to distinguish genuine “AI integrator” postings—i.e., those reflecting an active attempt to recruit workers tasked with adopting or implementing LLM-based systems—from false positives (appendix A.5.1 provides the exact prompt used). This procedure identifies 131,845 postings, corresponding to 0.066 percent of the full corpus, as AI integrator vacancies.

We then define a firm as an AI adopter if it posted at least one AI integrator vacancy. Based on this criterion, 10,599 firms (3.72 percent of our sample of 284,974 firms) are

⁴Keywords include: Copilot, Claude, Gemini, large language model, LLM, generative AI, ChatGPT, Gen AI, GPT, LangChain, RAG, retrieval-augmented generation, vector embeddings, vector database, transformer-based model, prompt engineering, prompt design, LlamaIndex, Pinecone, Weaviate, Milvus, OpenAI API, Anthropic Claude API, Azure OpenAI, Google Vertex AI Generative, HuggingFace Transformers, and RetrievalQA.

classified as AI adopters. Our methodology is related to [Acemoglu et al. \(2022\)](#), who identified AI-related vacancies in Burning Glass data by searching for 33 AI-related skills extracted from job postings. We extend their approach by applying LLM-based classification, which enables more direct and accurate detection of postings explicitly seeking AI integrators—a capability that was not available at the time of their study.

Below we present two illustrative job postings, identified by the LLM as clear examples of **AI integrator roles**.⁵

Role: Junior Product Manager (Computer and Network Security, Aryaka Networks)

*We are seeking a highly motivated **Junior Product Manager** with a strong understanding of **GenAI security challenges**, hands-on experience in **prompt engineering**, and preferably experience integrating with **GenAI security and safety products/services**. This role involves developing and documenting use cases and requires at least one year of Python programming.*

Key Responsibilities:

- *Collaborate with cross-functional teams to address **GenAI security challenges**.*
- *Apply **prompt engineering** techniques to optimize AI outputs.*
- *Integrate **GenAI security and safety products** into workflows.*
- *Develop and maintain use cases for GenAI applications.*
- *Assist in product features enhancing security and safety.*

⁵As shown in Appendix [A.5.1](#), we instructed the LLM to exclude AI producers. Upon manual inspection, however, we find that some vacancies are engaged in helping other firms integrate LLMs into their workflows, and we classify these postings as integrators. While this does not directly prove that such firms have embedded AI internally, we view it as a reasonable proxy, since firms offering integration services are highly likely to have adopted these technologies themselves.

Role: *Generative AI Developer Consultant (IT Services and IT Consulting, Genesis10)*

Summary: *We are seeking a talented and motivated **Software Engineer** to join our team, focusing on developing innovative applications using **Generative AI technologies**. You will play a key role in **designing, building, and deploying** solutions that leverage AI to transform user experiences.*

Responsibilities:

- *Design and develop scalable applications utilizing **Generative AI models**.*
- *Collaborate with cross-functional teams to deliver solutions.*
- ***Integrate AI models** into existing systems and applications.*
- *Optimize and fine-tune AI algorithms for performance and accuracy.*
- *Conduct code reviews and mentor junior team members.*

...

School Quality: To capture the quality of educational background, we construct the a school quality variable by assigning each position the rating of the school attended by the individual holding it. For each position, we assign the school of the most recent education that ended no later than one year after the job start; if no such spell exists, we fall back on the individual's first recorded education, provided it began before the position start date. We then merge these universities with an external school list in which each institution is rated on a 1–5 scale. The ratings were generated using OpenAI's GPT-4o-mini model, prompted to act as an academic evaluator and to assign a single integer score, where 1 corresponds to Ivy/elite global tier (e.g., Harvard, Stanford, MIT), 2 to very strong and internationally respected institutions, 3 to solid national or regional universities, 4 to lower-tier but standard universities, and 5 to very weak or diploma-mill institutions (see Appendix A.5.2 for the exact prompt).

Firm AI Exposure: Although not the primary focus of the paper, we also examine firms' exposure to generative AI. To construct this measure, we draw on the task-level exposure

indices developed by Eloundou et al. (2024). In their framework, exposure is defined using a rubric applied by both human annotators and GPT-4 to O*NET tasks: a task is considered exposed if an LLM can reduce the time required for completion by at least 50% while preserving or improving quality. Because our analysis focuses on relatively early effects of generative AI, we use the “alpha” measure, which captures exposure from an LLM alone or through a simple interface, rather than the “beta” or “gamma” measures that also incorporate tasks requiring complementary software to realize the 50 percent time savings. Following their weighting scheme, we compute occupation-level exposure as the weighted average of core tasks (full weight) and supplemental tasks (half weight). Finally, we aggregate to the firm level by taking the average exposure across all the vacancies posted by the firm in 2022, just prior to the widespread diffusion of generative AI.

3.3 Descriptive Patterns

Table 1 reports descriptive statistics for the full sample, AI adopters, and non-adopters. Several systematic differences emerge. Firm size and workforce structure stand out: AI adopters are far larger, averaging almost 500 employees compared to about 100 among non-adopters. Their workforces are more senior-intensive, with juniors representing just 42 percent of employment versus 55 percent in non-adopters. Consistently, adopters experience substantially higher hiring and separation volumes, and a larger share of these flows involve senior positions. Adopters also tend to hire from higher-quality universities, with fewer juniors from tier-5 (lowest quality) institutions and more from tier-1 and tier-2 schools. As expected, adopters score substantially higher on the Eloundou et al. (2024) alpha exposure index, indicating that they employ workers in occupations with a greater share of exposed tasks.

