

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: May 25, 2026

Abstract

We study whether generative AI (GenAI) constitutes *seniority-biased technological change*, disproportionately reducing demand for junior workers. We develop a conceptual framework in which GenAI adoption reduces junior labor demand through task displacement and labor-saving productivity gains. We test the framework’s mechanisms and implications using U.S. résumé data covering 65 million workers at more than 280,000 firms (2015–2025), allowing us to 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 trends remain largely unchanged. This decline is concentrated in occupations most exposed to GenAI, and within occupations, GenAI-exposed tasks are increasingly removed from junior job postings. The decline is driven primarily by slower hiring rather than increased separations.

*We are extremely grateful to Lawrence Katz, Jesse Shapiro, and Ludwig Straub 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 Heffetz, David Lagakos, Amanda Pallais, Santiago Medina, and Austin Zheng for their invaluable discussions and advice throughout various stages of this work. We additionally thank Netanel Ben-Porath, Shotaro Beppu, Martin Bernstein, Olena Bogdan, Bharat Chandar, Ruyuan 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 *Macro Lunch* at Harvard, *Digital Economy Lab* seminar series at Stanford University, the *Economics of AI Reading Group* at MIT, and the *Midwest Economics Association Annual Meetings 2026* 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 linked to GenAI adoption, drawing on insights 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. At the September 2025 FOMC press conference, Chair Powell was asked about the labor market impact of AI 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 also linked the diffusion of GenAI to the recent rise in unemployment among college graduates ([The New York Times, 2025a](#); [The Atlantic, 2025](#); [Forbes, 2025](#)). Others, however, dispute this interpretation, pointing instead to alternative factors such as macroeconomic uncertainty, post-pandemic adjustment, 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 and mobility ([Deming, 2023](#); [Guvonen 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

We begin with a simple conceptual framework, building on [Acemoglu and Autor \(2011\)](#) and [Acemoglu and Restrepo \(2018\)](#), to analyze why GenAI adoption may disproportionately reduce demand for junior workers and to clarify the mechanisms underlying this pattern. The framework is motivated by two broad patterns in the emerging evidence: micro-level experimental studies consistently find that GenAI disproportionately increases the productivity of less-experienced and less-skilled workers in a given task ([Brynjolfsson et al., 2025b](#); [Noy and Zhang, 2023](#); [Dell’Acqua et al., 2023](#); [Cui et al., 2025](#); [Cruces, 2026](#)), whereas labor market evidence, including our results, links GenAI adoption and exposure to declines in junior employment ([Brynjolfsson et al., 2025a](#); [Azar et al., 2025](#); [Klein Teeselink et al., 2025](#)). The framework reconciles these patterns.

In the framework, junior and senior workers sort into a continuum of tasks according to comparative advantage. GenAI both automates low-complexity tasks previously performed by juniors and augments productivity on remaining tasks, with larger gains for juniors. The change in junior labor demand decomposes into three forces: a *displacement effect*, capturing the loss of automated tasks; a *productivity effect*, whereby higher productivity reduces labor required per task; and a *reallocation effect*, whereby juniors expand into some previously senior tasks as their relative productivity rises. When tasks are complementary in production, both displacement and productivity effects reduce junior labor demand, while reallocation is the only offsetting force and depends on the extent of differential augmentation. In the framework, seniors have a different exposure structure: their tasks are not directly displaced, but some are reallocated to juniors via the reallocation channel. Whether junior employment declines more than senior employment depends on three primitives: the scope of automation, the gap in augmentation across worker types, and the steepness of comparative advantage near the task boundary.

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.

