

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

Seyed M. Hosseini[†]

Guy Lichtinger[‡]

First Version: August 31, 2025

This Version: October 6, 2025

Abstract

We study whether generative AI (GenAI) constitutes *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 GenAI adoption through text analysis that flags job postings for dedicated “GenAI integrator” roles, signaling active implementation of GenAI. 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. A staggered difference-in-differences design exploiting variation in adoption timing across firms supports these results. The junior decline is concentrated in high-exposure jobs, and is driven by slower hiring rather than increased separations or promotions. Overall, the results provide early evidence of a seniority-biased effect associated with GenAI adoption and its mechanisms.

*We are extremely grateful to Lawrence Katz and Jesse Shapiro for their continued guidance and feedback on this project. We are also deeply thankful to Erik Brynjolfsson, Gabriel Chodorow-Reich, Oren Danieli, Jonathon Hazell, David Lagakos, Amanda Pallais, and Ludwig Straub for their invaluable discussions and advice throughout various stages of this work. We additionally thank Abhishek Arora, Shotaro Beppu, Olena Bogdan, Bharat Chandar, Ruyu Chen, Fiona Chen, Rupsha Debnath, Cameron Deal, Aristotle Epanomeritakis, Sarah Gao, Eric Gold, Jay Garg, Ori Heffetz, Shao-Yu Jheng, Xianruo Kang, Iris Li, Santiago Medina, Saketh Prazad, Ewan Rawcliffe, Emanuel Schertz, Peyman Shahidi, Ragini Srinivasan, James Stratton, Katherine Wang, and Austin Zheng, as well as participants in the *Digital Economy Lab* seminar series at Stanford University and the *AI Reading Group* at MIT for their valuable feedback and discussion. All errors are our own. **Latest Version:** https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5425555.

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

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

*... You are seeing some effects [of AI on employment]... A particular focus is on **young people** coming out of college. Companies may be able to use AI more than they had in the past ... Hard to say how big it is.*

— Chair Powell, FOMC Press Conference, 17 Sept. 2025

1 Introduction

The impact of generative artificial intelligence (GenAI) 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 GenAI. 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; Ide and Talamas, 2025). If GenAI disproportionately substitutes for entry-level tasks, the lower rungs of these career ladders may be eroding.¹

Recent media reports reinforce these concerns (e.g., The New York Times, 2025b; The Wall Street Journal, 2025b). 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 (The Wall Street Journal, 2025a). 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.”²

Some observers have also linked GenAI adoption to rising unemployment among re-

¹The 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 GenAI disproportionately affects junior positions, it could have lasting consequences for the college wage premium, upward mobility, and income disparities.

²A survey from April 2024 of 804 U.S. hiring managers found that 78% anticipate laying off recent graduates due to AI, with the most vulnerable tasks including research, data entry, email writing, and other routine entry-level assignments (Intelligent, 2024).

cent college graduates (The New York Times, 2025a; The Atlantic, 2025; Forbes, 2025). Since late 2022, unemployment in this group has climbed sharply, even as the overall unemployment rate for young workers has remained stable (Appendix Figure A.1). Others, however, question the importance of GenAI 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 GenAI on the labor market. Specifically, we ask whether GenAI adoption by firms disproportionately affects junior roles relative to more senior positions. This perspective extends the classic literature on skill-biased technological change (e.g., Katz and Murphy, 1992; Autor et al., 2003; Acemoglu and Autor, 2011), which emphasizes shifts in labor demand across education or occupation groups, to a related but distinct dimension: *seniority*.³

Our analysis draws on a dataset that combines LinkedIn résumés and job-posting data from Revelio Labs (Revelio Labs, 2025). 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 almost 200 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.

We identify GenAI adoption by detecting job postings that explicitly recruit “GenAI integrator” roles. The method follows the approach of Hampole et al. (2025) and proceeds in two steps: first flagging postings with GenAI-related keywords, then using a large language model to determine whether the posting reflects a genuine integrator position—one dedicated to implementing or operating GenAI technology in the firm’s workflow. A firm is classified as an adopter if it has posted at least one such vacancy, thereby capturing firms that have actively initiated the integration of GenAI into their operations.⁴

³In a closely related theoretical framework, Ide and Talamas (2025) show that AI’s effects depend on its role and the distribution of knowledge within the organization. Non-autonomous AI that acts as a “co-pilot” tends to complement the least knowledgeable workers—by helping them solve problems—while substituting for more knowledgeable individuals who compete in advisory roles. By contrast, autonomous AI that functions as a *co-worker* substitutes for the least knowledgeable (routine) workers but complements the most knowledgeable, enabling them to scale their expertise.

⁴Our notion of a GenAI integrator is closely related to the “robot integrator” in Acemoglu and Restrepo (2020), which they describe as “companies that install, program, and maintain robots.” Because direct data on robot adoption at the commuting zone level are unavailable, they use the presence of robot inte-

By this definition, 10,599 firms in our sample adopted GenAI during the study period. Although adopters represent only about 3.7 percent of the 285,000 firms in our sample, they are substantially larger on average and account for 17.3 percent of total employment (positions) in our dataset. Our analysis shows that adoption of GenAI was minimal and relatively stable prior to 2023, but accelerated sharply thereafter, with a surge of new firms posting integrator roles following the launch of generative GenAI tools.⁵

We compare adopting and non-adopting firms using a difference-in-differences (DiD) design, tracking junior and senior employment quarterly. From 2015 through 2022, both groups followed parallel trends in junior employment. However, beginning in 2023Q1—coinciding with the sharp increase in GenAI adoption—junior employment in adopting firms declined steeply 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 showed no sign of a break in trend after 2022.

Our main empirical strategy is a triple-difference design that directly evaluates the “seniority-biased” effects. Importantly, this design incorporates firm-by-time fixed effects, which absorb any shocks or trajectories specific to a given firm in a given period, so identification comes solely from differences between juniors and seniors within the same firm and time. In addition, we include industry-by-time-by-seniority fixed effects, which capture shocks that affect juniors and seniors differently at the sector level, ensuring that our estimates are not confounded by industry-wide seniority dynamics. The triple-difference results are consistent with the difference-in-difference estimates. With the exception of a brief dip in early 2021, coefficients remain flat between 2018Q1 and 2022Q4. Starting in 2023Q1, however, they decline sharply, reaching roughly a 10 percent drop after six quarters.

A potential concern is that the observed decline in junior employment may be driven by broader economic shocks occurring around the same time, which are not captured

grators—drawing on data from [Leigh and Kraft \(2018\)](#)—as evidence of greater robot-related activity, and show that commuting zones more exposed to robots also host more integrators.

⁵Our adoption definition captures deliberate integration of GenAI into firm workflows, but does not account for “silent” adoption, such as employees using GenAI tools without firm integration. Therefore, one can think of our definition as conservative: more likely to misclassify adopters as non-adopters (false negatives) than the reverse (false positives). Such misclassification, by assigning treated firms to the control group, is likely to attenuate estimated effects toward zero. At the same time, this form of “silent” adoption may have distinct implications for employment dynamics from the deliberate, organization-level adoption we capture.

by the firm-by-time and industry-by-seniority-by-time fixed effects. We aim to address these concerns in two ways. First, we use an event study to examine employment dynamics around the adoption of GenAI, proxied by the first period with a GenAI integrator posting. This approach helps distinguish adoption effects from broader time-specific shocks by exploiting variation in adoption timing across firms, but it is sensitive to measurement error in the adoption proxy—for example, if firms begin using GenAI before posting for an integrator role. Our results show that junior employment begins to decline roughly three quarters after adoption, reaching an 8 percent reduction after eight quarters. When analyzing all firms collectively, however, the estimates display downward pre-trends, which may reflect firms that adopted GenAI and reduced junior employment before attempting to recruit an integrator. Consistent with this interpretation, excluding the information sector (NAICS 51)—the sector most likely to have adopted GenAI prior to posting—removes the pre-trend while leaving the post-adoption decline virtually unchanged.

Second, we merge the position-level data with occupational exposure measures from [Eloundou et al. \(2024\)](#) using O*NET occupation codes and classify positions into high-exposure (above the 75th percentile) and low-exposure (below the 25th percentile). We then estimate the difference-in-difference analysis separately for juniors in high- and low-exposure occupations. The results show that the number of highly exposed junior roles in adopting firms declined sharply beginning in 2022Q4 relative to non-adopting firms, whereas there was no significant change for juniors in low-exposure occupations. These results indicate that the relative contraction in junior employment in adopting firms is not a broad labor-market phenomenon but is concentrated in occupations most vulnerable to GenAI, consistent with a seniority-biased impact of the technology within firms.

We then turn to investigate the mechanisms behind the decline in junior employment. 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 smaller than the hiring contraction. Moreover, we find that promotions of juniors into more senior roles did not change significantly in adopting firms after early 2023.

We next examine heterogeneity in the effects by educational background and sector. Using a large language model to classify schools into five tiers, we uncover a U-shaped pattern: the steepest declines in junior hiring occur among graduates of solid (Tier 3) and less selective (Tier 4) institutions, while the declines are smaller for juniors from elite (Tier 1), strong (Tier 2), and lowest-tier (Tier 5) schools. At the industry level, the relative decline in junior hiring is evident across sectors and of comparable magnitude, indicating that the effect is not driven by industry-specific dynamics.