While non-adopters are widely dispersed across sectors, AI adopters are heavily concentrated in information (36 percent) and professional services (25 percent), both knowledge-intensive industries where AI adoption is most salient. Geographically, adopters are disproportionately headquartered in California (21 percent versus 14 percent overall), while being slightly less represented in Texas compared to non-adopters. Appendices A.2 and A.4 provide further detail on adopters’ distribution across industries and states.

Table 1: Descriptive Statistics: All Firms, AI Adopters, and Non-Adopters

Variable	All Firms	Non-Adopters	AI Adopters
Panel A. Workforce Composition and Characteristics			
Firm size (avg. employees)	116.0 (426.7)	101.3 (360.1)	498.1 (1177.4)
Share junior employees (Entry/Junior)	0.548 (0.228)	0.553 (0.227)	0.422 (0.211)
Share senior employees (Associate+)	0.454 (0.228)	0.450 (0.227)	0.581 (0.211)
Average number of new hires (per quarter)	6.5 (23.9)	5.7 (20.6)	27.1 (62.6)
Share of new hires junior	0.670 (0.351)	0.677 (0.351)	0.508 (0.307)
Share of new hires senior	0.330 (0.351)	0.323 (0.351)	0.492 (0.307)
Average number of separations (per quarter)	5.2 (21.2)	4.6 (18.5)	21.1 (53.8)
Share of separations junior	0.686 (0.351)	0.692 (0.351)	0.531 (0.311)
Share of separations senior	0.314 (0.351)	0.308 (0.351)	0.469 (0.311)
Average number of promotions for juniors (per quarter)	0.531 (2.4)	0.439 (1.8)	2.9 (7.6)
Share juniors from tier-1 universities (highest quality)	0.034 (0.091)	0.033 (0.090)	0.054 (0.109)
Share juniors from tier-2 universities	0.107 (0.142)	0.107 (0.142)	0.116 (0.130)
Share juniors from tier-3 universities	0.209 (0.175)	0.210 (0.176)	0.175 (0.148)
Share juniors from tier-4 universities	0.154 (0.142)	0.156 (0.143)	0.112 (0.113)
Share juniors from tier-5 universities (lowest quality)	0.094 (0.113)	0.095 (0.113)	0.064 (0.087)
Average exposure index (2022)	0.346 (0.165)	0.344 (0.167)	0.385 (0.093)
Panel B. Industry and Headquarters Location			
Share in NAICS sector 51 (Information)	0.121	0.110	0.360
Share in NAICS sector 52 (Finance and Insurance)	0.078	0.077	0.082
Share in NAICS sector 54 (Professional Services)	0.179	0.175	0.246
Share in NAICS sector 5 (Other)	0.063	0.063	0.068
Share in non-NAICS 5 sectors	0.559	0.575	0.244
HQ in California	0.140	0.138	0.210
HQ in Texas	0.075	0.076	0.060
HQ in New York	0.081	0.081	0.093
HQ in Other States	0.711	0.713	0.642
Observations	11,674,996	11,240,847	434,149
Number of firms	284,756	274,167	10,589

Notes: The table reports averages of the main variables across firm-by-quarter observations from 2015Q1 to 2025Q1, separately for the full sample, AI adopters, and non-adopters. Standard deviations (for non-binary variables) are reported in parentheses. Panel A reports workforce composition, such as hiring and separations, workers' education background, and automation exposure. Panel B reports industry and headquarters state distributions.

Taken together, the statistics depict AI adopters as larger, more senior-oriented firms, with higher volumes of worker flows, stronger recruitment from elite institutions, and greater presence in technology-intensive sectors and states—all features that align with the environments where AI technologies are most likely to take root.

Next, Figure 1 documents the timing of the first AI integrator vacancy posted by firms. As detailed in Section 3.2, we classify a firm as an AI adopter if it posted at least one AI integrator vacancy during the study period. Prior to 2023, the number of new firms posting such vacancies was small and stable, averaging about 30 per month. Beginning in early 2023, the count rose sharply, peaking at 456 in August 2023. Thereafter, it remained steady at around 400 firms per month through the end of 2024. In early 2025, the numbers increased again, reaching 574 new firms in March 2025. By the end of the sample period, the cumulative number of adopters exceeded 10,000 firms.




Figure 1: Timing of Firms' First AI-Integrator Vacancies

Notes: Panel (a) shows the monthly number of firms posting their first AI-integrator vacancy, while Panel (b) reports the cumulative total, covering the period from September 2021 to March 2025. See Section 3.2 for details on the identification of AI-integrator postings.

Finally, Figure 2 presents an aggregate time series of junior and senior employment in our data. We define “junior” workers as those in Entry- or Junior-level positions, and “senior” workers as Associate level and above (see Section 3.1 for details). The figure shows the average number of workers in each group (across all firms in our sample) over time, normalized to 1 in January 2015.

From the start of the period through early 2020, junior and senior employment tracked each other closely, growing at nearly identical rates. During the onset of the COVID-19

pandemic, junior employment fell slightly more than senior employment but recovered quickly, with both groups returning to a similar growth rate until mid-2022. Beginning in mid-2022, however, a marked divergence emerges. Senior employment continued to expand steadily, while junior employment flattened out. By mid-2023, the gap widened further: senior employment kept rising at roughly the same pace, but junior employment began to decline. This pattern is consistent with the findings of [Brynjolfsson et al. \(2025a\)](#), who documented a similar employment divergence in U.S. payroll data by worker age. The alignment between their results and ours provides external validation for the patterns observed in our LinkedIn-based dataset.




