

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 27, 2025

Abstract

We study whether generative AI (GenAI) constitutes *seniority-biased technological change*, disproportionately affecting junior relative to senior workers. Using U.S. résumé data covering 62 million workers across 285,000 firms (2015–2025), we track firm-level employment by seniority. GenAI adoption is identified through text analysis that detects “GenAI integrator” job postings, signaling active GenAI implementation by firms. Following adoption, junior employment declines sharply in adopting firms relative to non-adopters, while senior employment remains largely unchanged. The junior decline is concentrated in occupations most exposed to GenAI and is driven by slower hiring rather than increased separations or promotions.

*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 David Autor, Erik Brynjolfsson, Gabriel Chodorow-Reich, Oren Danieli, Eric Gold, Jonathon Hazell, Ori Hefetz, David Lagakos, Amanda Pallais, Ludwig Straub, Santiago Medina, and Austin Zheng for their invaluable discussions and advice throughout various stages of this work. We additionally thank Shotaro Beppu, Martin Bernstein, Olena Bogdan, Bharat Chandar, Ruyu Chen, Fiona Chen, Rupsha Debnath, Cameron Deal, Aristotle Epanomeritakis, Sarah Gao, Jay Garg, Shao-Yu Jheng, Xianruo Kang, Iris Li, Saketh Prazad, Ewan Rawcliffe, Emanuel Schertz, Peyman Shahidi, Ragini Srinivasan, James Stratton, and Katherine Wang 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: shosseinimaasoum@fas.harvard.edu

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

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—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 often 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 (Garicano and Rayo, 2025).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 The issue has also drawn attention from policymakers. During the September 2025 FOMC press conference, Chair Powell was asked about the impact of AI on the labor market and noted: “*You are seeing some effects ... A particular focus 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.*” (Federal Reserve Board, 2025).

Some observers have linked the recent rise in unemployment among college graduates to the diffusion of GenAI (*The New York Times*, 2025a; *The Atlantic*, 2025; *Forbes*, 2025). 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).

¹The implications can extend beyond short-term employment effects, as early-career earnings strongly influence lifetime income trajectories and, consequently, inequality (Deming, 2023; Guvenen et al., 2022).

²Moreover, a 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).

This paper aims 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 and senior employment within firms over time.

GenAI adoption is identified 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

By this definition, 10,599 firms in our sample adopted GenAI by March 2025. 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 release of advanced GenAI tools in late November 2022.

We begin our analysis by comparing adopting and non-adopting firms using a difference-in-differences (DiD) design, tracking junior and senior employment quarterly. From 2015 to 2022, adopters and non-adopters followed parallel trends in junior employment. How-

³See [Ide and Talamas \(2025\)](#) for a closely related theoretical framework, analyzing how GenAI may differentially affect workers within firms depending on their knowledge and role.

⁴Our concept of a GenAI integrator is closely related to the “robot integrator” in [Acemoglu and Restrepo \(2020\)](#), defined as “companies that install, program, and maintain robots.” They use the presence of robot integrators as a proxy for local robot adoption and show that commuting zones more exposed to robots have a higher number of integrators.

ever, beginning in 2023Q1—coinciding with the sharp increase in GenAI adoption—junior employment in adopting firms decreased steeply relative to controls, declining by about 9 percent after six quarters. Senior employment, by contrast, increased more quickly in adopting firms since 2015 and showed no sign of a break in trend after 2022.

To directly assess the “seniority-biased” effects, we use a triple-difference specification, comparing changes in junior versus senior employment within adopting firms relative to non-adopters. Importantly, this design incorporates firm-by-time fixed effects, which absorb any shocks or trajectories specific to a given firm in a given period. Additionally, we include industry-by-time-by-seniority fixed effects, which account for sector-level dynamics that differentially affect juniors and seniors. The results align with the DiD estimates: apart from a brief dip in early 2021, coefficients remain stable between 2018Q1 and 2022Q4, then decline sharply starting in 2023Q1, reaching roughly a 10 percent drop after six quarters.

Next, we examine heterogeneity by occupational exposure to GenAI. Merging our positions data with exposure measures from [Eloundou et al. \(2024\)](#), we find that the relative decline in junior employment at adopting firms is concentrated in high-exposure roles, while low-exposure roles show no significant change. The post-2022 decline in junior employment among adopting firms therefore reflects adjustments in roles most vulnerable to GenAI, rather than a broad contraction in junior employment.

We complement the DiD and triple-difference analyses with a staggered event study that traces employment dynamics around firms’ GenAI adoption date, proxied by the first period in which the firm posts a GenAI integrator role. This design helps distinguish adoption effects from broader time-specific shocks by exploiting variation in the timing of adoption across firms. However, it is sensitive to measurement error in the adoption timing proxy—for instance, if firms begin using GenAI before, or only several periods after, posting for an integrator role. The estimates show no significant differences between adopting and non-adopting firms in the eight quarters preceding adoption, consistent with parallel pre-trends. Roughly three quarters after adoption, junior employment in adopting firms begins to decline, reaching an 8 percent reduction after eight quarters. The absence of significant pre-trends provides additional reassurance that these post-adoption declines are unlikely to be driven by confounding shocks.

While our results indicate that GenAI adoption is associated with a decline in junior

employment, potential endogeneity concerns remain, as adopting firms tend to be larger, more technologically intensive, and employ a more highly educated workforce. We therefore discuss how our research design addresses key confounding factors that could bias the estimates, such as the 2022–2023 monetary tightening cycle or post-Covid corrections in the technology sector. Specifically: (i) sector-specific sensitivities to such shocks are absorbed by our industry-by-seniority-by-time fixed effects (ii) no corresponding pre-adoption trends appear in the event-study estimates; (iii) adopters and non-adopters followed similar trajectories during earlier tightening episodes (2015–2018); and (iv) larger, larger firms, which are more likely to adopt GenAI, tend to be less responsive to monetary policy (Chodorow-Reich, 2014; Gertler and Gilchrist, 1994). We also consider supply-side explanations—for example, if negative shocks to junior labor availability prompted adoption. Beyond the parallel pre-trends, which seem inconsistent with this story, we examine a measure that more directly captures labor demand: job postings. The decline in junior employment coincides with reduced postings, suggesting a contraction in labor demand rather than a constraint on labor supply.

We next examine the mechanisms underlying the decline in junior employment. Using our linked employer-employee data, we decompose workforce dynamics into inflows (hires), outflows (separations), and internal promotions. We find that the decline in junior employment at adopting firms is driven primarily by a substantial reduction in hiring. Separation rates for juniors in adopting firms also decreased relative to non-adopters, but the reduction in hiring was considerably larger, leading to a net decline in junior positions. Promotion rates, by contrast, remained broadly stable after early 2023.

We further investigate heterogeneity in the junior hiring decline by workers' educational background and by sector. By educational background, we find a U-shaped pattern: juniors from mid-tier institutions experienced the steepest relative declines in hiring, while those from the most and least prestigious schools also saw reductions, though of smaller magnitude. Results by sector show that the contraction in junior hiring is broad-based across industries and not disproportionately driven by a single sector.

Taken together, these results indicate that GenAI adoption is associated with reduced junior employment. However, importantly, this does not necessarily imply immediate task automation. A plausible alternative is that firms are making forward-looking adjustments: the rapid diffusion of GenAI 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. We present a simple dynamic model illustrating this mechanism, in which expectations of *future* automation—combined with labor-adjustment costs—can lead firms to reduce hiring *today*. Three findings are consistent with this interpretation: (i) the sharp acceleration in GenAI adoption in early 2023, consistent with an expectations shock following the release of GPT-3.5; (ii) the steep decline in junior employment that begins shortly thereafter; and (iii) the finding that slower hiring, rather than increased separations, drives the decline. Under this view, if firms have overreacted to the introduction of GenAI, the resulting decline in junior employment may be temporary. Nonetheless, our data do not allow us to directly test this mechanism.

It is worth noting that our 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, silent adoption may have distinct implications for employment dynamics from the deliberate, firm-level adoption we capture. Moreover, as with any empirical analysis based on cross-sectional variation, our results do not necessarily generalize to the aggregate economy without additional assumptions, due to the “missing intercept problem” (see [Acemoglu and Restrepo, 2020](#); [Wolf, 2023](#); [Moll and Hanney, 2025](#)).

The remainder of the paper is organized as follows. Section 2 reviews related literature. Section 3 describes the data and descriptive patterns. Section 4 presents our empirical strategy and main findings. Section 6 concludes.

2 Related Literature

Our study relates to three main strands of the literature: (i) skill-biased technological change, (ii) firm-level adoption of AI technologies, and (iii) the emerging evidence on the labor-market impacts of Generative AI.

Skill-Biased Technological Change: The classic literature on skill-biased technological change (SBTC) 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.⁵

Evidence from Firm-Level AI Adoption: A related strand of work—closest to our empirical approach—examines 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’ pro-

⁵While it is usually not the main focus, the SBTC literature has also examined variation by age/experience. For instance, Acemoglu and Autor (2011) show that the surge in the college wage premium during the 1980s was concentrated among less experienced workers, whereas from the mid-1990s onward, the increase was driven primarily by more experienced cohorts. They argue this pattern is consistent with changes over time in relative supply by experience and education (and imperfect substitution across experience groups).

ductivity gains offset job losses by expanding employment elsewhere, resulting in muted net changes in total headcount.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-adopting firms reduce junior employment while leaving senior employment unaffected.