We close by acknowledging the limitations of our empirical design. GenAI adoption is not random: adopting firms are systematically larger, more technologically oriented, and more concentrated in highly educated labor. While our empirical strategies address many observable and unobservable differences, we cannot fully rule out alternative explanations for the diverging junior–senior employment patterns. In the absence of a natural experiment, however, we believe our approach provides the most credible available evidence to date that the diffusion of GenAI 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 while the estimated effects for adopting 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” (see [Acemoglu and Restrepo, 2020](#); [Wolf, 2023](#); [Moll and Hanney, 2025](#)).

Finally, while our results suggest that GenAI adoption is associated with reduced junior hiring, this does not necessarily imply immediate task automation. A plausible alternative is that firms are making forward-looking adjustments: the rapid diffusion of GenAI since the end of 2022 may have shifted their expectations, leading them to scale back hiring for roles they predict will be automated in the near future. Firms might find such preemptive adjustments attractive if they view hiring cuts as less costly than future layoffs. Three findings align with this interpretation: (i) the sharp acceleration in GenAI adoption in early 2023, consistent with a sudden expectations shock following the introduction of GPT-3.5; (ii) the steep decline in junior employment that begins shortly thereafter; and (iii) the contraction is driven primarily by slower hiring rather than increased

separations. Importantly, under this view, if firms have overreacted to the introduction of AI, the resulting decline in junior employment may be temporary. Nonetheless, our data do not allow us to directly test this mechanism.

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

Skill-Biased Technological Change: 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 GenAI is a “*seniority-biased*” technological change, disproportionately affecting juniors who typically perform simpler, more routinized tasks even in high-skill fields.

Experimental Evidence on GenAI and Productivity: Since 2023, a rapidly growing empirical literature has examined GenAI’s labor-market effects. Experimental studies generally find that GenAI 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 GenAI 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. Moreover, Cui et al. (2025) find especially high productivity gains and adoption rates for less experienced software developers. Related field-experimental evidence

in Dell’Acqua et al. (2025) shows that adding a GenAI 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 GenAI can act as a “leveler,” narrowing productivity gaps between less and more experienced workers (Autor, 2024).

Employment Dynamics by GenAI Exposure: A second strand, closer to our study, uses broad-economy data to trace employment trajectories of occupations and industries by GenAI exposure, yielding mixed results. In a recent paper, Brynjolfsson et al. (2025a) show that since the late-2022 debut of GenAI, 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 GenAI 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 a more direct measure of firm-level adoption, identified through postings for explicit “GenAI integrator” roles. This design enables us to study realized adoption decisions and their within-firm consequences for junior versus senior employment. Moreover, unlike exposure-based measures, our adoption measure varies over time, allowing us to exploit variation in adoption timing with an event-study design.

Evidence from Firm-Level AI Adoption: A third strand, closest to our methodology, consists of studies that examine the implications of firm-level AI adoption. 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 similar data and NLP methods to develop measures of firm-level AI adoption and task-level exposure, which we closely follow in our analysis. They infer adoption from résumé text, using a large language model to extract in-house AI applications and map them to O*NET tasks via sentence embeddings, thereby identifying tasks exposed to AI. 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. Finally, Humlum and Vestergaard (2025) provide complementary evidence from Denmark, linking large-scale worker surveys on chatbot adoption to matched employer–employee data. Despite rapid adoption and firm investments, they do not find effects of AI chatbots on earnings and hours, highlighting that realized labor-market impacts in Denmark remain minimal two years after ChatGPT’s launch.⁶

These studies highlight that pre-2023 adoption of AI technology often entailed internal reallocation rather than aggregate job loss. In contrast, our paper provides U.S. evidence on the implications of 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 GenAI adoption reduces junior employment while leaving senior employment unaffected.

⁶Chen and Stratton (2025) analyze firm-level adoption of GitHub Copilot and Cursor using detailed engineering workflow data, documenting effects on productivity, task allocation, and organizational collaboration.

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 through WRDS. 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.⁷

A key feature of the dataset is the 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.⁸ In the analysis that follows, we group positions into two broad categories: *juniors* (Entry and Junior) and *seniors* (Associate and above). Appendix A.3 provides further detail and validation of this classification.

We complement the worker résumé data with Revelio’s database of job postings, which tracks recruitment activity by the firms since September 2021. Each posting contains the firm identifier, posting date, and the raw text description. We use these raw descriptions to construct our measure of firm-level GenAI adoption, as described in detail in Section 3.3.

In addition, although not the primary focus of the paper, we also incorporate the occupational exposure measures from Eloundou et al. (2024), merged with our position-level data via O*NET SOC codes. We rely on the GPT-4-based beta exposure measure (as in, e.g., Brynjolfsson et al., 2025a), and classify all the positions into three categories: low

⁷Roughly 5 percent of positions are “contained positions,” meaning that another position for the same individual in the same firm fully overlaps their reported work period. We treat these cases as follows: if the container has lower seniority than the contained position, we shorten the container’s end date to the contained position’s start date, treating the latter as a promotion. If the container has higher seniority, we drop the contained position.

⁸More details on Revelio’s seniority classification methodology are available at <https://www.data-dictionary.reveliolabs.com/methodology.html#seniority>.

exposure (0–25th percentile), medium exposure (25–75th percentile), and high exposure (75–100th percentile).

Finally, we construct a school-quality measure at the position level. For each position, we assign the institution where the worker most recently studied, provided that the education ended no later than one year after the job start. If no such record exists, we fall back on the individual’s first recorded education, provided it began before the job start date. We then merge these institutions with a GPT-4–based quality rating, in which each school is assigned a score from 1 to 5.⁹

Our final sample merges all U.S. positions in the Revelio Labs dataset with job postings at the firm level. The resulting dataset covers 284,974 firms that were successfully matched to both employee position data and job postings, and that were actively hiring between January 2021 and March 2025.¹⁰ For these firms, we observe 156,765,776 positions dating back to 2015 and 198,773,384 job postings since 2021, all with usable raw text descriptions.¹¹

3.2 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

⁹These ratings were produced using OpenAI’s GPT-4o-mini model, prompted to act as an academic evaluator (Appendix A.6.2 provides the full prompt used to generate these scores). A score of 1 corresponds to Ivy League and other globally elite institutions; 2 to highly respected international institutions; 3 to strong national or regional institutions; 4 to lower-tier but standard institutions; and 5 to weak or diploma-mill-type institutions.

¹⁰We define a firm as active if it recorded at least 20 new hires over this period.

¹¹We exclude the top 1 percent of firms with the highest postings-to-hires ratios. Manual inspection indicates that many of these firms are HR intermediaries recruiting on behalf of other employers.

who start a new position at the firm after previously holding a lower-seniority role within the same firm.

Figure 1 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 zero in January 2015.




Figure 1: Log Average Employment of Junior and Senior in Sample Firms

Notes: This figure plots the percentage change in the average number of junior- and senior-level workers across firms in our sample over time. Values are normalized to zero in January 2015, immediately after the launch of ChatGPT. “Junior” refers to Entry- and Junior-level positions, while “Senior” refers to Associate level and above (see Section 3.1 for details).

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 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 aligns 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.

3.3 GenAI Adoption

3.3.1 GenAI-Integrator Vacancies

We identify firms that adopted GenAI by detecting job postings that explicitly seek workers to integrate GenAI technologies into the organization. Inspired by the approach of [Hampole et al. \(2025\)](#), we proceed in two steps. First, we compile a list of GenAI-related keywords and flag all postings containing at least one of them.¹² 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 “GenAI integrator” postings—i.e., those reflecting an active attempt to recruit workers tasked with adopting or implementing GenAI in the firm’s workflows—from false positives (appendix [A.6.1](#) provides the exact prompt used). This procedure identifies 131,845 postings, corresponding to 0.066 percent of the full corpus, as GenAI integrator vacancies. A graphical overview of this procedure is shown in Appendix [A.7](#).

Below we present two illustrative job postings: one that our LLM classifier identifies as a *GenAI integrator* (green box) and one that it does not (red box). Both postings pass the initial screening step due to the presence of the keyword “*generative AI*.”

The first example (green box) shows a correct classification, as the role explicitly includes the responsibility to “**integrate AI models** into existing systems and applications,” and the job title—*Generative AI Developer Consultant*—closely fits our notion of a GenAI integrator. The second example highlights the value of an LLM-based classifier that goes beyond simple keyword search. Although the posting is from a Generative AI company, it describes a standard *customer support* role unrelated to integrating AI into workflows.

¹²Keywords 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.

The model correctly classifies it as a non-integrator position. See Appendix A.2 for additional examples.¹³

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.

...

Role: Customer Service Representative (HireQuotient)

Summary: HireQuotient is a pioneering company in the Software Development industry, transforming recruitment processes through Generative AI and Skill Intelligence. The position is a Mid-Level customer service representative.... focused on ensuring high-quality support and satisfaction for customers ...