Figure 2: Time Series of Junior and Senior Employment in Sample Firms

Notes: This figure plots the average number of junior-level workers and senior-level workers in our sample of firms over time, normalized to 1 in January 2015. We define “junior” workers as those in Entry- or Junior-level positions, and “senior” workers as Associate level and above (see Section 3.1 for details).

4 Empirical Strategy and Results

4.1 Workforce Dynamics by Adoption

We examine how the adoption of generative AI technologies affects employment dynamics by comparing adopting and non-adopting firms.⁶ Our first approach is a difference-in-differences (DiD) specification, estimated separately for junior and senior workers:

$$\log(\text{Employment}_{it}) = \alpha + \sum_{j=2015Q1}^{2025Q1} \beta_j \mathbf{1}\{t = j\} \times \text{Adopt}_i + \delta_t + \text{Adopt}_i + \varepsilon_{it}, \quad (1)$$

where the dependent variable $\log(\text{Employment}_{it})$ denotes the log employment of junior (or senior) workers at firm i in period t . The term $\mathbf{1}\{t = j\}$ is an indicator function that equals one if $t = j$ and zero otherwise, so that the coefficients β_j capture the differential evolution of employment for adopters relative to non-adopters in each period j . The variable Adopt_i is a dummy equal to one for firms that adopt generative AI, as defined in Section 3.2. Time fixed effects δ_t absorb aggregate shocks common to all firms, while Adopt_i controls for time-invariant differences between adopters and non-adopters. The error term ε_{it} captures unobserved idiosyncratic determinants of employment.

Figure 3 reports the estimated coefficients β_j . For junior workers, the coefficients are flat and indistinguishable from zero through 2022Q4, consistent with parallel pre-trends. Starting in 2023Q1, they turn sharply negative, indicating that junior employment in adopting firms fell by 7.7 percent relative to controls six quarters after the diffusion of generative AI. By contrast, coefficients for senior workers show a persistent upward trajectory throughout the sample, suggesting that adopting firms expanded senior employment more strongly than non-adopters over the last decade.

⁶For computational efficiency we aggregate our panel data to quarterly frequency in all analyses.




Figure 3: Employment Differences Between Adopters and Non-Adopters Over Time

Notes: The graph present the estimated coefficients β_j from Equation 1, ran separately to juniors and seniors. Standard errors are clustered in firm level.

We next account for firm-specific employment trajectories using a triple-differences specification:

$$\begin{aligned}
 \log(\text{Employment}_{ist}) = & \alpha + \sum_{j=2015Q1}^{2025Q1} \beta_j \mathbf{1}\{t = j\} \times \text{Adopt}_i \times \text{Junior}_s \\
 & + \sum_{j=2015Q1}^{2025Q1} \pi_j \mathbf{1}\{t = j\} \times \text{Adopt}_i + \sum_{j=2015Q1}^{2025Q1} \rho_j \mathbf{1}\{t = j\} \times \text{Junior}_s \\
 & + \kappa (\text{Adopt}_i \times \text{Junior}_s) + \gamma_{it} + \varepsilon_{ist},
 \end{aligned} \tag{2}$$

where $\log(\text{Employment}_{ist})$ denotes the log employment of workers in seniority group $s \in \{\text{junior}, \text{senior}\}$ at firm i in period t . The indicator $\mathbf{1}\{t = j\}$ equals one in period j and zero otherwise. Adopt_i is a firm-level dummy equal to one for firms that adopt generative AI (see Section 3.2 for the definition), and Junior_s is an indicator equal to one for juniors and zero for seniors. The coefficients β_j form a triple-differences event-time profile, comparing juniors to seniors *within* firm i in period j , for adopters relative

to non-adopters. Firm-by-time fixed effects γ_{it} absorb arbitrary shocks at the firm-time level, ensuring identification comes from the within-firm-time junior-senior contrast. The main identification assumption is that, after accounting for firm-by-time fixed effects and lower-order interactions, no other factors systematically affect juniors and seniors differently in adopting versus non-adopting firms.

Figure 4 presents the results. Because of computational constraints, estimation begins in 2018Q1. Between 2018Q1 and 2022Q4, the coefficients show a mild downward trend. Beginning in 2023Q1, however, we observe a clear break: the coefficients decline more steeply, indicating that the relative employment of juniors fell within adopting firms compared to seniors by 12%. This pattern aligns with the DiD results in Figure 3: senior employment in adopting firms grew consistently throughout the sample, while junior employment was flat until 2022Q4 and then dropped sharply. The trend break coincides with the diffusion of generative AI in early 2023, providing suggestive evidence that AI constitutes a seniority-biased technological change—reducing junior employment while leaving senior employment unaffected.




Figure 4: Triple Differences Results

Notes: The graph present the estimated coefficients β_j from Equation 2. Standard errors are clustered in firm level.

4.2 Hires, Exits, and Promotions

To decompose these adjustments, we estimate separate difference-in-differences regressions of the form

$$y_{it} = \alpha + \beta (\text{Adopt}_i \times \text{Post}_t) + \gamma_i + \delta_t + \varepsilon_{it}, \quad (3)$$

where y_{it} denotes the number of hires, separations, promotions, or net changes in firm i at time t . The variable Adopt_i is an indicator equal to one for AI-adopting firms, and Post_t equals one for periods beginning in 2023Q1 and zero otherwise.⁷ γ_i denotes firm fixed effects and δ_t are time fixed effects.