The Effects of Generative AI on the Labor Market: Finally, a rapidly growing empirical literature examines the labor-market effects of GenAI. 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](#)).

A second wave of studies uses large-scale labor-market data to track employment dynamics across occupations and industries by GenAI exposure. Closest to our work, [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

⁶Two related studies examine the effects of *Generative* AI adoption. [Humlum and Vestergaard \(2025\)](#) provide evidence from Denmark, linking large-scale worker surveys on GenAI adoption to matched employer-employee data. Despite rapid adoption, they do not find effects of adoption on earnings and hours, suggesting that labor-market impacts in Denmark remain minimal. [Chen 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.

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. Moreover, unlike exposure-based measures, our adoption measure varies over time, allowing us to exploit variation in adoption timing with an event-study design.

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.

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 and then categorized into seven standardized seniority levels: Entry Level, Junior Level, Associate

Level, Manager Level, Director Level, Executive Level, and Senior Executive Level. In the analysis that follows, we group positions into two broad categories: *juniors* (Entry and Junior) and *seniors* (Associate and above).⁷ Appendix A.1 provides further detail and validation of the seniority variable.

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

Finally, although not the primary focus of the paper, we also incorporate the occupational GenAI 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 exposure (0–25th percentile), medium exposure (25–75th percentile), and high exposure (75–100th percentile).

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 2025.⁸ 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.⁹

3.2 Workforce Dynamics by Seniority

We construct a monthly panel at the firm level. For each firm-by-month, 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, enabling us to track the workforce composition over time. Additionally, we identify monthly inflows and outflows by seniority. For

⁷More details on Revelio’s seniority classification methodology are available at <https://www.data-dictionary.reveliolabs.com/methodology.html#seniority>.

⁸We define a firm as active if it recorded at least 20 new hires over this period.

⁹We 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.

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 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.¹⁰

Figure 1 presents the aggregate time series of average junior and senior employment across firms. Between 2015 and 2022, employment for both groups expanded at a similar pace, aside from a temporary decline in junior employment during the Covid period. Beginning in mid-2022, however, a clear divergence emerges: senior employment continues to expand steadily, while junior employment plateaus and then, by mid-2023, begins to fall. This divergence is consistent with the findings of Brynjolfsson et al. (2025a), who document a comparable pattern in U.S. payroll data by worker age. The consistency 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 GenAI adoption by detecting job postings that explicitly seek workers to implement or integrate GenAI technologies into firm workflows. 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.¹¹ Out of 198.8 million postings with usable text descriptions, 603,152 (0.30 percent) include at least one keyword. Second, we apply a large language model (LLM) classifier to this subset to

¹⁰Roughly 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 such positions 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.

¹¹Keywords 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.

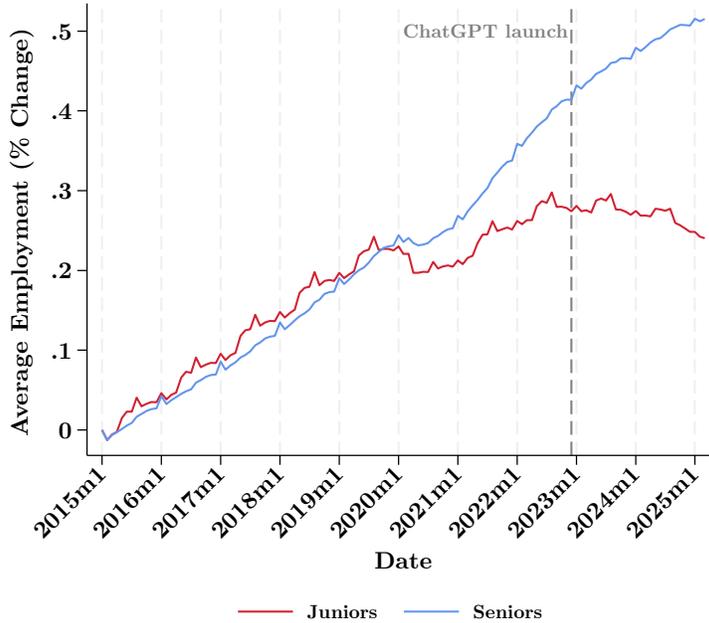


Figure 1: Log Average Employment of Junior and Senior in Sample Firms

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. “Junior” refers to Entry- and Junior-level positions, while “Senior” refers to Associate level and above (see Section 3.1 for details).

identify genuine “GenAI integrator” postings—vacancies reflecting an active attempt to recruit workers tasked with adopting or implementing GenAI in the firm’s workflows—from false positives (appendix A.1.1 provides the exact prompt used). This process yields 131,845 postings, 0.066 percent of the full corpus, classified as GenAI integrator roles (the accompanying online appendix provides a graphical overview of this procedure). Appendix A.2 presents illustrative examples of postings identified as GenAI-integrator roles, alongside examples that were flagged by the keyword search but classified as non-integrators by the LLM.¹²

¹²As shown in Appendix A.1.1, 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 Appendix A.3) and account for 17.3 percent of the employment (positions) in our dataset.

Figure A.4 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 launch of ChatGPT—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.

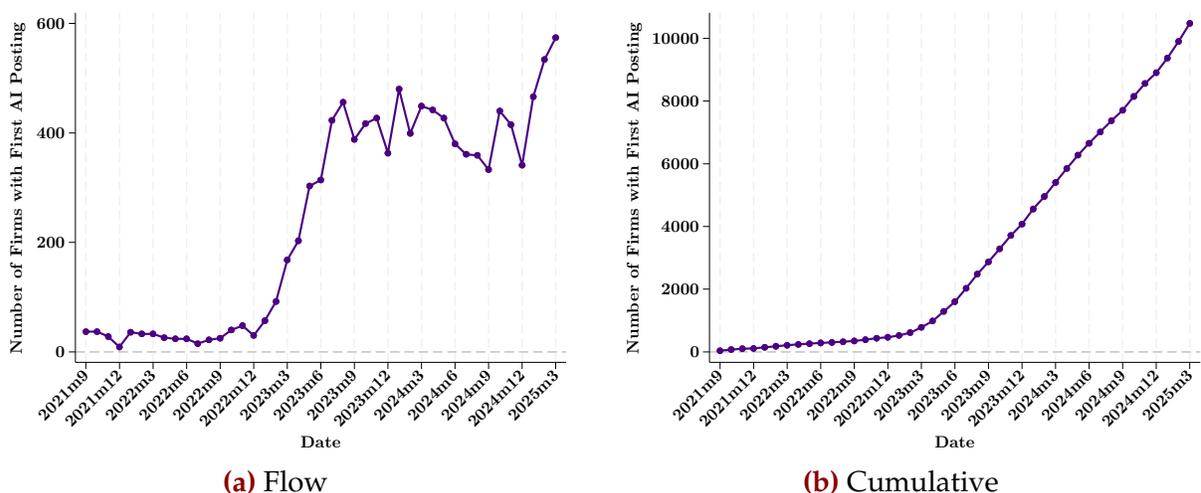


Figure 2: Adoption Distribution Over Time

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

Additionally, 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.

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.

4 Results

4.1 Employment Dynamics by Adoption and Seniority

Descriptive Time Trends: We begin by comparing the evolution of junior and senior employment over time in GenAI-adopting versus non-adopting firms. Figure 3a displays the average employment of junior workers across the two groups. From 2018 through the end of 2022, junior employment followed closely parallel trajectories. Starting in late 2022, however, the trends diverged sharply: junior employment in adopting firms began to decline markedly, while employment in non-adopting firms remained relatively stable. By contrast, Figure 3b shows that since 2018, senior employment in adopting and non-adopting firms grew at a similar pace—aside from a brief acceleration in adopting firms in 2021—with no apparent break in trend around 2023.¹³

¹³These patterns are consistent with Brynjolfsson et al. (2025a), who document a similar divergence by occupational exposure among younger workers in U.S. payroll data, with no corresponding shift for older workers.

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.

Difference-in-Difference: 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. Specifically, we estimate the following 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 quarter t .¹⁴ 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.¹⁵ 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 3c 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 about 9 percent relative to controls six quarters after the diffusion of GenAI. 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.

Triple-Difference: To directly assess the “seniority-biased” effects associated with GenAI adoption, we estimate a triple-difference specification that compares changes in junior versus senior employment within adopting firms relative to non-adopters:

¹⁴For computational efficiency we aggregate our panel data to quarterly frequency in all analyses.

¹⁵Note 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 4.3, we relax this definition by exploiting variation in the timing of adoption across firms using an event-study design.

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

Here, $\log(\text{Employment}_{ist})$ 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, 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 .