Responsibilities:

- Manage customer inquiries, complaints, and feedback through various channels, ensuring a high level of satisfaction.
- Provide proactive support through live chat, email, and phone.
- Remain informed about product updates and company policies to deliver accurate information.

...

¹³As shown in Appendix A.6.1, 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.

3.3.2 GenAI-Adopting Firms

We define a firm as a GenAI adopter if it has posted at least one GenAI integrator vacancy. By this criterion, 10,599 firms qualify as adopters. While they make up only 3.72 percent of the 284,974 firms in our sample, adopters are disproportionately large (see more details below) and account for 17.3 percent of the employment (positions) in our dataset.

Figure 2 plots the timing of GenAI adoption, defined as the posting date of each firm’s first GenAI integrator vacancy. Prior to 2023, adoption was minimal and stable, with roughly 30 new adopters per month. Beginning in early 2023—shortly after the release of GPT-3.5—the number of new adopters rose sharply, peaking at 456 in August 2023. Adoption then stabilized at around 400 firms per month through the end of 2024 before accelerating again in early 2025, reaching 574 new firms in March. By the end of the sample period, the cumulative number of adopters had surpassed 10,000.




Figure 2: Adoption Distribution Over Time

Notes: Panel (a) shows the monthly number of firms posting their first GenAI-integrator vacancy, while Panel (b) reports the cumulative total, covering the period from September 2021 to March 2025.

Next, Table 1 presents descriptive statistics for the full sample, adopters, and non-adopters. Several systematic differences stand out. Adopting firms are much larger, averaging over 500 employees (median: 82) compared to roughly 100 (median: 33) for non-adopters. Their workforces are more senior-heavy, with juniors comprising only 42 percent of employment versus 55 percent in non-adopting firms. Consistent with this pat-

tern, adopters exhibit substantially higher hiring and separation volumes, with a smaller share involving junior positions. As expected, adopters also employ a significantly larger share of workers in highly exposed occupations. In addition, workers in adopting firms are more likely to have graduated from higher-quality institutions.

While non-adopters are widely dispersed across sectors, GenAI adopters are heavily concentrated in professional services (28 percent) and information (24 percent), both knowledge-intensive industries where GenAI adoption is most salient. Geographically, adopters are disproportionately headquartered in California (20 percent versus 14 percent overall), while being slightly less represented in Texas compared to non-adopters. Appendices A.4 and A.5 provide further detail on adopters' distribution across industries and states.

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

4 Results

4.1 Employment Dynamics by Adoption

4.1.1 Comparing Employment of Adopting vs. Non-Adopting Firms

We begin by comparing the evolution of junior and senior employment over time in adopting versus non-adopting firms. Figure 3 plots the average employment of juniors (Panel 3a) and seniors (Panel 3b). From 2018 through the end of 2022, junior employment followed similar trajectories across both groups. Starting in late 2022, however, the paths diverged: junior employment in adopting firms declined markedly, while non-adopters remained relatively stable. By contrast, senior employment in adopting firms began to grow slightly faster than in non-adopters from early 2022 onward, with no evident trend

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

Variable	All Firms	Non-Adopters	GenAI Adopters
Panel A. Workforce Composition and Characteristics			
Firm size (average)	122.9	107.2	542.6
Firm size (median)	33.7	33.0	82.3
	(438)	(370)	(1219)
Share junior employees (Entry/Junior)	0.548	0.553	0.422
	(0.228)	(0.227)	(0.211)
Share senior employees (Associate+)	0.454	0.450	0.581
	(0.228)	(0.227)	(0.211)
Average number of new hires (per quarter)	6.9	6.0	29.6
	(24.6)	(21.2)	(64.8)
Share of new hires junior	0.670	0.677	0.508
	(0.351)	(0.351)	(0.307)
Average number of separations (per quarter)	5.5	4.9	23.2
	(21.7)	(19.0)	(55.9)
Share of separations junior	0.684	0.691	0.528
	(0.352)	(0.352)	(0.311)
Average number of promotions for juniors (per quarter)	0.489	0.404	2.8
	(2.1)	(1.6)	(6.8)
Juniors in high-exposure jobs (% of all juniors)	0.274	0.267	0.443
	(0.248)	(0.245)	(0.260)
Juniors in low-exposure jobs (% of all juniors)	0.281	0.287	0.121
	(0.270)	(0.271)	(0.168)
Seniors in high-exposure jobs (% of all seniors)	0.187	0.181	0.332
	(0.192)	(0.190)	(0.210)
Seniors in low-exposure jobs (% of all seniors)	0.153	0.156	0.064
	(0.191)	(0.193)	(0.102)
Juniors from college tier-1—highest (% of all juniors)	0.052	0.051	0.098
	(0.126)	(0.124)	(0.169)
Juniors from college tier-2 (% of all juniors)	0.169	0.168	0.213
	(0.192)	(0.192)	(0.193)
Juniors from college tier-3 (% of all juniors)	0.341	0.341	0.334
	(0.228)	(0.229)	(0.213)
Juniors from college tier-4 (% of all juniors)	0.267	0.269	0.223
	(0.213)	(0.213)	(0.191)
Juniors from college tier-5—lowest (% of all juniors)	0.171	0.172	0.132
	(0.187)	(0.188)	(0.159)
Panel B. Industry and Headquarters Location			
Share in NAICS sector 51 (Information)	0.070	0.064	0.240
Share in NAICS sector 52 (Finance and Insurance)	0.067	0.067	0.087
Share in NAICS sector 54 (Professional Services)	0.156	0.152	0.279
Share in NAICS sector 5 (Other)	0.066	0.065	0.093
Share in non-NAICS 5 sectors	0.640	0.653	0.302
HQ in California	0.137	0.135	0.199
HQ in Texas	0.074	0.075	0.060
HQ in New York	0.080	0.079	0.090
HQ in Other States	0.716	0.718	0.654
Observations	11,021,214	10,622,695	398,519
Number of firms	284,756	274,168	10,588

Notes: The table reports averages (unless otherwise indicated) of the main variables across firm-by-quarter observations from 2015Q1 to 2025Q1, separately for the full sample, GenAI 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.

break around 2023.¹⁴

To place these descriptive patterns in a more formal framework, we estimate a difference-in-differences (DiD) specification, comparing employment in adopting and non-adopting firms.¹⁵ Specifically, we estimate this specification separately for junior and senior workers:

$$\log(\text{Employment}_{it}) = \alpha + \sum_{j=2015Q2}^{2025Q1} \beta_j \mathbf{1}\{t = j\} \times \text{Adopt}_i + \delta_t + \text{Adopt}_i + \varepsilon_{it}, \quad (1)$$

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 GenAI.¹⁶ 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.

Panel 5a 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.

¹⁴These patterns are consistent with Brynjolfsson et al. (2025a), who document a similar divergence by occupational exposure among younger workers, with no corresponding shift for older workers.

¹⁵For computational efficiency we aggregate our panel data to quarterly frequency in all analyses.

¹⁶Note that Adopt_i is time-invariant: a firm is defined as an adopter if it posted at least one GenAI integrator vacancy at any point during the sample (see Section 3.3 for more details). In Section 4.1.3, we relax this definition by exploiting variation in the timing of adoption across firms using an event-study design.




Figure 3: Average Employment by Adoption Status

Notes: Panels (a) and (b) shows the average number of workers for junior-level and senior-level employees, respectively, in adopting versus non-adopting firms (as percentage change relative to December 2022, immediately following the launch of ChatGPT). Panel (c) presents the estimated coefficients β_j from Equation 1, estimated separately for juniors and seniors.

4.1.2 Triple-Difference

Our main empirical strategy is a triple-difference design that directly evaluates the “seniority-biased” effects. Specifically, we estimate the following specification:

$$\begin{aligned} \log(\text{Employment}_{ist}) = & \sum_{j=2015Q2}^{2025Q1} \beta_j \mathbf{1}\{t = j\} \times \text{Adopt}_i \times \text{Junior}_s \\ & + \sum_{j=2015Q2}^{2025Q1} \pi_j \mathbf{1}\{t = j\} \times \text{Adopt}_i + \sum_{j=2015Q2}^{2025Q1} \rho_j \mathbf{1}\{t = j\} \times \text{Junior}_s \\ & + \kappa (\text{Adopt}_i \times \text{Junior}_s) + \gamma_{it} + \xi_{p(i)st} + \varepsilon_{ist}, \end{aligned} \quad (2)$$

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 GenAI (see Section 3.3 for the definition), and Junior_s is an indicator equal to one for juniors and zero for seniors. $p(i)$ denotes the sector (NAICS 2 digit) of the firm i .