Table 2 presents the results of Equation 3. The results reveal that the sharp contraction in junior employment observed among adopters is primarily driven by a reduction in hiring rather than by an increase in exits. Specifically, the coefficient on *Hiring* indicates that, relative to non-adopters, AI-adopting firms hired on average 3.7 fewer junior workers after 2023Q1. Interestingly, separation rates for juniors also declined among adopters relative to non-adopters, but the effect is small, less than one worker on average, and much weaker than the contraction in hiring.

Table 2: Effects of AI Adoption on Hiring, Separation, Promotion, and Total Change

	Hiring	Separation	Promotion	Total Change
<i>Panel A: Juniors</i>				
Treat × Post	−3.694*** (0.227)	−0.788*** (0.173)	0.391*** (0.033)	−3.436*** (0.146)
<i>Panel B: Seniors</i>				
Treat × Post	0.821*** (0.141)	2.271*** (0.114)	—	−0.728*** (0.137)
Observations	11,683,934	11,683,934	11,683,934	11,398,960
Clusters (firms)	284,974			

Notes: Standard errors clustered by firm in parentheses. We control for firm and time fixed effects in all columns.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

⁷We abstract from heterogeneous adoption timing and assume all adopters as treated beginning in 2023Q1. This approach may attenuate the estimated effects, since actual adoption rates rise over time starting in 2023.

Note that the magnitudes of these effects are economically large. Before the first quarter of 2023, the average number of junior hires in AI-adopting firms was 17.45. Hence, our estimates imply that adopting firms reduced their junior hiring relative to non-adopters by roughly 22% of their average hiring. By contrast, AI adopters *increased* their hiring of senior workers by about 0.8 workers relative to non-adopters, but simultaneously experienced an additional 2.2 senior separations, resulting in a small net decline in relative senior employment growth.

Taken together, these findings imply that the reduction in junior headcount within adopting firms is not the result of layoffs or elevated attrition, but instead reflects a slowdown in new entry. This pattern is consistent with firms strategically scaling back junior hiring once generative AI tools become available, while continuing to retain their existing workforce.

Interestingly, the promotion regressions point to a complementary adjustment margin. Adopting firms increased the number of promotions granted to junior workers by about 0.4 on average, relative to non-adopters. This suggests that while fewer juniors are being recruited, those who remain may face enhanced opportunities for internal advancement. One interpretation is that generative AI substitutes for entry-level tasks, reducing demand for new juniors, but simultaneously raises the relative value of experienced juniors, who are more likely to be promoted into senior roles.

4.3 Impact on Hiring Across Sectors

To complement the employment results, we estimate the effect of generative AI adoption on the number of hires by sector. Specifically, we run regressions of the form:

$$(\text{Hires}_{it}) = \alpha + \beta (\text{Adopt}_i \times \text{Post}_t) + \delta_t + \gamma_i + \varepsilon_{it}, \quad (4)$$

where Hires_{it} denotes the number of new junior or senior workers hired by firm i at time t , Adopt_i is an indicator for adopters of generative AI, and Post_t is an indicator for the post-adoption period, defined as after the first quarter of 2023. We also control for time and firm fixed effects. Figure 5 reports the results across major sectors.

We find that, consistently across all industries, adopting firms significantly reduced




Figure 5: Estimated Effects of Generative AI Adoption on Hiring by Sector

Note: Sectors correspond to the following NAICS classifications: Manufacturing (31–33), Wholesale/Retail (42, 44–45), Information (51), Finance & Insurance (52), Professional Services (54), Education (61), and Health Care & Social Assistance (62). All coefficients are normalized by the pre-2023 average number of hires in each sector, in order to account for differences in baseline labor turnover across industries. Standard errors are clustered at the firm level.

the hiring of junior workers, while there is no statistically significant impact on the hiring of senior workers. Interestingly, we observe that adopting firms in the wholesale and retail sector experienced the largest decrease in hiring, reducing junior hires by around 40% per quarter compared to non-adopters. This pattern may reflect the greater substitutability of junior tasks in these sectors with generative AI tools, which can automate routine communication, customer service, and documentation.

4.4 Heterogeneity by Human Capital

In this section, we explore heterogeneity in employment declines among junior workers by distinguishing them according to educational background. Specifically, we use the quality of the institutions they attended prior to employment as a proxy for human capital (see Section 3.2 for details on the school quality measure). To capture these differences, we

estimate the following regression separately for each of five categories of juniors, defined by school prestige:

$$\log(\text{Junior Employment}_{it}) = \alpha + \beta (\text{Adopt}_i \times \text{Post}_t) + \delta_t + \varepsilon_{it}. \quad (5)$$

The results, presented in Figure 6, reveal a U-shaped pattern. Juniors from tier-2 and tier-3 universities experienced the sharpest relative declines in employment. Those from tier-1 (elite) and tier-4 schools also saw reductions, but of smaller magnitude. By contrast, juniors from tier-5 (lower-prestige) institutions experienced only a modest decline, statistically indistinguishable from zero.




Figure 6: Results by School Quality

Notes: The figure presents the estimates of Equation 5, run separately for juniors by university prestige category. The coefficients represent post-adoption changes in junior employment for AI adopters relative to non-adopters. Standard errors are clustered at the firm level.

Figure 7 provides additional context by plotting average pre-adoption salaries of juniors across school-quality categories. Not surprisingly, Juniors from elite (tier-1) schools earn the highest salaries, while juniors from lower-prestige (tier-5) institutions earn roughly half as much.