We then move to our main empirical analysis, drawing on a dataset that combines LinkedIn résumés and job-posting data from Revelio Labs (Revelio Labs, 2025). The dataset covers more than 280,000 U.S. firms, more than 160 million employment spells from roughly 66 million unique workers between 2015 and 2025, and almost 200 million job postings. A key advantage of this 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.⁴

We identify GenAI adoption by detecting job postings that explicitly recruit “GenAI integrator” roles. The method follows the approach of Hampole et al. (2025) and proceeds in two steps: first flagging postings with GenAI-related keywords, then using a large language model to determine whether the posting reflects a genuine integrator position—one dedicated to implementing or operating GenAI technology in the firm’s workflow. A firm is classified as an adopter if it has posted at least one such vacancy, thereby capturing firms that have actively initiated the integration of GenAI into their operations.⁵

By this definition, 10,433 firms in our sample adopted GenAI by March 2025. Although adopters represent only about 3.7 percent of the 281,000 firms in our sample, they are substantially larger on average and account for roughly 16 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 ChatGPT in late 2022. We validate this firm-level classification on the worker side: using a list of GenAI keywords applied to position descriptions, we find a sharp, differential rise in AI mentions at adopting firms after the launch of ChatGPT, supporting the interpretation that our adoption measure captures genuine changes in workplace technology use.

The main empirical exercise of the paper is to test the framework’s aggregate implication: that GenAI adoption disproportionately reduces junior labor demand relative to

⁴We subject the Revelio data and our key variables to several validation exercises (Appendices A.5 and A.6), showing that (i) the LinkedIn sample’s composition relative to QCEW, OEWS, and QWI benchmarks is stable across the GenAI diffusion period, and (ii) the seniority classification shows no meaningful differential drift, based on both job-title keyword distributions and worker experience profiles.

⁵Our 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.

senior labor demand. Before turning to this aggregate test, we provide direct empirical evidence on the framework’s underlying channels. We use a large language model to match a representative sample of our own job postings to standardized O*NET tasks, and find that AI-exposed tasks become disproportionately less likely to appear in junior task bundles at adopting firms after 2023, with no similar pattern for seniors—direct evidence of the displacement channel. By contrast, ex-ante senior-intensive tasks do not measurably migrate into junior bundles, suggesting limited support for the reallocation channel. We complement these tests with external task-level data from the Anthropic Economic Index (Appel et al., 2026), documenting that real-world GenAI usage tilts toward outright automation for tasks characteristic of junior occupations and toward augmentation for senior occupations, consistent with the framework’s assumption that junior tasks are disproportionately displaceable.

We then turn to the aggregate response. We compare adopting and non-adopting firms using a difference-in-differences (DiD) design to track junior and senior employment quarterly. From 2015 to 2022, adopters and non-adopters followed parallel trends in junior employment. However, beginning in 2023Q1—coinciding with the sharp increase in GenAI adoption—junior employment in adopting firms starts to decline 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 start to decline sharply starting in 2023Q1, reaching roughly a 9 percent drop after six quarters.

We then examine which types of jobs experienced the most pronounced declines, exploring heterogeneity by occupational exposure to GenAI based on the widely used exposure measure from Eloundou et al. (2024). Specifically, we estimate a triple-difference

specification that compares employment changes in high- versus low-exposure occupations, within adopting versus non-adopting firms, estimated separately for juniors and seniors. We find that junior employment in high-exposure occupations contracts by roughly 7 percent relative to low-exposure occupations between 2022Q4 and 2025Q1, while the corresponding senior series continues to rise throughout the same period. This divergence is consistent with GenAI substituting for exposed junior tasks.⁶

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 two 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.

We then address potential endogeneity concerns. Adopting and non-adopting firms differ systematically, which could bias our estimates if these differences are correlated with their responsiveness to other shocks affecting employment, such as the 2022–2023 monetary tightening cycle or post-COVID boom-and-bust dynamics in the technology sector. We provide several pieces of evidence that such shocks do not drive our results: (i) including industry-by-seniority-by-time fixed effects, which absorb sector-specific shocks, does not affect the junior decline; (ii) no corresponding pre-adoption trends appear in the event-study estimates; and (iii) a formal test of firms' differential sensitivity to monetary policy, in the spirit of [Ottonello and Winberry \(2020\)](#) and using the high-frequency identified monetary-policy surprises of [Jarocinski and Karadi \(2020\)](#), finds no evidence that junior workers at adopter firms respond differently to monetary shocks, consistent with the