The coefficients β_j trace a triple-differences event-time profile: they capture how junior employment evolves relative to senior employment *within the same firm and period*, comparing adopters to non-adopters. Firm-by-time fixed effects γ_{it} absorb firm-specific shocks in a given period, ensuring identification comes from the within-firm-time junior-senior contrast. $\xi_{p(i)st}$ are industry-by-seniority-by-time fixed effects. These absorb broad trends in junior and senior employment at the sector level over time, ensuring that identification comes from within-industry deviations across firms. Intuitively, they remove the possibility that our results are driven by sector-wide shifts in junior versus senior employment unrelated to adoption. See section 4.1.3 for a detailed discussion on the importance of this variable in our estimation. 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 displays the estimated coefficients β_j from Equation 2. Estimation begins in 2018Q1 due to computational constraints. Aside from a brief dip in early 2021, the

coefficients are essentially flat through 2022Q4. Starting in 2023Q1, however, the coefficients decline sharply, reaching roughly a 10 percent drop after six quarters. This break—coinciding with the rapid diffusion of GenAI—provides suggestive evidence that adoption is associated with seniority-biased employment patterns, reducing junior employment relative to senior employment within firms.¹⁷




Figure 4: Triple Differences

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

The sharp and relatively early decline in the DiD and triple-difference coefficients—beginning in 2023Q1, only shortly after the release of GPT-3.5—may seem surprising, as one might expect the automation effects of GenAI to unfold more gradually. However, the rapid surge in adoption from early 2023 onward (see Figure 2) suggests that firms perceived GPT-3.5 as a discrete shock. In response, they may have adjusted in a forward-looking manner, reducing junior roles they anticipated would be automated in the near future. This interpretation is consistent with the evidence in Section 4.2.1, where we show that the decline in junior employment is driven by reduced hiring rather than increased

¹⁷Appendix A.8 reports triple-difference estimates excluding the industry-by-seniority-by-time fixed effects. The results are very similar to those in Figure 4.

separations. If firms view hiring freezes as less costly than subsequent layoffs, they may prefer to adjust early in response to revised expectations about automation following the introduction of new GenAI tools. That said, our data cannot directly test this mechanism, so we offer it only as a plausible interpretation.

4.1.3 Addressing Threats to Identification

A potential concern with our findings is that the observed decline in junior employment could be driven by broader economic shocks occurring around the same time. While the long and relatively flat pre-trends, combined with rich set of controls—most notably firm-by-time and industry-by-seniority-by-time fixed effects—help mitigate this concern, they cannot fully eliminate it. We therefore implement several complementary checks.

Event-Study Evidence Around Adoption: Our first approach is an event study that traces junior employment dynamics around the timing of GenAI adoption, proxied by the first period in which a firm posts a GenAI integrator vacancy. This design helps separate adoption effects from broader time-specific shocks by exploiting variation in adoption timing across firms. However, it is sensitive to measurement error in the adoption proxy—for example, if firms begin using GenAI before posting for an integrator role. Specifically, we estimate:

$$\log(\text{JuniorEmployment}_{it}) = \sum_{j=2}^J \beta_j (\text{Lag}_j)_{it} + \sum_{k=1}^K \gamma_k (\text{Lead}_k)_{it} + \mu_i + \lambda_t + \varepsilon_{it}, \quad (3)$$

where $\log(\text{JuniorEmployment}_{it})$ denotes the log number of junior workers at firm i at time t ; $(\text{Lag}_j)_{it}$ is an indicator equal to one if the current period t is j periods before adoption; and $(\text{Lead}_k)_{it}$ is defined analogously for periods after adoption. μ_i and λ_t are firm and time fixed effects, and ε_{it} is an error term.

Panel (a) of Figure 5 reports the results for all firms in our sample. Junior employment begins to decline roughly three quarters after adoption, reaching an 8 percent reduction after eight quarters. However, the estimates also show downward pre-trends, which may reflect firms that adopted GenAI and reduced junior employment before formally posting


an integrator vacancy. Consistent with this interpretation, Panel (b) shows that excluding the Information sector (NAICS 51)—the sector most likely to adopt before posting—largely removes the pre-trends, while leaving the post-adoption decline essentially unchanged. In addition, Appendix A.9 reports estimates of this specification separately by firm size. The results confirm that our findings are not driven by comparisons between large adopter firms and small non-adopters, which could otherwise reflect differential dynamics or heterogeneous treatment effects across the firm size distribution.

DiD by Occupational Exposure We next re-estimate the DiD specification (Equation 1) separately for junior positions with high and low exposure to GenAI (see Section 3.1 for details on the exposure measure).¹⁸ Figure 6 reports the results. For high-exposed occupations, the coefficients increase from 2015 through 2022Q3, indicating that adopting firms were expanding junior employment in these roles relative to non-adopters. However, beginning in 2022Q4, the trend reverses sharply, with high-exposure junior employment in adopters declining significantly relative to non-adopters. In contrast, coefficients for low-exposure junior occupations decline gradually between 2015 and 2019 and then stabilize, with no evident break around late 2022. Taken together, these patterns suggest that the post-2022 decline in junior employment we documented is concentrated in high-exposure jobs. The relative contraction is therefore not a broad labor-market phenomenon but rather reflects reductions in roles most vulnerable to GenAI in the adopting firms, consistent with a seniority-biased impact of this technology within firms.


Discussion on Specific Potential Confounders: We discuss two potential confounding shocks: the rise in interest rates beginning in March 2022, and the tech sector’s overhiring during 2021–2022 followed by its subsequent correction. We provide additional evidence—beyond the event study and the DiD by occupational exposure—indicating that these factors do not drive our findings.

First, one might worry that adopters are inherently more sensitive to monetary policy cycles and therefore responded more strongly to the interest rate hikes of 2022–2023. Although our firm-by-time and industry-by-seniority-by-time fixed effects should absorb

¹⁸Appendix A.11 lists the most common high- and low-exposure occupations by industry (by number of positions).



(a) All Firms



(b) Without Information Sector (NAICS 51)

Figure 5: Event Study

Notes: The graph presents the estimated coefficients β_j from Equation 3 using the method of Callaway and Sant'Anna (2021) for staggered adoption. Panel (a) shows the results for the entire sample, while Panel (b) excludes the Information sector (NAICS 51). Standard errors are clustered in firm level.




Figure 6: DiD Estimates for High- and Low-Exposure Occupations (Juniors)

Notes: This figure reports separate difference-in-differences estimates (β_j from Equation 1) for juniors in high- and low-exposure occupations. Standard errors are clustered in firm level.

such shocks, it is worth considering whether they could still bias the results. Several pieces of evidence suggest this is unlikely. Most importantly, as shown in Sections 4.1.1 and 4.1.2, the decline is concentrated exclusively among junior workers. While younger workers are typically more cyclical, it is implausible that monetary tightening would affect only juniors without also impacting seniors. In addition, our pre-trends extend back to 2015, covering the 2015–2018 tightening cycle. As Figure A.11 shows, credit and financial markets tightened considerably in this period. Yet Figure 5a reveals no relative decline in junior employment among adopters during those years. Finally, because adoption is positively correlated with firm size (Section 3.3.2), and larger firms tend to be less sensitive to interest rate shocks (see Chodorow-Reich, 2014; Gertler and Gilchrist, 1994), it is unlikely that interest rate exposure explains our results.

Second, one might be concerned about the post-Covid hiring boom in the tech sector (U.S. Bureau of Labor Statistics, 2025). If this boom led to a subsequent correction, it could disproportionately affect adopters given their higher representation in the Infor-

mation sector (see Section 3.3.2). However, these patterns are likely to be absorbed by the industry-by-seniority-by-time fixed effects in our main specifications. Moreover, if boom–bust dynamics affected our results, we would expect to see a post-Covid relative increase in junior employment among adopters before the decline in 2023Q1, which we do not observe. Finally, as shown in Figure 5, the decline in junior employment following GenAI adoption is essentially unchanged when the Information sector is excluded, reinforcing that the results are not driven by this sector.

In sum, while we cannot completely rule out confounding shocks in the absence of a natural experiment, the event-study analysis, the DiD by occupational exposure, and the additional evidence discussed here—together with the rich set of fixed effects in our main analysis—provide strong suggestive evidence that the observed decline in junior employment is indeed associated with GenAI adoption.

4.2 Decomposing Decline in Junior Employment

4.2.1 Hires, Exits, and Promotions

A decline in the headcount of junior workers can result from any of the following channels: (i) a decrease in junior hiring, (ii) an increase in separations or layoffs of juniors from adopting firms, or (iii) an increase in promotions of juniors to senior positions. Our online résumé data can be thought of as a detailed matched employer-employee dataset, allowing us to closely track all of these flows. To decompose these dynamics, we estimate separate difference-in-differences regressions of the following form:

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

where y_{it} denotes the number of hires, number of separations, probability of promotion,¹⁹ or net changes in firm i at time t . γ_i denotes firm fixed effects, δ_t are time fixed effects and ξ_{pt} are industry-by-time fixed effects.

Table 2 reports the estimated coefficients on the interaction term from Equation 4. The results indicate that the sharp contraction in junior employment among adopters is driven

¹⁹Probability of promotion is defined as $\frac{\text{Promotions}_t}{\text{Junior Employment}_{t-1}}$.

Table 2: Hiring, Separations, Promotions, and Net Employment

	Hiring	Separation	Promotion	Total Change
<i>Panel A: Juniors</i>				
Treat \times Post	−5.029*** (0.225)	−1.781*** (0.166)	0.018 (0.015)	−3.721*** (0.149)
<i>Panel B: Seniors</i>				
Treat \times Post	−0.256* (0.140)	1.484*** (0.108)	–	−1.255*** (0.146)
Observations	8,027,376	8,027,376	7,772,524	7,998,378
Clusters (firms)			284,500	

Notes: This table reports the estimated β from Equation 4. Standard errors clustered by firm in parentheses. We control for firm and time fixed effects in all columns.