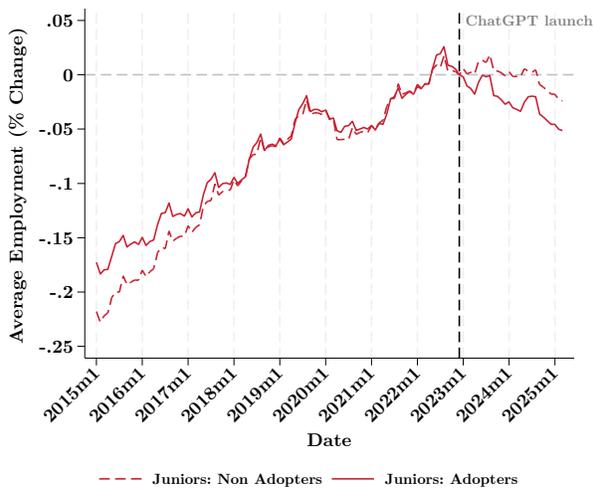
The coefficients β_j trace a triple-differences event-time profile: they capture how junior employment evolves relative to senior employment *within the same firm and period*, comparing adopters to non-adopters. Firm-by-time fixed effects γ_{it} absorb any shocks or trajectories specific to a given firm in a given period. The industry-by-seniority-by-time fixed effects $\xi_{p(i)st}$ control for broad sectoral trends that differentially affect juniors and seniors over time, ensuring that our results are not driven by sector-wide shifts in junior versus senior employment.

Figure 3d presents the estimated coefficients β_j from Equation 2. Estimation begins in 2018Q1 due to computational constraints. The coefficients remain essentially flat through 2022Q4, aside from a brief dip in early 2021. The patterns shown in Panels 3a and 3b indicate that this temporary decline likely reflects a brief acceleration in senior employment among adopting firms in 2021, rather than a decrease in junior employment.

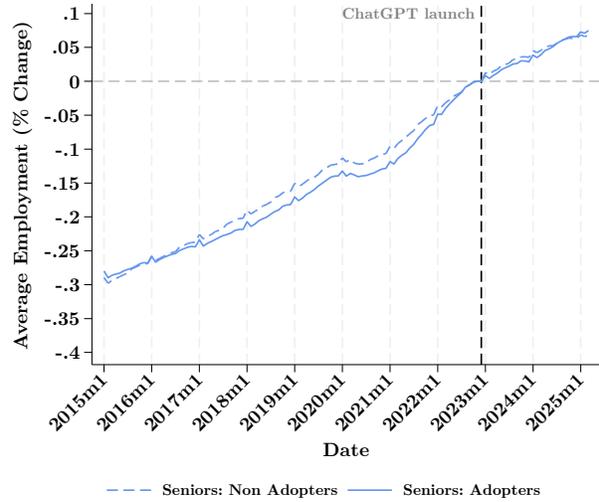
Beginning in 2023Q1, the coefficients drop sharply, reaching roughly a 10 percent decrease after six quarters. This break—coinciding with the rapid diffusion of GenAI—provides suggestive evidence that adoption is associated with increasingly seniority-biased patterns, reducing junior employment relative to senior employment within firms.¹⁶

The β_j coefficients can be interpreted as reflecting the differential impact of GenAI

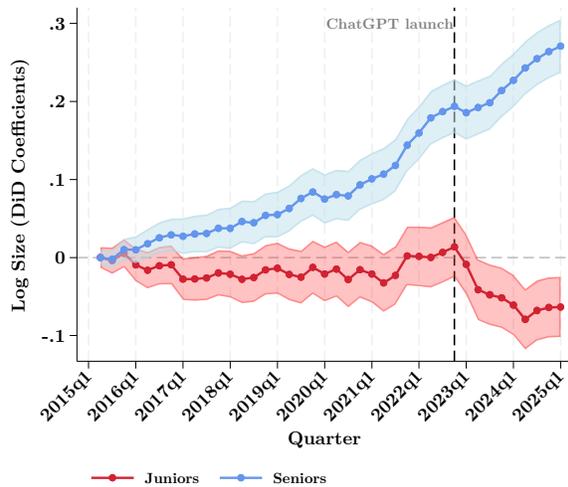
¹⁶In the accompanying online appendix, we present triple-difference estimates excluding the industry-by-seniority-by-time fixed effects. The results are very similar to those in Figure 3d.



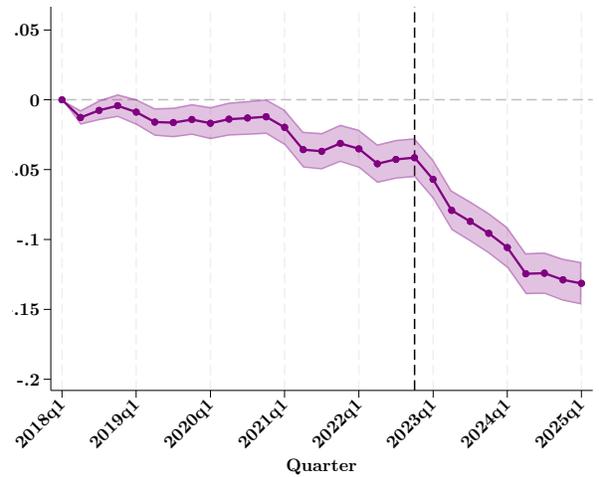
(a) Average Employment—Juniors



(b) Average Employment—Seniors



(c) Difference-in-Difference



(d) Triple-Difference

Figure 3: Employment Patterns by Adoption and Seniority

Notes: Panels (a) and (b) show the average number of junior and senior workers, 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. Panel (d) presents the estimated coefficients β_j from Equation 2. The adoption measure used in this analysis is time-invariant; hence, the figures simply compare adopters and non-adopters. For results using a time-varying adoption measure, see Section 4.3. Standard errors are clustered at the firm level.

adoption on junior relative to senior employment, provided that—after conditioning on firm-by-time and industry-by-seniority-by-time effects—no other factors since 2023 systematically affected juniors and seniors differently across adopting and non-adopting firms. This might be a concern due to the potential endogeneity of GenAI adoption. The close similarity in junior employment dynamics between adopters and non-adopters from 2015 through 2022 helps alleviate this concern by supporting the plausibility of parallel pre-trends. We provide further support for our interpretation through analyses by occupational exposure (Section 4.2) and an event-study design (Section 4.3). Section 4.4 further discusses possible forces that could challenge this interpretation and how we address them.

Finally, the early and pronounced decline in the DiD and triple-difference coefficients—beginning in 2023Q1, shortly after the release of GPT-3.5—may appear surprisingly abrupt, given that automation effects typically unfold more gradually. However, the rapid surge in adoption from early 2023 onward—as shown in Figure A.4 and documented by studies such as Bick et al. (2024)—suggests that firms perceived the rapid diffusion of GenAI as a discrete shock. This sudden acceleration may have shifted firms’ expectations about future automation, prompting them to adjust in a forward-looking manner by reducing junior roles they anticipated would be automated in the near future. Such preemptive adjustments could be optimal for firms if they view slower hiring as less costly than future layoffs. Section 5 provides a simple dynamic model illustrating this mechanism, in which expectations of *future* automation—combined with labor-adjustment costs—lead firms to reduce hiring *today*. That said, our data cannot directly test this mechanism, so we offer it only as a plausible interpretation.

4.2 DiD by Occupational Exposure

We next re-estimate the DiD specification (Equation 1) separately for junior positions with high and low exposure to GenAI.¹⁷ Figure 4 reports the results. For high-exposed occupations, the estimated coefficients rise steadily from 2015 through 2022Q3, indicating that adopting firms were expanding junior employment in these roles relative to non-adopters prior to the diffusion of GenAI. Beginning in 2022Q4, however, this trend re-

¹⁷Appendix A.4 shows the most common high- and low-exposure occupations by industry, ranked by number of positions.

verses sharply, with high-exposure junior employment in adopting firms declining significantly relative to non-adopters. In contrast, coefficients for low-exposure junior occupations decline gradually between 2015 and 2019 and then remain stable through early 2025, with no discernible break around late 2022.