The combination of Figures 6 and 7 implies a trade-off: at the very top, high productivity/quality may discourage firms from reducing hiring aggressively despite generative

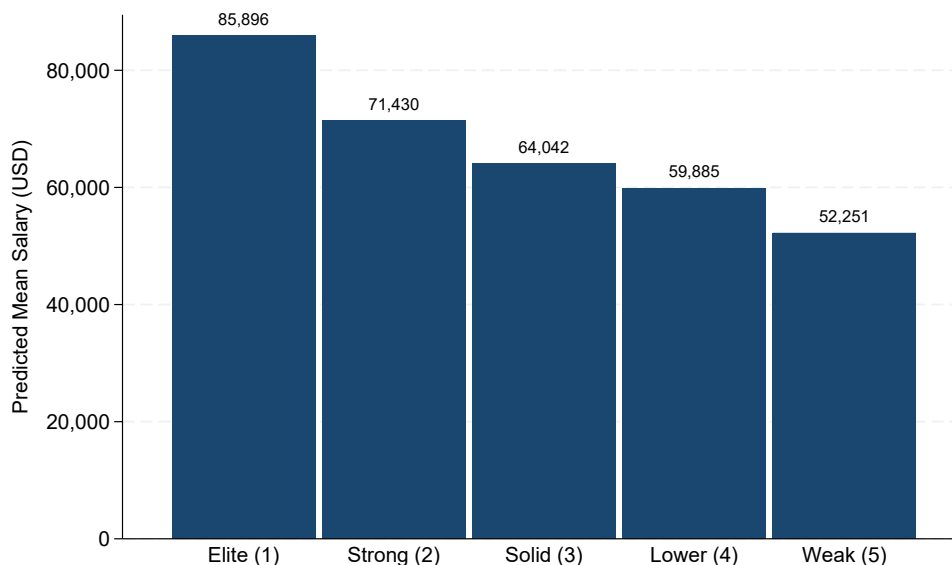


Figure 7: Predicted Salary by School Quality (Juniors, 2022)

Notes: Bars report average predicted salaries (in USD) for juniors employed in 2022, by university prestige category. Salaries decline monotonically with school prestige, with juniors from elite schools (tier-1) earning the most and juniors from lower-prestige schools (tier-5) earning roughly half as much.

AI adoption, while at the very bottom, low costs shield lower-prestige juniors from displacement. By contrast, mid-tier juniors—who combine moderate cost with moderate quality—appear most vulnerable to substitution, experiencing the steepest reductions in employment.

5 Conclusion

This paper provides early, broad-based evidence that the diffusion of generative AI since 2023 is associated with *seniority-biased technological change* within firms. Using résumé-posting data linked to nearly 285,000 U.S. firms and a direct measure of adoption based on “AI integrator” vacancies, we document a sharp relative decline in junior employment at adopting firms, alongside continued growth in senior employment. Our difference-in-differences estimates indicate flat pre-trends for juniors from 2015 to 2023, followed by a discrete break in 2023Q1. A triple-difference specification with firm-by-time fixed effects confirms that the within-firm junior-senior gap widens precisely as generative AI diffuses.

The decline in junior employment is driven primarily by reductions in *hiring*. Adopting firms substantially curtailed junior hiring after 2023Q1, with exits also decreasing modestly, implying that net declines occurred mainly through slower entry rather than layoffs. At the same time, promotions of existing juniors increased, suggesting that while firms recruit fewer new entrants, those already in the firm may benefit from greater opportunities to move into senior roles. Sector-level evidence shows that this pattern is not limited to information technology: wholesale and retail exhibit the largest cuts in junior hiring, with smaller but still meaningful declines in several other major sectors.

We also find heterogeneity by educational background. The decline in junior employment follows a U-shaped pattern: graduates from mid-tier institutions are most affected, while those from elite and lower-tier schools experience smaller reductions, and the lowest tier shows negligible effects. Together with pre-adoption salary differentials, this pattern is consistent with a simple trade-off: at the top, higher productivity discourages firms from reducing hiring too sharply; at the bottom, lower costs shield juniors from displacement; in the middle, juniors appear most vulnerable to substitution.

These findings should be interpreted with caution. AI adoption is not random, and adopters differ systematically in size, workforce composition, and sector. Although the triple-difference design absorbs firm-specific shocks and trajectories, confounding factors cannot be ruled out. Our adoption measure, based on integrator postings, captures realized organizational uptake but may miss forms of adoption that occur without explicit hiring. Finally, the window of analysis is relatively short (2023–2025), and longer-run adjustments in training, task design, and internal career ladders may either mitigate or amplify these initial effects.

Even with these caveats, our results point to important implications. Generative AI appears to shift work away from entry-level tasks, narrowing the “bottom rungs” of internal career ladders. Because early-career jobs play a critical role in lifetime wage growth and mobility, these dynamics may carry lasting consequences for inequality and the college wage premium. Firms, for their part, may adapt by relying more heavily on experienced workers and by accelerating the promotion of incumbents. Taken together, our evidence indicates that generative AI is *seniority-biased*, with significant implications for how careers begin, how firms develop talent, and how the benefits of new technologies are distributed.