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

primarily by a slowdown in hiring, rather than by increased exits. Specifically, the coefficient on *Hiring* implies that, relative to non-adopters, GenAI-adopting firms hired on average 5.0 fewer junior workers per quarter after 2023Q1. Interestingly, separation rates for juniors also fell among adopters, though the magnitude of this decline is smaller than the reduction in hiring.²⁰ The relative number of junior promotions appears unchanged following 2023Q1. For senior employees, by contrast, hiring shows little change, while separations rise modestly, leading to a small net decline in senior headcount.

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 curbing junior recruitment once GenAI tools become available, while maintaining their existing workforce.

4.2.2 Heterogeneity

Heterogeneity by Human Capital We next examine heterogeneity in the decline of junior employment by workers' educational background. To capture educational quality, we use a school-prestige measure described in Section 3. Specifically, we re-estimate Equation 4 for junior hires separately for each of the five school-quality tiers.

²⁰See Appendix A.12 for the time-series difference-in-differences results for junior hires and exits.




Figure 7: Results by School Quality

Notes: The figure presents the estimates of Equation 4 for junior hires, run separately by university prestige category. The coefficients represent post-adoption changes in junior employment for GenAI adopters relative to non-adopters. Standard errors are clustered at the firm level.

The results, shown in Figure 7, reveal a pronounced U-shaped pattern. Juniors from tier-3 and tier-4 universities experienced the steepest relative declines in employment, while juniors from tiers 1, 2 and 5 also saw reductions, but of smaller magnitude.

Appendix A.13 provides additional context by plotting average salaries and AI-exposure levels of junior positions across school-quality tiers. As expected, Figure A.13 shows a monotonic positive relationship between juniors' salaries and the prestige of their alma mater. Interestingly, AI exposure also rises monotonically with school quality. This pattern implies that the stronger declines observed for tiers 3 and 4 in Figure 7 cannot be attributed to differences in exposure levels, suggesting instead that GenAI adoption may be reshaping demand unevenly across the human capital distribution.

Heterogeneity by Sector Finally, we examine heterogeneity in the effects by sector. For this, we re-estimate Equation 4 for junior hires separately by sector. Results are presented in Figure 8. Across all sectors, adopting firms exhibit a sharp and statistically significant relative decline in junior hiring after 2023Q1, while senior hiring remains stable or increases slightly. This pattern indicates that the contraction in junior hiring is broad-

based across industries and not driven by any single sector disproportionately reducing demand for junior workers.




Figure 8: 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-2022 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.

5 Conclusion

This paper provides early, broad-based evidence that the diffusion of generative artificial intelligence (GenAI) since 2023 is associated with *seniority-biased* employment effects within firms. Using résumé-posting data linked to nearly 285,000 U.S. firms and a direct measure of adoption based on “GenAI 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 and industry-by-seniority-by-time fixed effects confirms that the within-firm gap

between junior and senior employment widens precisely as GenAI diffuses. Consistent with these results, a staggered difference-in-differences design exploiting variation in adoption timing across firms shows that the contraction in junior employment emerges after GenAI adoption by firms. Moreover, we find that the decline in junior employment among adopting firms is almost entirely driven by occupations with high exposure to GenAI, while low-exposure occupations show no comparable decline.

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. Examining heterogeneity in the junior hiring decline by educational background, we reveal a U-shaped pattern: graduates from mid-tier institutions are most affected, while those from strong and lowest-tier schools experience smaller reductions. Moreover, we find that the decline in junior hiring among GenAI adopters is broad-based and not limited to the information technology sector.

These findings should be interpreted with caution. GenAI adoption is not random, and adopting firms differ systematically in size, workforce composition, and industry. Although the triple-difference design accounts for firm-specific and seniority-specific shocks, unobserved confounding factors may remain. Our adoption measure, based on integrator postings, captures deliberate organizational uptake but may miss informal or “silent” adoption within firms. Moreover, the analysis covers a relatively short period (2023–2025); longer-run adjustments in training, task allocation, and internal career ladders could either attenuate or amplify these initial effects.

Even with these caveats, the results point to important implications. GenAI adoption appears to shift work away from entry-level tasks, narrowing the bottom rungs of internal career ladders. Because early-career jobs are central to lifetime wage growth and mobility, such shifts may have lasting consequences for inequality and the college wage premium. Taken together, our evidence suggests that GenAI diffusion constitutes a form of seniority-biased technological change, with far-reaching implications for how careers begin, how firms cultivate talent, and how the gains from 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. (2020). Robots and Jobs: Evidence from U.S. Labor Markets. *Journal of Political Economy*, 128(6):2188–2244.
- 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., 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.
- 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.
- 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.
- Brave, S. A. and Butters, R. A. (2012). Diagnosing the financial system: Financial conditions and financial stress. *International Journal of Central Banking*, 8(2):191–239.
- 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.
- Callaway, B. and Sant’Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.
- Chandar, B. (2025). Tracking Employment Changes in AI-Exposed Jobs. *Available at SSRN* 5384519.
- Chen, F. and Stratton, J. (2025). Generative AI and Organizational Structure. Unpublished manuscript.
- Chodorow-Reich, G. (2014). The employment effects of credit market disruptions: Firm-level evidence from the 2008–09 financial crisis. *Quarterly Journal of Economics*, 129(1):1–59.
- Cui, Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S., and Salz, T. (2025). The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers. *Available at SSRN*. <https://ssrn.com/abstract=4945566> or <http://dx.doi.org/10.2139/ssrn.4945566>.
- 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 Management 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.
- Federal Reserve Bank of Chicago (2023). National Financial Conditions Index: Frequently Asked Questions. <https://www.chicagofed.org/research/data/nfci/about>. Accessed October 2025.
- 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.
- Forbes (2025). As AI Reduces New Grad Hiring, Apprenticeships Will Become Essential. Accessed: 2025-09-25.
- Garicano, L. (2000). Hierarchies and the Organization of Knowledge in Production. *Journal of Political Economy*, 108(5):874–904.
- Gertler, M. and Gilchrist, S. (1994). Monetary policy, business cycles, and the behavior of small manufacturing firms. *Quarterly Journal of Economics*, 109(2):309–340.
- Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720.
- 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.
- Hatzius, J., Hooper, P., Mishkin, F. S., Schoenholtz, K. L., and Watson, M. W. (2010). Financial conditions indexes: A fresh look after the financial crisis. *NBER Working Paper No. 16150*.
- Humlum, A. and Vestergaard, E. (2025). Large Language Models, Small Labor Market Effects. Technical report, National Bureau of Economic Research.

Ide, E. and Talamas, E. (2025). Artificial intelligence in the knowledge economy. *Journal of Political Economy*. Forthcoming.

Intelligent (2024). 8 in 10 Companies Plan to Layoff Recent College Grads This Year Due to AI. <https://www.intelligent.com/8-in-10-companies-plan-to-layoff-recent-college-grads-this-year-due-to-ai/>. Survey of 804 U.S. hiring managers conducted via Pollfish, April 2024.

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

Leigh, N. G. and Kraft, B. R. (2018). Emerging Robotic Regions in the United States: Insights for Regional Economic Evolution. *Regional Studies*, 52(6):804–815.

Moll, B. and Hanney, O. (2025). The ‘Missing Intercept’ Problem with Going from Micro to Macro. <https://voxdev.org/topic/methods-measurement/missing-intercept-problem-going-micro-macro>. VoxDev. Accessed: 2025-09-28.

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.

Revelio Labs (2025). Revelio Labs Workforce Data. Accessed via Wharton Research Data Services (WRDS).

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 (2025a). AI Is Wrecking an Already Fragile Job Market for College Graduates. The Wall Street Journal. Accessed September 25, 2025.
- The Wall Street Journal (2025b). The ‘Great Hesitation’ That’s Making It Harder to Get a Tech Job. Accessed: 2025-06-20.
- U.S. Bureau of Labor Statistics (2025). All Employees, Information [CEU5000000001]. Retrieved from FRED, Federal Reserve Bank of St. Louis. <https://fred.stlouisfed.org/series/CEU5000000001>.
- Wolf, C. K. (2023). The Missing Intercept: A Demand Equivalence Approach. *American Economic Review*, 113(8):2232–2269.