⁶Supplemental Appendix [A.19](#) reports a complementary heterogeneity exercise by juniors' educational background. Measuring school quality by the prestige of juniors' alma mater, we find a U-shaped pattern: juniors from mid-tier universities experienced the largest relative declines, while those from the most and least prestigious institutions saw smaller reductions.

literature documenting that larger firms—which are more likely to adopt GenAI—tend to be less responsive to monetary policy (Chodorow-Reich, 2014; Gertler and Gilchrist, 1994). Beyond labor-demand-side explanations, we also consider potential supply-side channels—for example, if negative shocks to junior labor availability prompted adoption. However, the parallel pre-trends are inconsistent with this mechanism. Moreover, an analysis of job postings—an indicator more directly reflecting labor demand—shows that the decline in junior employment coincides with reduced postings, pointing to a contraction in labor demand rather than constraints 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 rose by a small but statistically significant amount, though the magnitude is economically small relative to the hiring channel.

Taken together, these results indicate that GenAI adoption is associated with reduced junior employment. However, the relatively early and pronounced declines—starting to emerge soon after the release of ChatGPT—may seem surprising, as automation impacts typically materialize more slowly. This suggests that the decline may not reflect immediate task automation but rather forward-looking adjustments by firms: the rapid diffusion of GenAI may have shifted firms’ expectations, leading them to scale back hiring for roles they predict will be automated in the near future, consistent with broader evidence on asymmetric employment adjustment (Ilut et al., 2018). We formalize this mechanism in a simple dynamic model in which expectations of *future* automation—combined with labor-adjustment costs—lead firms to reduce hiring *today* (Supplemental Appendix A.1). Several patterns are consistent with this interpretation: GenAI adoption accelerated sharply in early 2023; mentions of “AI” in firms’ earnings calls began rising in 2022Q4; prominent U.S. companies publicly adjusted hiring practices in early 2023, explicitly attributing these changes to GenAI; and the junior employment decline is driven primarily by slower hiring rather than increased separations. While these patterns align with forward-looking adjustments, our data do not allow us to test this channel di-

rectly. One potential implication is that if firms overestimate the near-term automation capabilities of GenAI, the resulting decline in junior employment may ultimately prove temporary.

We conclude by noting several limitations. 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, our definition can be thought of 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 estimates toward zero. However, silent adoption may have distinct implications for employment dynamics from the deliberate, firm-level adoption we capture.⁷ An additional caveat is that, 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 develops the conceptual framework. Section 4 describes the data and descriptive patterns. Section 5 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 GenAI.

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

⁷We partly address this concern in Supplemental Appendix A.8 by constructing an augmented adoption measure that combines integrator postings with worker position descriptions classified as integrator roles; the headline result is robust to this alternative definition.

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.⁸

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’ productivity gains offset job losses by expanding employment elsewhere, resulting in muted net changes in total headcount.⁹

⁸While 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).

⁹Two 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, sug-

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 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 about 13 percent relative to trend, while more experienced workers in those occupations saw stable or rising employment. [Simon \(2025\)](#) documents 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,

gesting 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.

while senior roles in those same occupations increase 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. More recently, [Lodefalk et al. \(2026\)](#) use Swedish job-posting and register data in an employer-level difference-in-differences design and find a widening age gradient in hiring: since the release of ChatGPT, employment of workers aged 22–25 in high-exposure occupations has fallen by about 5.5 percent relative to less-exposed occupations within the same employers, while workers over 50 have seen a slight increase. [Crane and Soto \(2026\)](#) focus on computer-programming occupations—among the most AI-exposed—and, using a within- versus between-industry decomposition of CPS data, find that coder employment growth is roughly 3 percentage points lower post-ChatGPT than pre-ChatGPT, even after controlling for industry-level shocks.¹⁰ By contrast, [Chandar \(2025\)](#) and [Financial Times \(2025\)](#) 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 Conceptual Framework

3.1 Overview

This section develops a simple formal framework to analyze the “seniority-biased” labor demand effects of GenAI. The objective is to explain why GenAI adoption may disproportionately reduce demand for junior workers and to clarify potential mechanisms under-