References

- Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, volume 4, pages 1043–1171. Elsevier.
- Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.
- Acemoglu, D. and Restrepo, P. (2022). Tasks, automation, and the rise in us wage inequality. *Econometrica*, 90(5):1973–2016.
- Autor, D. (2024). Applying ai to rebuild middle class jobs. Technical report, National Bureau of Economic Research.
- Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the u.s. labor market. *American Economic Review*, 103(5):1553–1597.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4):1279–1333.
- Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.
- Becker, G. (1966). Human capital: A theoretical and empirical analysis, with special reference to education the residual factor and economic growth econometric models of education.
- Brynjolfsson, E., Chandar, B., and Chen, R. (2025a). Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Working paper. Latest version available at <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/>.
- Brynjolfsson, E., Li, D., and Raymond, L. (2025b). Generative ai at work. *The Quarterly Journal of Economics*, page qjae044.
- Chandar, B. (2025). Tracking employment changes in ai-exposed jobs. Available at SSRN 5384519.

- Dell'Acqua, F., Ayoubi, C., Lifshitz, H., Sadun, R., Mollick, E., Mollick, L., Han, Y., Goldman, J., Nair, H., Taub, S., et al. (2025). The cybernetic teammate: A field experiment on generative ai reshaping teamwork and expertise. Technical report, National Bureau of Economic Research.
- Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., and Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (24-013).
- Deming, D. J. (2023). Why do wages grow faster for educated workers? Technical report, National Bureau of Economic Research.
- Dominski, J. and Lee, Y. S. (2025). Advancing ai capabilities and evolving labor outcomes. *arXiv preprint arXiv:2507.08244*.
- Eckhardt, S. and Goldschlag, N. (2025). AI and jobs: The final word (until the next one). Economic Innovation Group. Accessed 2025-08-30.
- Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308.
- Financial Times (2025). Is ai killing graduate jobs? <https://www.ft.com/content/996b6acb7-a079-4f57-a7bd-8317c1fbb728?shareType=nongift>. By Clara Murray, Delphine Strauss, John Burn-Murdoch, and Sarah Lim. Published July 24, 2025.
- Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of political economy*, 108(5):874–904.
- Guvenen, F., Kaplan, G., Song, J., and Weidner, J. (2022). Lifetime earnings in the united states over six decades. *American Economic Journal: Applied Economics*, 14(4):446–479.
- Hampole, M., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2025). Artificial intelligence and the labor market. Technical report, National Bureau of Economic Research.
- Independent (2025). Ai is wrecking an already fragile job market for college graduates. <https://www.wsj.com/lifestyle/careers/>

- [ai-entry-level-jobs-graduates-b224d624](#). Published in *The Wall Street Journal*, July 29, 2025. Appeared in print as "AI Is Killing the Entry-Level Job" on July 31, 2025.
- Murray, C., Strauss, D., Burn-Murdoch, J., and Lim, S. (2025). Is ai killing graduate jobs? *Financial Times*. Accessed via <https://www.ft.com/content/99b6acb7-a079-4f57-a7bd-8317c1fbb728>.
- Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.
- Simon, L. K. (2025). Is ai responsible for the rise in entry-level unemployment? <https://www.reveliolabs.com/news/macro/is-ai-responsible-for-the-rise-in-entry-level-unemployment/>. Revelio Labs, macro section.
- The Atlantic (2025). Something alarming is happening to the job market. <https://www.theatlantic.com/economy/archive/2025/04/job-market-youth/682641/>. By Derek Thompson. Published in *The Atlantic*, April 2025.
- The New York Times (2025a). For some recent graduates, the a.i. job apocalypse may already be here. <https://www.nytimes.com/2025/05/30/technology/ai-jobs-college-graduates.html>. By Kevin Roose. Published May 30, 2025. Appeared in print as "Foot in Door? Not With A.I. Doing the Job."
- The New York Times (2025b). LinkedIn executive: A.i. is coming for entry-level jobs. Accessed: 2025-06-20.
- The Wall Street Journal (2025). The 'great hesitation' that's making it harder to get a tech job. Accessed: 2025-06-20.

A Appendix

A.1 College Graduates Unemployment




Figure A.1: Unemployment Rates for Recent College Graduates vs. All Young Workers

Notes: The orange line shows the unemployment rate for recent U.S. college graduates (aged 22–27 with a bachelor’s degree), and the blue line shows the unemployment rate for all U.S. workers aged 22–27, on a monthly basis. Since late 2022, the college-graduate rate has risen even as the overall young worker rate remained flat. Source: Federal Reserve Bank of New York, *Labor Market for Recent College Graduates*.

A.2 Geography of Adopters

Figure A.2 shows the distribution of adopters across U.S. states. As expected, adoption is highly concentrated in California, which alone accounts for about 26% of all adopters. Importantly, this share, while large, is not a majority. The top five adopter states are California, New York, Texas, Massachusetts, and Virginia, all of which are technology-intensive regions. This pattern highlights that AI adoption is not exclusively a Silicon Valley phenomenon but instead spans several large, tech-heavy states.