A Appendix

A.1 College Graduates Unemployment

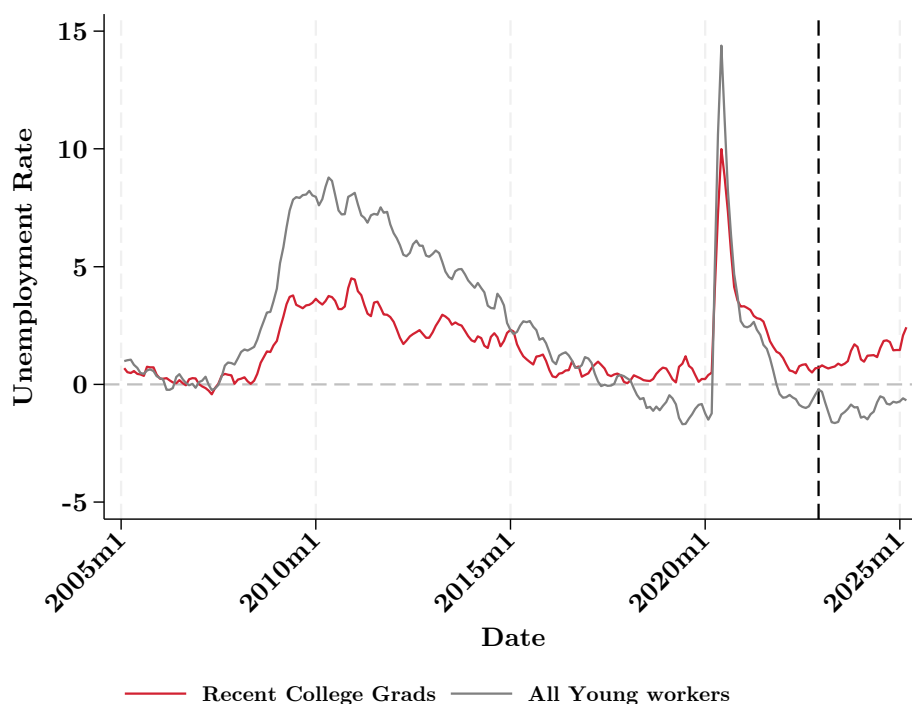


Figure A.1: Unemployment Rates for Recent College Graduates vs. All Young Workers

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 Additional Job Postings Examples

The green box presents another example correctly classified by the LLM as a *GenAI integrator*, while the red box shows a posting that, despite containing the related keyword *Large Language Model*, is not about integrating GenAI into workflows.

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.

Role: Senior Security Engineer (Offensive Security, BytePlus)

Summary: The team builds infrastructures, platforms, and technologies to protect users, products, and systems. You will contribute to key security initiatives, developing scalable and secure-by-design solutions.

Responsibilities:

- Responsible for risk discovery and penetration testing of cloud products and infrastructure.
- Conduct risk analysis and threat modeling; provide systematic solutions to business lines.
- Research cutting-edge technologies including cloud-native, microservices, zero trust, big data, and **large language models**.
- Support the development of secure business technologies and architectures.

...

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.2–A.5 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.

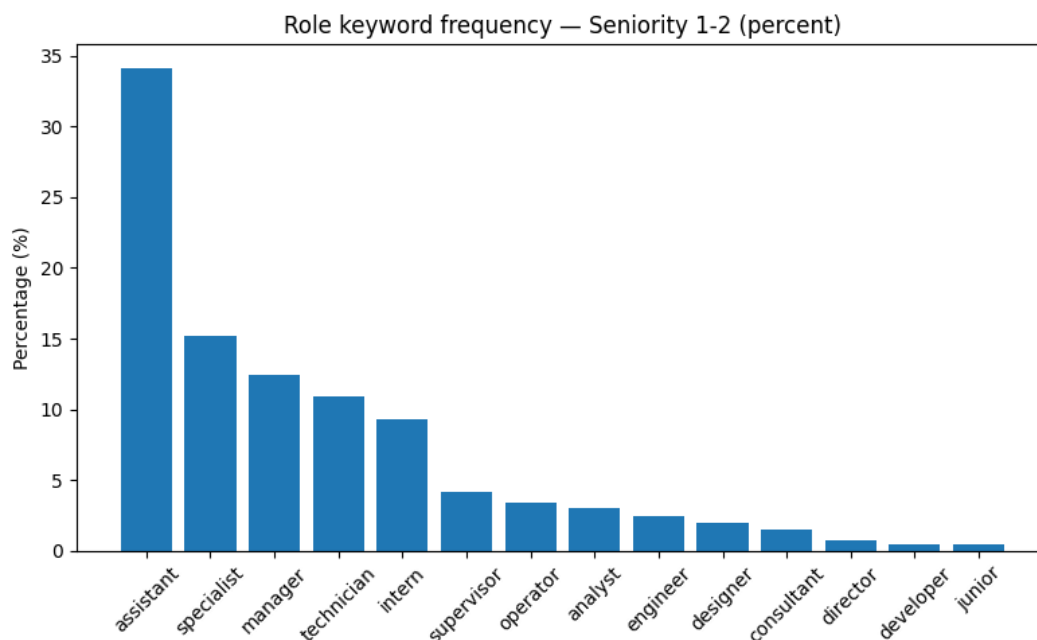


Figure A.2: Role keyword frequency — Seniority 1 and 2 (percent).




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




Figure A.4: Role keywords — Seniority 1-2 (percent share).

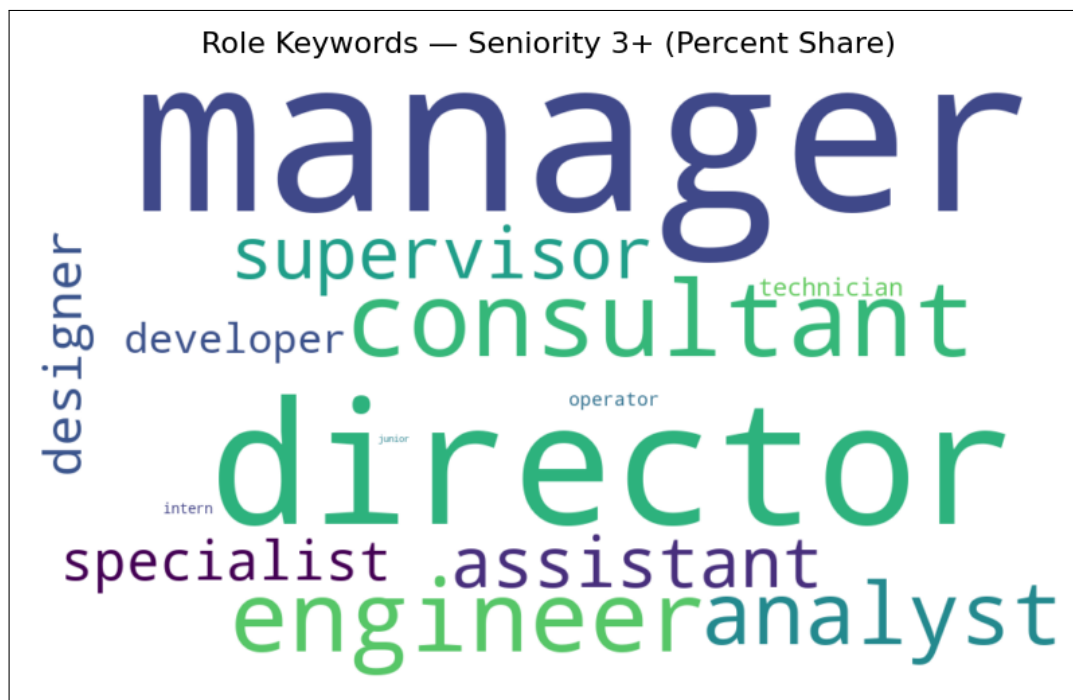


Figure A.5: Role keywords — Seniority 3+ (percent share).

A.4 Sectoral Distribution of AI Adopters

In this appendix, we provide descriptive evidence on the sectoral distribution of AI adoption. Figure A.6 documents the share of firms in each major sector that have adopted AI, while Figure A.7 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.6: Share of Adopting Firms by Industry

Notes: The figure reports the share adopting firms in each 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.7: Distribution of Adopters Across Sectors

Notes: The figure 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 Geography of Adopters

Figure A.8 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.

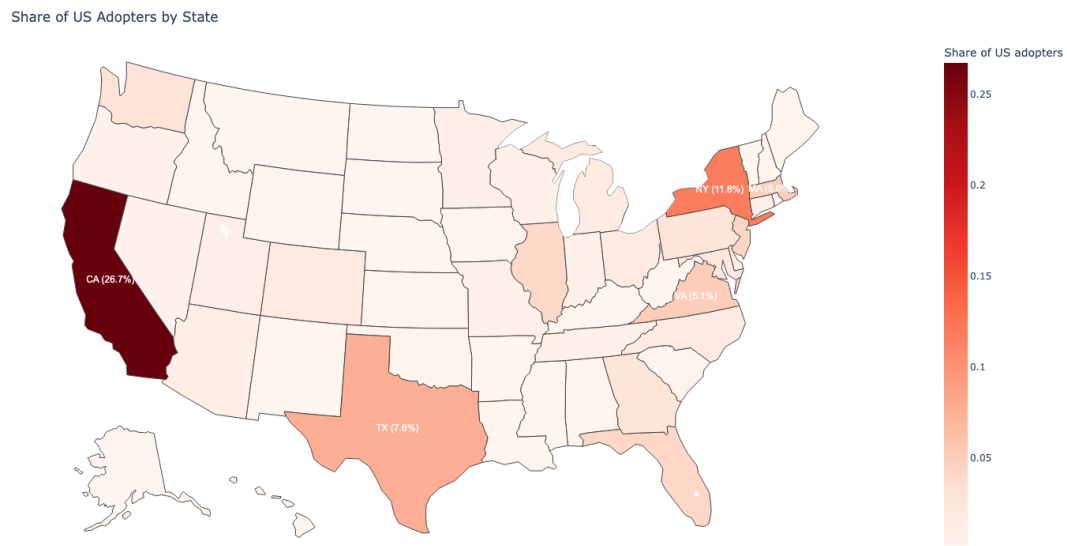


Figure A.8: Share of U.S. adopters by state.