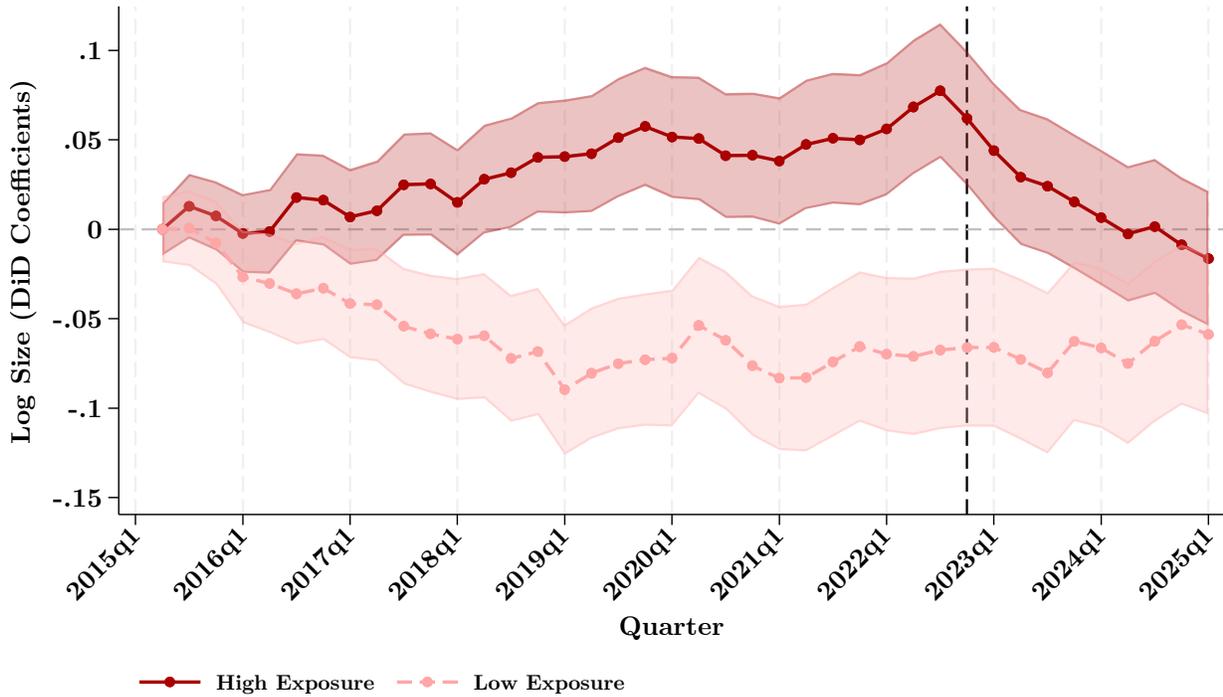


Figure 4: DiD Estimates for High- and Low-Exposure Occupations (Juniors)

Notes: This figure reports separate difference-in-differences estimates (β_j from Equation 1) for juniors in high- and low-exposure occupations. Details on the exposure measure are provided in Section 3.1. Standard errors are clustered at the firm level.

Taken together, these results provide two insights. First, the post-2022 decline in junior employment among adopting firms is concentrated in high-exposure occupations, rather than reflecting a broad contraction in junior employment across these firms. Second, the design helps rule out a purely compositional explanation—namely, that high-exposure occupations declined economy-wide after 2023, and adopting firms simply employed a larger share of such roles. By estimating DiD separately by exposure categories, we compare like with like—junior employment in adopting and non-adopting firms within occupations of the same exposure level. The observed post-2022 decline in the high-exposure group therefore reflects a contraction within similarly exposed occu-

pations among adopters relative to non-adopters.

4.3 Staggered Event-Study of Junior Employment Around Adoption

We complement the DiD and triple-difference analyses with a staggered 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 distinguish 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 or only several periods afterward. Specifically, we estimate:

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

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.

Figure 5 reports the results. The coefficients remain flat during the pre-adoption period, supporting the validity of the parallel-trends assumption. Approximately three quarters after adoption, employment begins to decline, reaching an 8 percent reduction after eight quarters. The absence of significant pre-trends provides additional reassurance that these post-adoption declines are not driven by confounding shocks.

4.4 Addressing Endogeneity Concerns

This section examines potential endogeneity concerns that could confound our interpretation of the estimated effects of GenAI adoption on junior employment. We discuss both demand- and supply-side shocks that may coincide with the diffusion of GenAI and explain how our empirical design mitigates these risks.

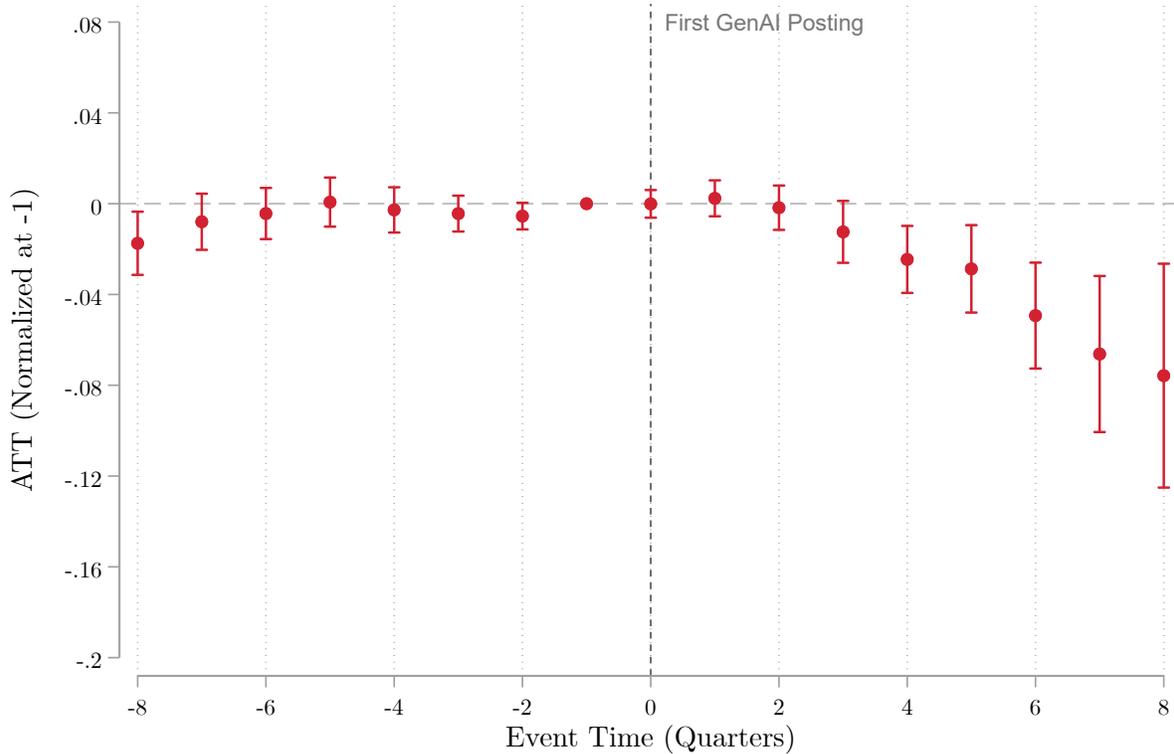


Figure 5: Event Study

Notes: The graph presents the estimated coefficients β_j from Equation 3 using the method of Callaway and Sant’Anna (2021) for staggered adoption. Firms with zero recorded employment in 2021Q1—eight quarters before the diffusion of GenAI—are excluded (2.3 percent of firms). Standard errors are clustered at the firm level.

4.4.1 Confounding Demand Shocks

A key concern is that other contemporaneous shocks may have affected labor demand for junior workers among GenAI-adopting firms, potentially biasing our estimates. We focus on two commonly discussed shocks and explain why they are unlikely to drive our results.

The first concern is that GenAI adopters may be more sensitive to monetary policy cycles and thus may have reacted more strongly to the 2022–2023 interest rate hikes. However, several pieces of evidence suggest that this mechanism does not drive our findings. First, if monetary policy shocks had played a role, we would expect to see effects in the pre-trend coefficients of Figure 5, yet none are observed. Second, even if interest rate

changes disproportionately affected industries where adopters are concentrated or had stronger impacts on junior roles, these dynamics should be absorbed by the industry-by-seniority-by-time fixed effects in the triple-difference estimates in Section 4.1. Third, in the difference-in-differences results in this section, our pre-trends extend back to 2015, encompassing the 2015–2018 tightening cycle. As the online appendix shows, credit and financial conditions tightened substantially during that period, yet Figure 3c reveals no relative decline in junior employment among adopters. Finally, because adoption is positively correlated with firm size (Section 3.3), and larger firms tend to be less sensitive to monetary policy shocks (see Chodorow-Reich, 2014; Gertler and Gilchrist, 1994), it is unlikely that differential exposure to interest rate changes explains our results.

The second potential confounder is the post-COVID hiring boom in the technology sector (U.S. Bureau of Labor Statistics, 2025). If this boom was followed by a correction, it could disproportionately affect GenAI adopters, who are more heavily represented in the Information sector (Section 3.3). However, these dynamics are also absorbed by the industry-by-seniority-by-time fixed effects. Moreover, if boom-bust dynamics were driving our results, we would expect to observe a relative increase in junior employment among adopters prior to GenAI adoption, which we do not (see Figure 5).

4.4.2 Supply Shocks and Reverse Causality

Another concern is reverse causality—namely, that firms adopting GenAI were disproportionately affected by negative supply shocks to junior labor, which prompted adoption. To assess this possibility, we examine a more direct measure of labor demand: job postings. Specifically, we re-estimate Equation 1 using the number of job postings as the dependent variable rather than employment. Since the job postings data do not include seniority information, this analysis captures total demand for workers rather than its distribution by seniority. Moreover, the job postings data begin only in 2021Q4.

Figure 6 presents the results. Starting in 2023Q1, the number of job postings declined significantly in adopting firms relative to non-adopters. This finding provides suggestive evidence that the observed decline in junior employment reflects a contraction in labor demand rather than a supply-driven shock.