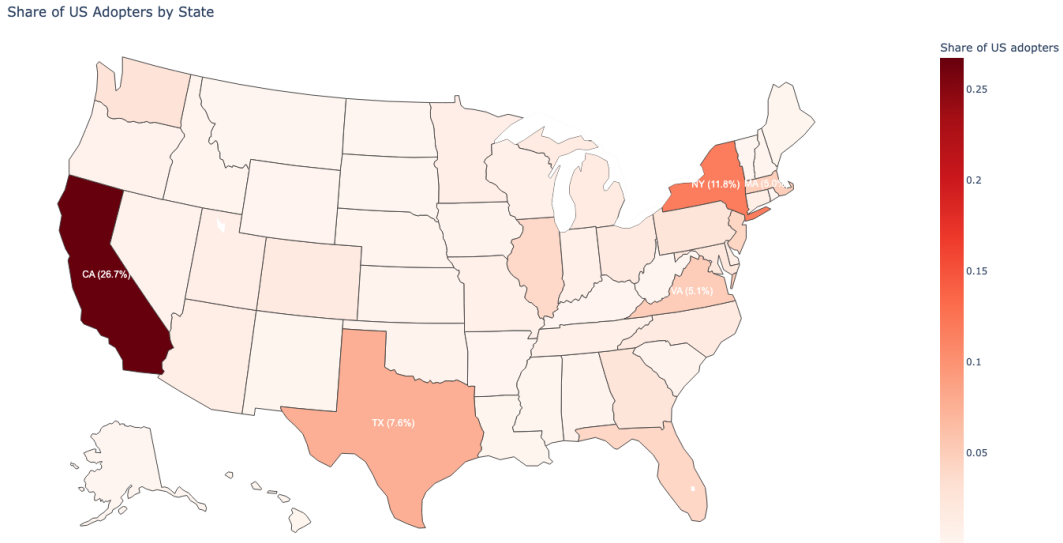


Figure A.2: Share of U.S. adopters by state.

A.3 Verification of the Seniority Variable

To validate the seniority variable provided by Revelio Labs, we conducted a role keyword frequency analysis. The goal was to check whether the job titles associated with lower-seniority workers (“Seniority 1 and 2”) and higher-seniority workers (“Seniority 3+”) match intuitive expectations.

For junior roles (Seniority 1 and 2), the most frequent keywords are *assistant*, *specialist*, *technician*, and *intern*, consistent with entry-level or supporting positions. By contrast, for senior roles (Seniority 3+), the dominant keywords are *manager*, *director*, and *consultant*, which are associated with leadership and higher-responsibility positions. This pattern provides strong support for the validity of the seniority classification.

Figures A.3–A.6 illustrate these distributions using both bar charts and word clouds. The bar charts report the percentage share of each role keyword, while the word clouds visualize their relative frequencies in a more intuitive manner.




Figure A.3: Role keyword frequency — Seniority 1 and 2 (percent).




Figure A.4: Role keyword frequency — Seniority 3+ (percent).

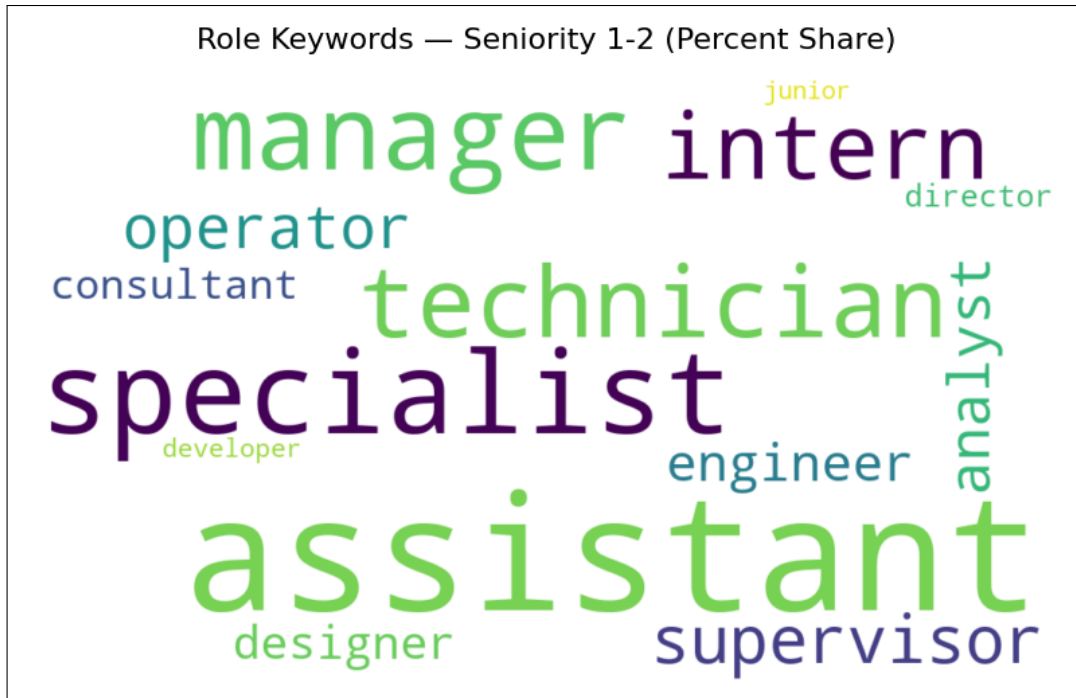


Figure A.5: Role keywords — Seniority 1-2 (percent share).