A.6 API Prompts

A.6.1 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/LlamaIndex, 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.6.2 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):

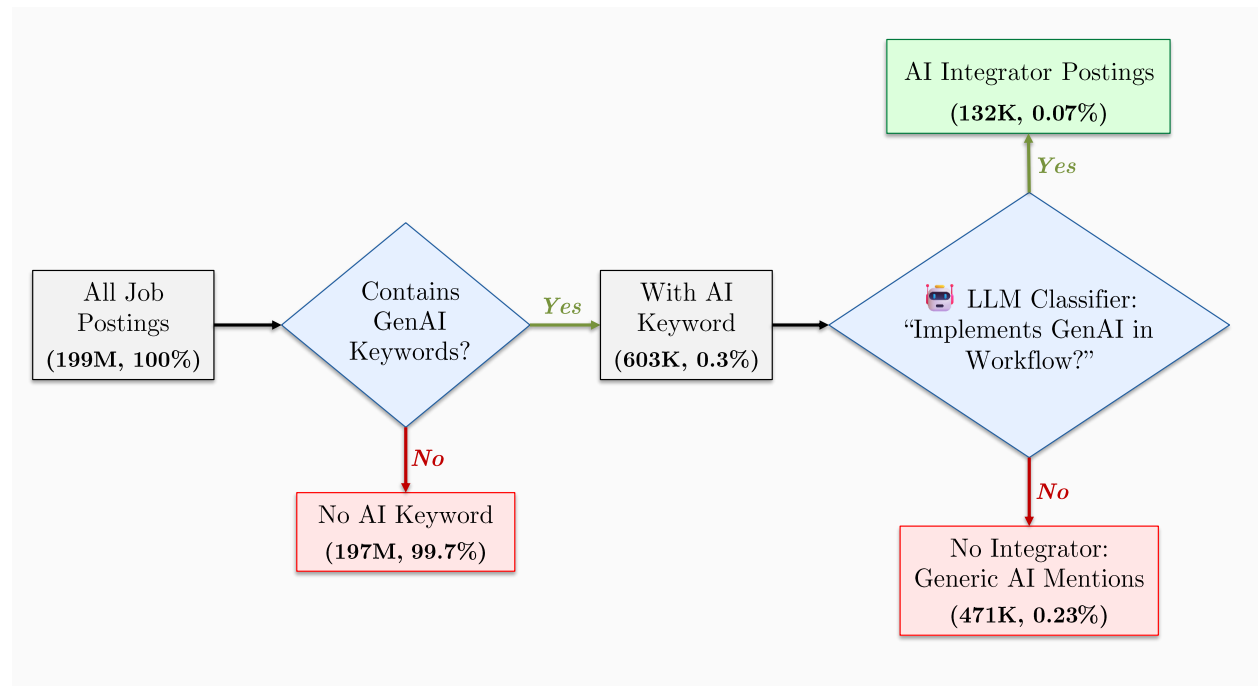
- Return EXACTLY {n} lines.*
- Each line contains ONLY one integer in 1..5 for the corresponding line below.*
- Do NOT include any keys, bullets, indexes, punctuation, or extra text.*
- Do NOT include blank lines.*
- STOP OUTPUT immediately after printing the {n}th line.*

NAMES (one per line, in order):

{names_block}

'''

A.7 Detecting GenAI Integrator Postings—Graphical Illustration



A.8 Triple-Difference Without Industry-by-Time-by-Seniority Fixed Effects

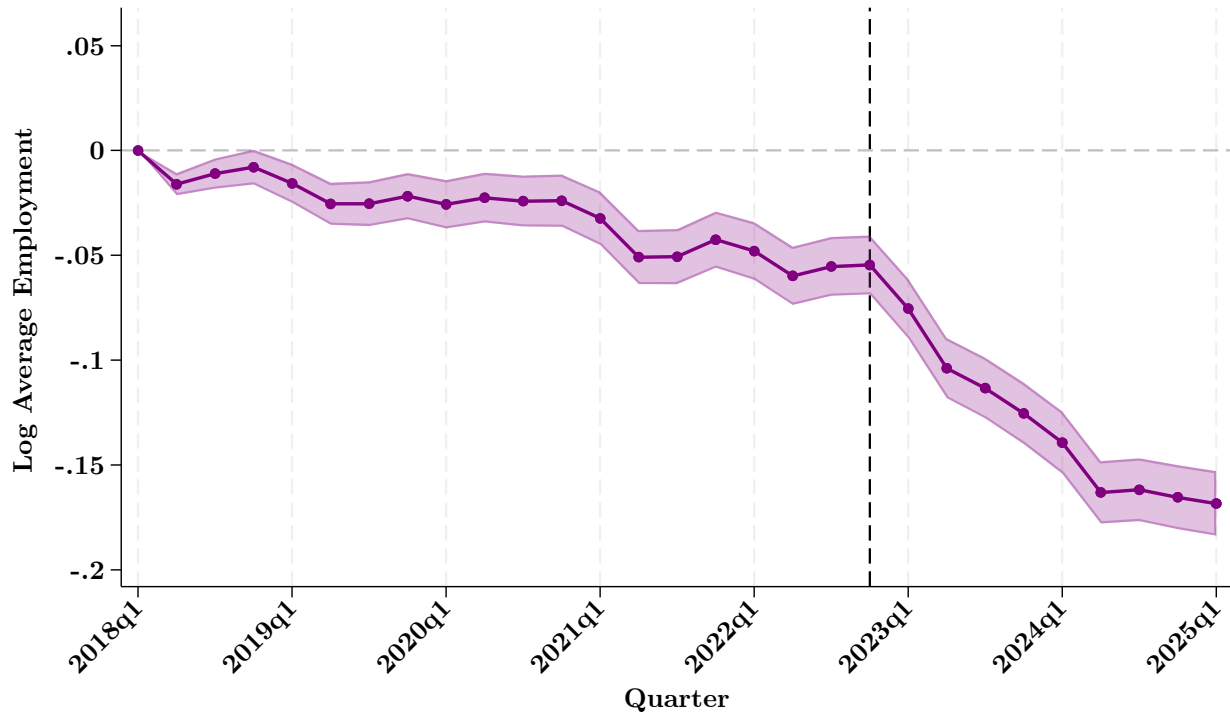


Figure A.9: Triple Differences (Without Industry-by-Time-by-Seniority Fixed Effects)

Notes: This graph shows the results of the same exercise as in Panel (b) of Figure 4, excluding the industry-by-time-by-seniority fixed effects.

A.9 Robustness to Firm Size Differences

Although our main specification (triple difference, equation 2) includes firm-by-time fixed effects that account for differing dynamics across firms with varying characteristics (e.g., size, managerial ability), the significant size gap between adopters and non-adopters may still pose concerns. Larger firms (see Table 1) could exhibit distinct adjustment patterns to shocks compared with smaller firms or may have a heterogeneous treatment effect. To address this, we re-estimate our staggered difference-in-differences specification (3) by restricting the sample to larger firms.

Panel A.12a reports results for firms above the median size in 2018Q1 (pre-period), while Panel A.12b focuses on the top 25 percent of the size distribution in 2018. In both cases, we find precisely estimated zero pre-trends for six quarters prior to AI adoption and a significant post-adoption decline in junior employment, consistent with Figure 5.




Figure A.10: Event Study by Size

Notes: The graph presents the estimated coefficients β_j from Equation 3 using the method of Callaway and Sant'Anna (2021) for staggered adoption. Panel (a) shows the results for the top 50% of firms in size distribution in 2018Q1, while Panel (b) depicts the results for firms above 75th percentile. Standard errors are clustered in firm level.

A.10 Macroeconomic Conditions

As discussed in the main text, one potential threat to our identification is that changes in the macro-financial environment may exert differential pressures on adopter versus non-adopter firms. In particular, the sharp rise in interest rates beginning in early 2022 may have tightened credit and financial conditions in a way that disproportionately affected capital-intensive or borrowing-constrained firms. To guard against this confounder, we incorporate a broad measure of financial conditions in our robustness checks: the *Chicago Fed National Financial Conditions Index (NFCI)* and its adjusted variant (ANFCI).




Figure A.11: Chicago FED Financial Conditions Index

The NFCI is a comprehensive weekly index that aggregates more than 100 indicators spanning money markets, debt and equity markets, and both traditional and shadow banking sectors.²¹ The index is standardized to have mean zero and unit variance (over the sample starting in 1971), so that positive values indicate *tighter-than-average* financial conditions and negative values indicate *looser-than-average* conditions.

The NFCI and its adjusted version have been widely used to track financial stress, fore-

²¹For more information, see the Chicago Fed's NFCI FAQ (2023).

cast macroeconomic activity, and study the transmission of monetary policy and credit shocks (e.g., [Brave and Butters, 2012](#); [Hatzius et al., 2010](#); [Gilchrist and Zakrajšek, 2012](#)).