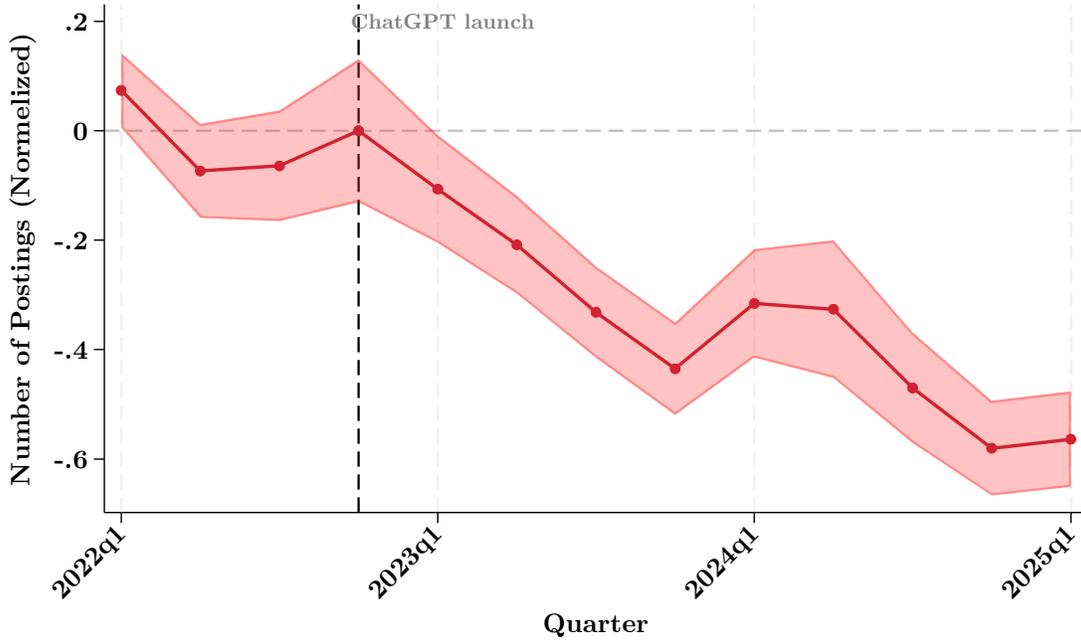


Figure 6: DiD—Job Postings

Notes: This figure reports the difference-in-differences estimates (β_j from Equation 1) for the quarterly number of job postings. The estimates are normalized to represent percentage changes relative to the firm’s average number of postings before 2022. Standard errors are clustered at the firm level.

4.5 Decomposing Decline in Junior Employment Into Flows

4.5.1 Hires, Exits, and Promotions

The decline in junior employment in adopting firms can arise through three channels: (i) reduced hiring, (ii) increased separations, or (iii) increased promotions to senior positions. Our LinkedIn-based résumé data can be viewed as a detailed matched employer-employee dataset, allowing us to track these worker flows over time. Specifically, we estimate separate DiD regressions of the following form:

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

where y_{it} denotes the number of hires, number of separations, probability of promo-

tion,¹⁸ or net changes in juniors in firm i at quarter t . γ_i denotes firm fixed effects, δ_t are time fixed effects and ζ_{pt} are industry-by-time fixed effects.

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

	Hiring	Separation	Promotion	Total Change
Treat \times Post	−5.029*** (0.225)	−1.781*** (0.166)	0.018 (0.015)	−3.721*** (0.149)
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.

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

Table 2 reports the estimated β from Equation 4. The results indicate that the sharp contraction in junior employment among adopters is driven by a slowdown in hiring, rather than by increased exits or promotions. 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. Separation rates for juniors also declined among adopters, though the magnitude of this reduction is smaller than that of hiring, resulting in a negative net hiring effect.¹⁹ The relative number of junior promotions appears unchanged following 2023Q1.

4.5.2 The Junior Hiring Decline—Heterogeneity

Heterogeneity by Human Capital We examine heterogeneity in the decline of junior employment by workers’ educational background. To capture educational quality, 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 (for most workers, this is a college or university, though for some it is a high school). 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

¹⁸Probability of promotion is defined as $\frac{\text{Promotions}_t}{\text{Junior Employment}_{t-1}}$.

¹⁹See Appendix A.5 for the time-series DiD results for junior hires and exits.

score from 1 to 5.²⁰

Then, we re-estimate Equation 4 for junior hires separately for each of the five school-quality tiers. 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.

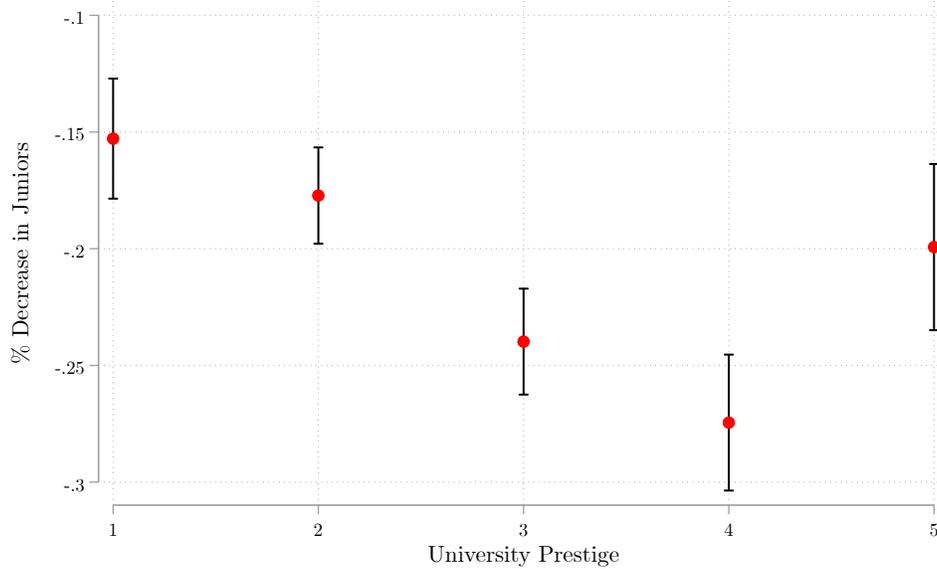


Figure 7: Results by School Quality

Notes: The figure presents the estimates of Equation 4 for junior hires, run separately by university prestige category. The coefficients represent post-2023Q1 changes in junior employment for GenAI adopters relative to non-adopters. Standard errors are clustered at the firm level.

Appendix A.6 provides additional context by plotting average salaries and GenAI-exposure levels of junior positions across school-quality tiers. As expected, there is 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

²⁰These ratings were produced using OpenAI's GPT-4o-mini model, prompted to act as an academic evaluator (Appendix A.1.2 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.

be reshaping demand unevenly across the human capital distribution.

Heterogeneity by Sector: We also examine heterogeneity in the junior decline 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.

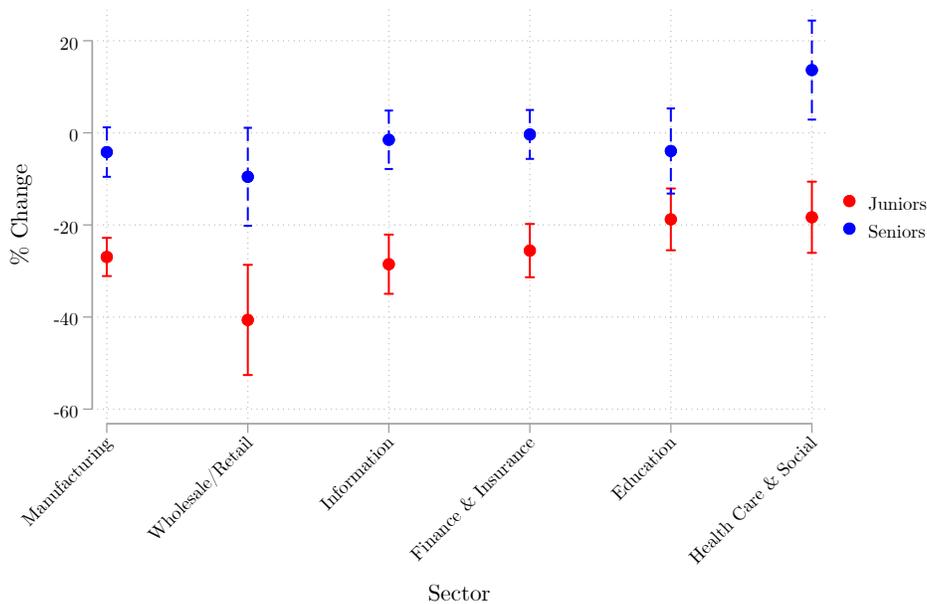


Figure 8: Estimated Effects of Generative AI Adoption on Hiring by Sector

Notes: 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 Automation Tomorrow, Hiring Reductions Today

This section develops a parsimonious framework in which expectations of *future* labor-saving productivity gains depress *current* hiring when employment is costly to adjust. We first present a two-period model that delivers an explicit “firing wedge” in the period-0 hiring condition. We then embed the same mechanism in a task-based production structure in the spirit of [Acemoglu and Restrepo \(2019\)](#).