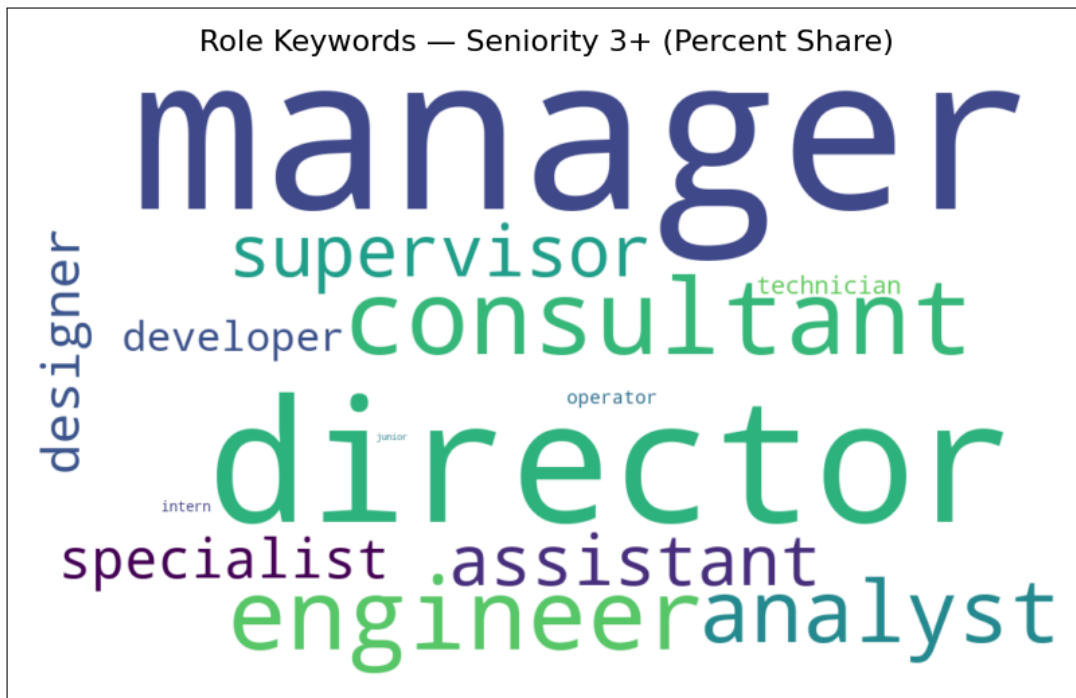


Figure A.6: Role keywords — Seniority 3+ (percent share).

A.4 Sectoral Distribution of AI Adopters

In this appendix, we provide descriptive evidence on the sectoral distribution of AI adoption. Figure A.7 documents the share of firms in each major sector that have adopted AI, while Figure A.8 shows the distribution of *adopters only*, i.e., the fraction of adopting firms that belong to each sector. These figures highlight that adoption is not concentrated in a single industry, but rather spread across information, professional services, finance, manufacturing, and other sectors. As expected, adoption is somewhat higher in knowledge- and technology-intensive industries, but traditional sectors such as manufacturing and wholesale/retail are also represented.




Figure A.7: Share of adopters across sectors. Notes: Reports the distribution of all firms across sectors. Sectors correspond to the following NAICS codes: Manufacturing (3), Wholesale/Retail (4), Information (51), Finance and Insurance (52), Professional Services (54), Education (61), and Health Care and Social Assistance (62).




Figure A.8: Distribution of adopters across sectors. Notes: Reports the distribution restricted to AI adopters. Sectors correspond to the following NAICS codes: Manufacturing (3), Wholesale/Retail (4), Information (51), Finance and Insurance (52), Professional Services (54), Education (61), and Health Care and Social Assistance (62).

A.5 API Prompts

A.5.1 Prompt: Identify Job Postings for AI Integrators or Users

We use llama-3.1-8b-instant model through groq api.

`SYSTEM = '''`

You are a precise classifier for job postings. Output ONLY compact JSON.

We distinguish two categories:

A) LLM INTEGRATOR = roles that build/operate LLM-powered systems or embed LLMs into workflows.

Signals: RAG (retrieval-augmented generation), embeddings/vector DB (FAISS/Milvus/Pinecone),

prompt engineering at system level, orchestration/agents/LLMOps, LangChain/Lla-

maIndex,
fine-tuning/adapters, model serving/inference, evaluation/guardrails/red-teaming,
API integration of LLMs into products or internal processes.

B) LLM USER = roles primarily using LLM tools (ChatGPT, Gemini, Copilot, etc.) to perform tasks
such as drafting, summarizing, coding assistance, customer responses—without building systems.

NOT in-scope for integrator unless integration is explicit:
– Foundation-model pretraining/research scientist roles at model labs (OpenAI/DeepMind/etc.).
– Generic ML/NLP with no explicit LLM signals.
– Pure labeling/annotation.

Edge rules:
– If both integration and user aspects appear, set role_type="both".
– If acronyms like "RAG" appear, assume the LLM meaning unless context contradicts.
– Prefer TRUE for integrator/user when listed signals appear.
– Output JSON only; no prose.

'''

A.5.2 Prompt: School Quality Rating

We use 4o-mini model through openai api.

SYSTEM_PROMPT = '''You are an academic evaluator.
Assign each input university a single integer rating on this scale:
1 = Ivy/elite global tier (e.g., Harvard, Stanford, Oxford, MIT)
2 = Very strong, internationally respected
3 = Solid national/regional reputation

4 = Lower tier/less selective but standard university

5 = Very weak / diploma-mill territory

Return ONLY what is requested. No commentary, no markdown.

When uncertain, choose the closest reasonable tier using overall global reputation.

'''

*USER_PROMPT_TEMPLATE = '''Rate the following {n} institutions on the
1–5 scale described.*

INSTRUCTIONS (STRICT):

- Return EXACTLY {n} lines.*
- Each line contains ONLY one integer in 1..5 for the corresponding line below.*
- Do NOT include any keys, bullets, indexes, punctuation, or extra text.*
- Do NOT include blank lines.*
- STOP OUTPUT immediately after printing the {n}th line.*

NAMES (one per line, in order):

{names_block}

'''