To account for shifts in financial conditions over time, we extend our pre-treatment window back to 2015, which includes earlier episodes of credit-market tightening. [Figure A.11](#) plots the quarterly average of the NFCI over our sample period. As shown, financial conditions tightened notably in 2015–2016 while our estimates display flat pre-trends in that period. Since 2022, despite a sharp rise in nominal interest rates, the NFCI has eased, suggesting that higher policy rates do not necessarily translate into tighter credit availability.

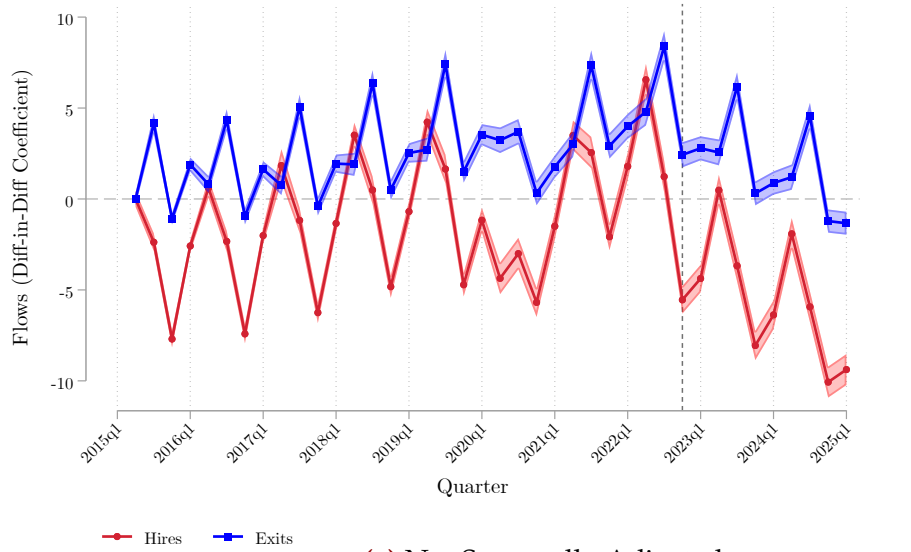
A.11 Most Common High- and Low-Exposed Occupations by Industry

Table A.1: Most Common Low/High-Exposure Occupations by Industry (Share of All ONET Roles in Industry)

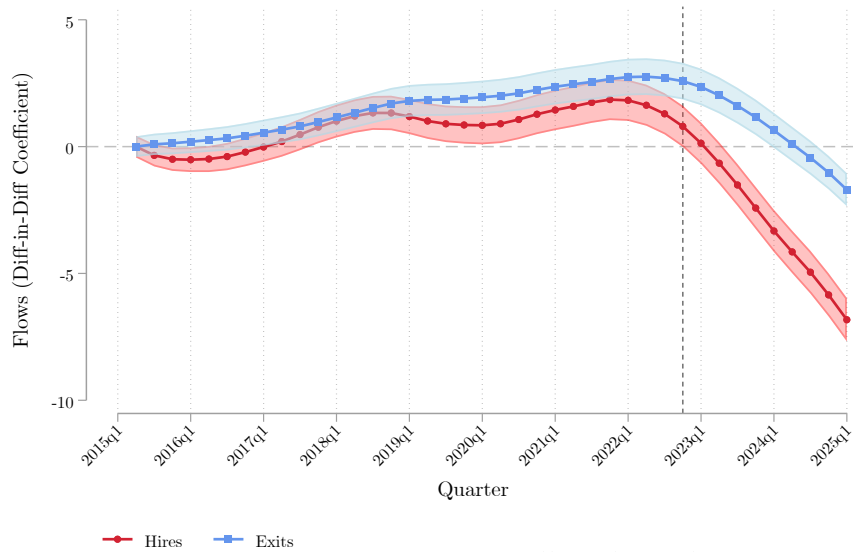
NAICS	Low Exposure (Top 5)	High Exposure (Top 5)
3 <i>Manufacturing</i>	<ul style="list-style-type: none"> – Door-to-Door Sales Workers, News and Street Vendors, and Related Workers (5.4%) – Retail Salespersons (2.2%) – Maintenance Workers, Machinery (1.8%) – Machinists (1.4%) – Biofuels Processing Technicians (1.4%) 	<ul style="list-style-type: none"> – Software Developers (6.8%) – Computer User Support Specialists (2.7%) – Customer Service Representatives (2.5%) – Bioengineers and Biomedical Engineers (2.0%) – Validation Engineers (1.9%)
4 <i>Trade / Retail</i>	<ul style="list-style-type: none"> – Retail Salespersons (9.9%) – Gambling Change Persons and Booth Cashiers (5.8%) – Cashiers (4.6%) – Stockers and Order Fillers (4.0%) – Merchandise Displayers and Window Trimmers (3.5%) 	<ul style="list-style-type: none"> – Customer Service Representatives (5.8%) – Computer User Support Specialists (1.7%) – Bookkeeping, Accounting, and Auditing Clerks (1.6%) – Software Developers (1.6%) – Sales Representatives, Wholesale and Manufacturing, Except Technical and Scientific Products (1.6%)
51 <i>Information</i>	<ul style="list-style-type: none"> – Actors (2.3%) – Retail Salespersons (1.7%) – Door-to-Door Sales Workers, News and Street Vendors, and Related Workers (1.2%) – Nannies (1.1%) – Career/Technical Education Teachers, Secondary School (1.0%) 	<ul style="list-style-type: none"> – Software Developers (13.1%) – Writers and Authors (6.0%) – News Analysts, Reporters, and Journalists (5.6%) – Editors (4.2%) – Customer Service Representatives (4.0%)
52 <i>Finance & Insurance</i>	<ul style="list-style-type: none"> – Gambling Change Persons and Booth Cashiers (0.9%) – Door-to-Door Sales Workers, News and Street Vendors, and Related Workers (0.8%) – Phlebotomists (0.6%) – Retail Salespersons (0.6%) – Nannies (0.5%) 	<ul style="list-style-type: none"> – Loan Officers (7.9%) – Customer Service Representatives (7.2%) – Securities, Commodities, and Financial Services Sales Agents (6.1%) – Loan Interviewers and Clerks (5.5%) – Software Developers (4.4%)
54 <i>Professional Services</i>	<ul style="list-style-type: none"> – Demonstrators and Product Promoters (1.0%) – Medical and Clinical Laboratory Technicians (1.0%) – Door-to-Door Sales Workers, News and Street Vendors, and Related Workers (0.9%) – Retail Salespersons (0.9%) – Career/Technical Education Teachers, Secondary School (0.7%) 	<ul style="list-style-type: none"> – Software Developers (7.4%) – Writers and Authors (5.7%) – Accountants and Auditors (4.1%) – Computer User Support Specialists (4.1%) – Public Relations Specialists (3.2%)
61 <i>Educational Services</i>	<ul style="list-style-type: none"> – Substitute Teachers, Short-Term (20.0%) – Career/Technical Education Teachers, Secondary School (19.1%) – Coaches and Scouts (3.8%) – Lifeguards, Ski Patrol, and Other Recreational Protective Service Workers (3.5%) – Nannies (3.4%) 	<ul style="list-style-type: none"> – Public Relations Specialists (3.1%) – Computer User Support Specialists (2.6%) – Executive Secretaries and Executive Administrative Assistants (1.8%) – Software Developers (1.8%) – Writers and Authors (1.7%)
62 <i>Health Care & Social Assistance</i>	<ul style="list-style-type: none"> – Acute Care Nurses (14.3%) – Licensed Practical and Licensed Vocational Nurses (11.5%) – Phlebotomists (7.9%) – Home Health Aides (6.2%) – Nannies (5.4%) 	<ul style="list-style-type: none"> – Eligibility Interviewers, Government Programs (2.2%) – Executive Secretaries and Executive Administrative Assistants (1.9%) – Public Relations Specialists (1.8%) – Computer User Support Specialists (1.7%) – Customer Service Representatives (1.5%)

A.12 DiD for Hires and Exits—Time Series (Juniors)

$$y_{it} = + \sum_{j=2015Q2}^{2025Q1} \beta_j \mathbf{1}\{t = j\} \times \text{Adopt}_i + \delta_t + \gamma_i + \varepsilon_{it}, \quad (5)$$



(a) Not Seasonally Adjusted



(b) Seasonally Adjusted

Figure A.12: DiD for Hires and Exits—Time Series (Juniors)

Notes: This Figure plots β_j from Equation 5. For seasonal adjustment, the coefficients are normalized to zero at 2015Q2 and are smoothed using LOWESS (bandwidth = 0.5) before plotting.

A.13 Average Salary and Exposure by Educational Background




Figure A.13: Predicted Salary by School Quality (Juniors, 2022)

Notes: Bars report average predicted salaries (in USD) for juniors employed in 2022 by university prestige category.




Figure A.14: Average Exposure by School Quality (Juniors, 2022)

Notes: Bars report exposure for juniors employed in 2022, by university prestige category. The standard deviation of the exposure variables is 0.21.