5.1 A Baseline Two-Period Model

Time is $t \in \{0, 1\}$. A representative firm chooses employment n_t each period and produces with a concave, increasing production function $f(\cdot)$; per-worker wages w_t are exogenous. The firm discounts with factor $\beta \in (0, 1]$ and faces a per-unit hiring cost $c_h \geq 0$ in $t = 0$. Reductions in employment incur a per-worker firing cost $\phi > 0$ on separated workers, so that $\Phi = \phi(n_0 - n_1)_+$. Let $A_t > 0$ denote period- t productivity. Anticipated labor-saving progress in period 1 is captured by a higher A_1 that lowers the amount of labor required per unit of output.

The firm’s objective is

$$\max_{n_0, n_1} \Pi = A_0 f(n_0) - (w_0 + c_h)n_0 + \beta_0 [A_1 f(n_1) - w_1 n_1 - \phi(n_0 - n_1)_+]. \quad (5)$$

In period 1, conditional on n_0 , the first-order condition (FOC) equates the marginal product of labor (MPL) to the effective marginal cost:

$$A_1 f'(n_1) = \begin{cases} w_1, & \text{if } n_1 \geq n_0 \text{ (no firing),} \\ w_1 - \phi, & \text{if } n_1 < n_0 \text{ (firing).} \end{cases}$$

In period 0, the hiring decision internalizes that raising n_0 today can trigger firing tomorrow. The FOC is

$$A_0 f'(n_0) = \begin{cases} w_0 + c_h, & \text{if no firing is expected,} \\ w_0 + c_h + \beta\phi, & \text{if firing is anticipated,} \end{cases} \quad (6)$$

so expectations of layoffs add a *firing wedge* $\beta\phi$ to the effective marginal cost of labor at $t = 0$, lowering the optimal n_0 .

Closed-form comparative statics (Cobb–Douglas). With $f(n) = n^\alpha$ and $0 < \alpha < 1$, the period-0 choice under the contraction (anticipated-firing) regime satisfies

$$n_0^C = \left(\frac{\alpha A_0}{w_0 + c_h + \beta\phi} \right)^{\frac{1}{1-\alpha}}. \quad (7)$$

At $t = 1$, the firing-region choice solves

$$n_1^- = \left(\frac{\alpha A_1}{w_1 - \phi} \right)^{\frac{1}{1-\alpha}} \quad (\text{assuming } w_1 > \phi). \quad (8)$$

The contraction regime is relevant whenever

$$\frac{A_1}{w_1 - \phi} \leq \frac{A_0}{w_0 + c_h + \beta\phi}. \quad (9)$$

Differentiating (7) yields

$$\frac{\partial n_0^C}{\partial \phi} = -\frac{\beta}{1-\alpha} \cdot \frac{n_0^C}{w_0 + c_h + \beta\phi} < 0, \quad (10)$$

which captures the *option value of restraint*: hiring less today mitigates expected firing costs tomorrow.

5.2 Embedding the Mechanism in a Task-Based Model

To connect with canonical theories of task displacement, suppose production in each period requires a unit measure of tasks $z \in [0, 1]$. Tasks $z < I_t$ are automated (performed by capital) and $z \geq I_t$ are performed by labor; write $\Gamma_t \equiv 1 - I_t$ for the labor share of tasks. Output aggregates capital- and labor-performed tasks with CES elasticity $\sigma > 1$,

$$Y_t = A_t \left[\Gamma_t^{1/\sigma} n_t^{(\sigma-1)/\sigma} + (1 - \Gamma_t)^{1/\sigma} (A^K K)^{(\sigma-1)/\sigma} \right]^{\sigma/(\sigma-1)}, \quad (11)$$

where capital K is fixed and the constants are absorbed into Ω_t below to simplify notation. The MPL implied by (11) is

$$\frac{\partial Y_t}{\partial n_t} = A_t \left[\Gamma_t^{1/\sigma} n_t^{(\sigma-1)/\sigma} + \Omega_t \right]^{1/(\sigma-1)} \Gamma_t^{1/\sigma} n_t^{-1/\sigma}. \quad (12)$$

The firm now maximizes (5) with $A_t f(n_t)$ replaced by Y_t . The optimality conditions take exactly the same form as in the baseline:

$$\frac{\partial Y_1}{\partial n_1} = \begin{cases} w_1, & n_1 \geq n_0, \\ w_1 - \phi, & n_1 < n_0, \end{cases} \quad \frac{\partial Y_0}{\partial n_0} = \begin{cases} w_0 + c_h, & \text{no firing expected,} \\ w_0 + c_h + \beta\phi, & \text{firing anticipated.} \end{cases} \quad (13)$$

Implications. Current automation lowers Γ_0 and hence the MPL of labor, reducing n_0 . Anticipated future automation (higher I_1 and lower Γ_1) makes separations more likely at $t = 1$. When the firm expects to be in the firing regime tomorrow, the period-0 condition (13) features the same firing wedge $\beta\phi$, which depresses current hiring. Larger firing costs ϕ strengthen this effect. Consequently, in equilibrium $n_0^C < n_0^E$ (where E denotes the no-firing regime), $\partial n_0 / \partial I_0 < 0$, and $\partial n_0 / \partial ({}_0I_1) < 0$.

In both formulations, the key force is the intertemporal link created by convex adjustment on the downside: the possibility of paying ϕ tomorrow raises the effective marginal cost of hiring by $\beta\phi$ today. Expectations of labor-saving progress therefore reduce current employment even before automation is realized.

6 Conclusion

This paper provides early, large-scale evidence that the diffusion of GenAI since 2023 is associated with *seniority-biased* employment effects within firms. Using résumé and job posting data linked to nearly 285,000 U.S. firms, together with a direct measure of adoption based on “GenAI integrator” vacancies, we document that GenAI adoption coincides with a pronounced decline in junior employment, while senior employment remains unchanged. Difference-in-differences, triple-difference, and staggered event-study estimates consistently point to this pattern. The decline in junior employment is concen-

trated in occupations highly exposed to GenAI, with low-exposure occupations showing no comparable change. Decomposing employment flows reveals that this effect stems primarily from a slowdown in hiring rather than increased exits or promotions.

These findings should be interpreted with caution. Adopting firms differ systematically in size, workforce composition, and industry. Although our designs account for many observable and unobservable differences by adoption, unobserved confounding factors may remain. Additionally, 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 evidence suggests that GenAI adoption may be shifting work away from entry-level tasks, potentially narrowing the bottom rungs of internal career ladders. Because early-career jobs play a central role in skill development and lifetime wage growth, such shifts could have lasting implications for inequality and mobility. These patterns raise several important questions for future research. Understanding whether the observed adjustments persist, and how firms and workers adapt through training, task design, or career development, remains an open and important area for further study.

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. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2):3–30.
- 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.
- Bick, A., Blandin, A., and Deming, D. J. (2024). The rapid adoption of generative AI. Technical report, National Bureau of Economic Research.
- 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., Krayer, 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 Board (2025). Transcript of Chair Powell’s Press Conference. <https://www.federalreserve.gov/mediacenter/files/FOMCpresconf20250917.pdf>. Federal Open Market Committee Press Conference Transcript.

- 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.
- Garicano, L. and Rayo, L. (2025). Training in the Age of AI: A Theory of Apprenticeship Viability. CEPR Discussion Paper DP20634, Centre for Economic Policy Research.
- 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.
- 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.
- 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.
- 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 Verification of the Seniority Variable

To validate the seniority variable provided by Revelio Labs, we conducted a detailed analysis of job title keywords associated with workers of different seniority levels. The main purpose of this exercise is twofold: (i) to verify that titles assigned to junior and senior workers align with intuitive expectations, and (ii) to examine whether these patterns remain stable before and after the emergence of generative AI around 2023.

Seniority patterns. We first examine the distribution of role keywords separately for junior and senior workers. The results confirm that the seniority classification is meaningful. Job titles associated with junior workers are dominated by keywords such as *assistant*, *specialist*, *technician*, and *intern*. In contrast, job titles among senior workers are concentrated around terms like *manager*, *director*, and *consultant*, which are typically used in leadership and high-responsibility positions.

Stability over time. A potential concern is that firms may have changed how they label positions after 2023—for instance, by rebranding junior roles as “specialists” or “associates” to make them appear more senior, or vice versa. To address this, we repeat the keyword frequency analysis separately for the pre-2023 and post-2023 periods. The results show that the relative prominence of key role terms remains remarkably stable over time within each seniority group. This stability suggests that firms are not systematically re-labeling positions in response to technological change, and that our seniority measure consistently reflects actual job hierarchy before and after the diffusion of GenAI technologies.

Figures [A.1–A.2](#) present the distribution of role keywords using grouped bar charts, which display the pre- and post-2023 shares side by side. Figures [A.3a–A.3d](#) complement this evidence with word clouds that visually represent the most common job titles by seniority and period.

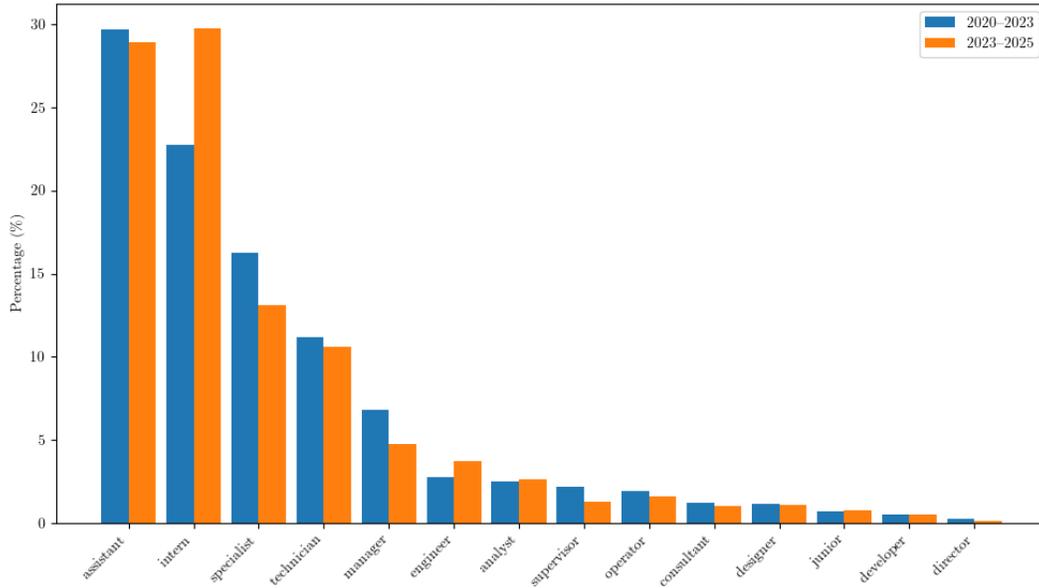


Figure A.1: Role keyword frequency—Seniority 1 and 2 (percent, pre- vs. post-2023).

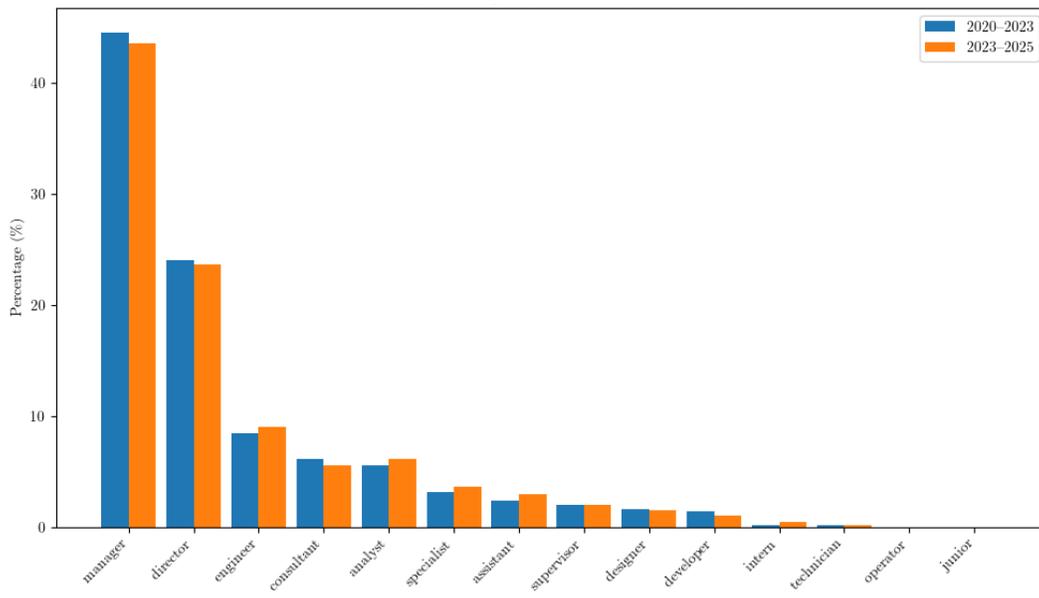
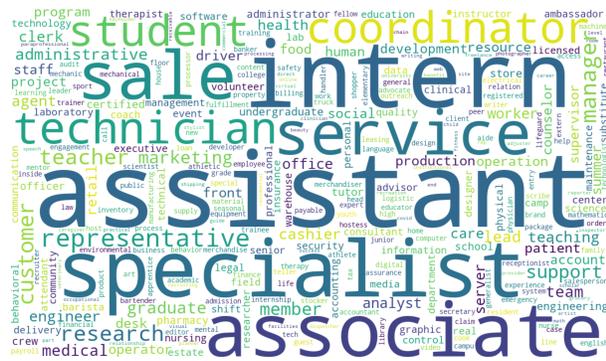
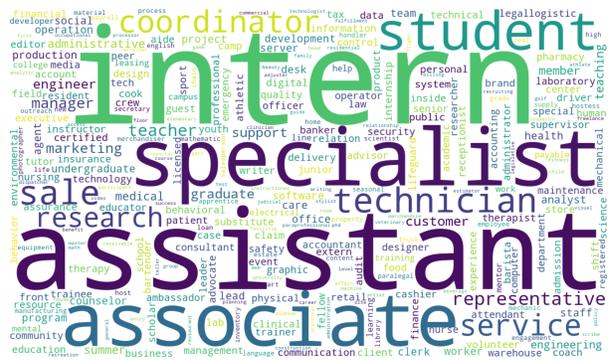


Figure A.2: Role keyword frequency—Seniority 3+ (percent, pre- vs. post-2023).



(a) Seniority 1-2, 2020-2023



(b) Seniority 1-2, 2023-2025



(c) Seniority 3+, 2020-2023



(d) Seniority 3+, 2023-2025

Figure A.3: Employment Patterns by Adoption and Seniority

Notes: Panels (a) and (b) show the average number of junior and senior workers, 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. Panel (d) presents the estimated coefficients β_j from Equation 2. The adoption measure used in this analysis is time-invariant; hence, the figures simply compare adopters and non-adopters. For results using a time-varying adoption measure, see Section 4.3. Standard errors are clustered at the firm level.

A.1.1 Prompt: Identify Job Postings for AI Integrators or Users

We use llama-3.1-8b-instant model through Groq API.

`SYSTEM_PROMPT = '''You are a classifier for job postings. Output ONLY compact JSON.`

role_type:

- *integrator: builds/operates LLM systems (RAG, embeddings/vector DB, agents, LangChain/LlamaIndex, fine-tune/adapters, serving/inference, eval/guardrails, API integration).*
- *user: mainly uses LLM tools (ChatGPT, Gemini, Copilot, etc.) without building systems.*
- *both: both apply.*
- *none: neither.*

Exclusions: not integrator if only foundation-model research, generic AI/ML, or developer roles at AI labs (e.g., OpenAI, DeepMind), or labeling/annotation.

department (choose one):

- *Technology*
- *Operations*
- *Marketing*
- *HR*

Rules:

- *If both integrator + user → role_type="both"*
- *Acronyms like "RAG" = LLM context*
- *Prefer 1 when signals appear*
- *JSON only; no prose*

Output format:

```
{  
  "integrator": 0/1,  
  "user": 0/1,  
  "role_type": "integrator"/"user"/"both"/"none",  
  "department": "Technology"/"Operations"/"Marketing"/"HR",  
  "confidence": 0.0-1.0  
}
```

'''

A.1.2 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.2 Job Postings Examples

The first two boxes (green) show two illustrative examples of job postings that were classified by the LLM as “GenAI integrator” postings. For example, the first example is a posting that explicitly includes the responsibility to “**integrate AI models** into existing systems and applications.” Moreover, the job title—*GenAI Developer Consultant*—closely fits our notion of a GenAI integrator.

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: 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.

On the other hand, the red boxes show a posting that, despite containing the related keywords *Gen-*

erative AI of Large Language Model, are not about integrating GenAI into workflows. For example, the first example highlights the value of an LLM-based classifier that goes beyond simple keyword search. Although the posting is from a GenAI company, it describes a *customer support* role unrelated to integrating AI into workflows. The model correctly classifies it as a non-integrator position. See Appendix A.2 for additional examples.

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.

...

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 Descriptive Statistics

Sectoral Distribution of AI Adopters: We provide here more detailed evidence on the sectoral distribution of GenAI adopting firms. Figure A.4a documents the share of firms in each major sector that have adopted AI, while Figure A.4b 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.

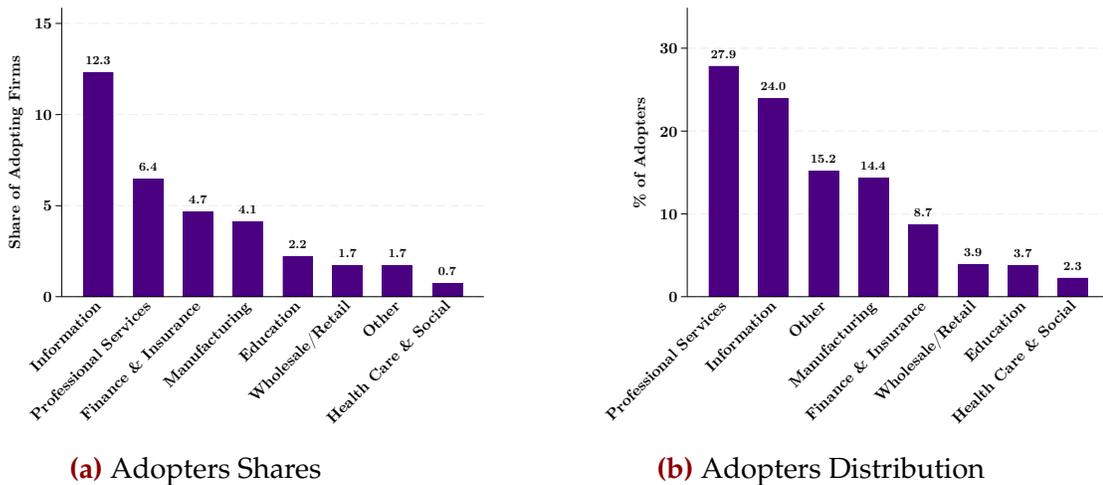


Figure A.4: Adoption Distribution Over Time

Notes: Panel (a) reports the share of adoption within each sector, while Panel (b) reports the distribution of adopting firms across sector. Sectors are: Manufacturing (3), Wholesale/Retail (4), Information (51), Finance and Insurance (52), Professional Services (54), Education (61), and Health Care and Social Assistance (62).

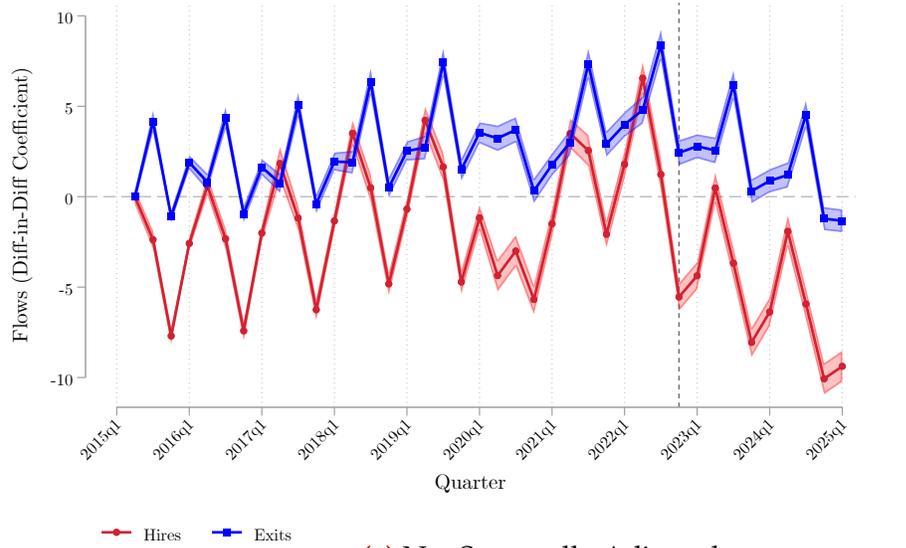
A.4 Most Common High- and Low-Exposed Occupations by Industry

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

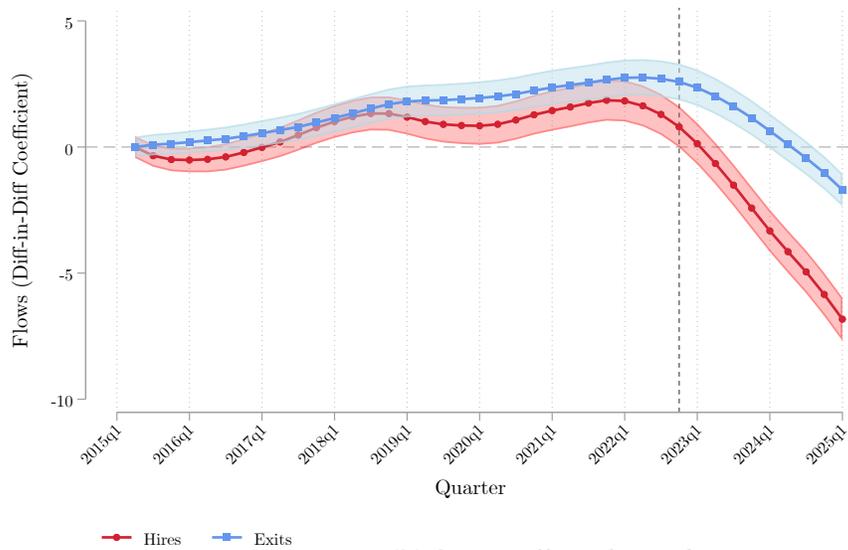
NAICS	Low Exposure (Top 5)	High Exposure (Top 5)
3 <i>Manufacturing</i>	<ul style="list-style-type: none"> – Door-to-Door Sales Workers, News and Street Vendors, and Related Workers (5.4%) – Retail Salespersons (2.2%) – Maintenance Workers, Machinery (1.8%) – Machinists (1.4%) – Biofuels Processing Technicians (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.5 DiD for Hires and Exits—Time Series (Juniors)

$$y_{it} = + \sum_{j=2015Q2}^{2025Q1} \beta_j \mathbf{1}\{t = j\} \times \text{Adopt}_i + \delta_t + \gamma_i + \varepsilon_{it}, \quad (14)$$



(a) Not Seasonally Adjusted



(b) Seasonally Adjusted

Figure A.5: DiD for Hires and Exits—Time Series (Juniors)

Notes: This figure plots β_j from Equation 14. For seasonal adjustment, the coefficients are normalized to zero at 2015Q2 and are smoothed using LOWESS (bandwidth = 0.5) before plotting.

A.6 Predicted Salary and Exposure by Educational Background

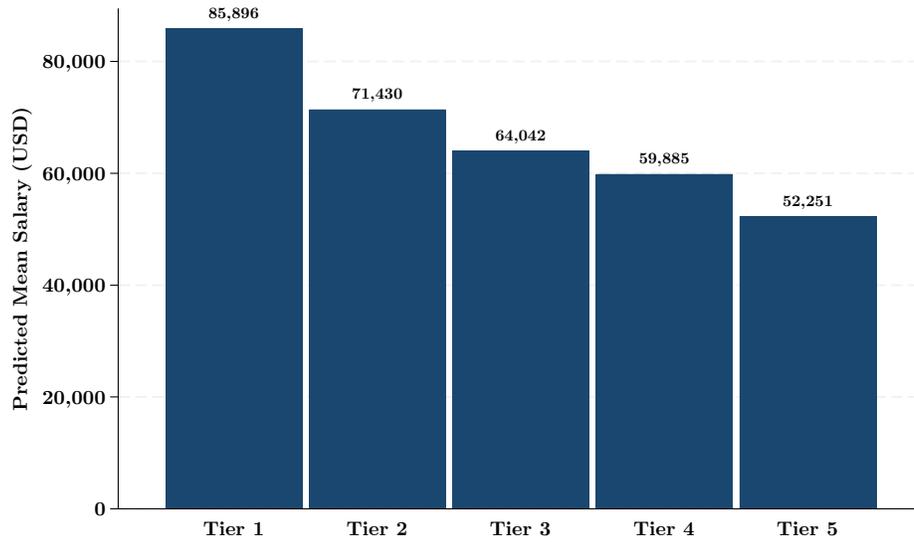


Figure A.6: Predicted Salary by School Quality (Juniors, 2022)

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

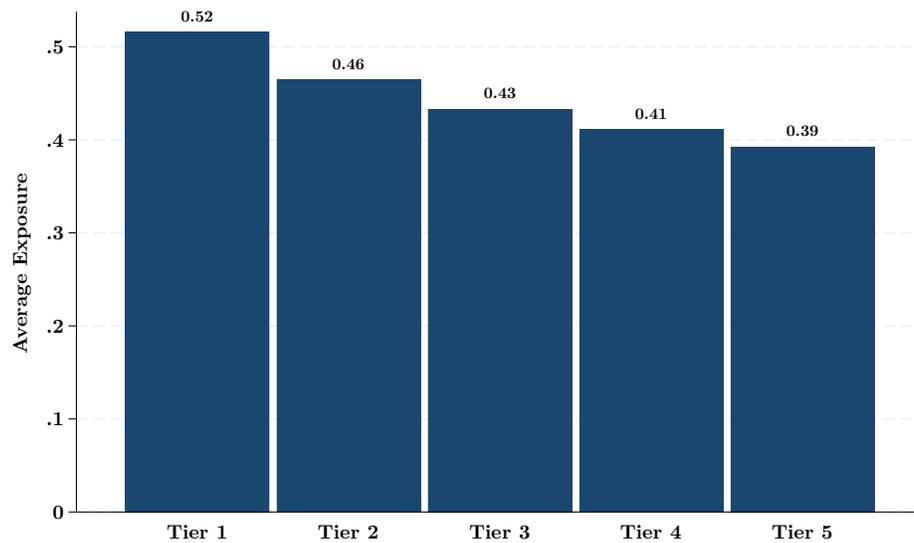


Figure A.7: Average Exposure by School Quality (Juniors, 2022)

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