¹⁰Motivated by this body of evidence, [Friebel et al. \(2026\)](#) model the dynamic implications of reduced junior hiring for the short- and long-run structure of firms’ workforce. Relatedly, [Ide \(2025\)](#) develop an overlapping-generations model in which AI-driven entry-level automation weakens the intergenerational transmission of tacit knowledge from experts to novices, potentially reducing long-run growth.

lying this pattern.¹¹ The framework is motivated by two broad patterns in the emerging evidence: micro-level experimental studies consistently find that GenAI disproportionately increases the productivity of less-experienced workers (Brynjolfsson et al., 2025b; Noy and Zhang, 2023; Dell’Acqua et al., 2023; Cui et al., 2025; Cruces, 2026), whereas labor market evidence, including our results, suggests that GenAI adoption and exposure are associated with declines in junior employment (Brynjolfsson et al., 2025a; Azar et al., 2025; Klein Teeselink et al., 2025).

The model examines how GenAI adoption by firms reshapes the composition of labor demand between junior and senior workers. We adopt a partial equilibrium approach, taking output prices and wages as given and treating GenAI adoption as exogenous. The framework highlights three mechanisms. First, a *displacement effect*: although GenAI may help juniors more within any given task (as suggested by the experimental evidence), juniors and seniors specialize in different tasks, and the routine tasks performed by juniors are likely most susceptible to automation. This channel unambiguously reduces demand for junior labor and is not captured by experimental settings that hold task assignment fixed. Second, a *productivity effect*: even for tasks that remain, higher within-task productivity can reduce the number of juniors required. When tasks are complementary in production, firms demand a relatively fixed mix of task outputs, so higher productivity per junior implies fewer juniors are needed. The sign of this channel depends on the elasticity of substitution across tasks—an intuition related to Bessen (2018). Third, a *reallocation effect*: because GenAI augments juniors more than seniors, juniors may expand into some previously-senior tasks, which works to *increase* junior labor demand. The magnitude of this positive force depends on the differential productivity gain across worker types and the steepness of comparative advantage near the boundary. The main result is that when tasks are complementary or Cobb-Douglas ($\sigma \leq 1$) and the junior augmentation advantage is not too large, the net effect on junior labor demand is negative.

¹¹This framework explains the *incidence* of GenAI across seniority levels—why juniors may bear the labor demand cost while seniors do not. It is distinct from, and complementary to, the forward-looking hiring model in Supplemental Appendix A.1, which explains the *timing* of the employment response—why firms adjust junior hiring early, before automation is fully realized.

3.2 Setup

We build on the task-based framework of [Acemoglu and Autor \(2011\)](#) and [Acemoglu and Restrepo \(2018\)](#). We consider a short-run partial-equilibrium setting in which firm output Y and wages are held fixed. The goal is to isolate how adoption of GenAI by a firm changes the composition of labor demand across junior and senior workers, holding firm scale fixed, rather than to model total employment in general equilibrium. A firm produces Y by combining a continuum of tasks indexed by $x \in [0, 1]$ with a CES technology:

$$Y = \left(\int_0^1 y(x)^{\frac{\sigma-1}{\sigma}} dx \right)^{\frac{\sigma}{\sigma-1}}, \quad \sigma > 0, \sigma \neq 1, \quad (1)$$

where $y(x)$ is the output of task x and σ is the elasticity of substitution across tasks (the case $\sigma = 1$ is the Cobb-Douglas limit). The firm employs junior workers L_J and senior workers L_S . Tasks are ordered by complexity, reflecting the hierarchical structure of knowledge-based organizations ([Garicano, 2000](#); [Garicano and Rossi-Hansberg, 2006](#)). When task x is performed by a worker of type $k \in \{J, S\}$, output is

$$y(x) = \gamma_k(x) \cdot l_k(x), \quad (2)$$

where $\gamma_k(x)$ is type- k productivity and $l_k(x)$ is the labor allocated. We assume:

Assumption 1 (Comparative advantage). *The productivity ratio $\gamma_S(x)/\gamma_J(x)$ is continuously differentiable and strictly increasing in x , with $\frac{d}{dx} \ln\left(\frac{\gamma_S(x)}{\gamma_J(x)}\right) > 0$ for all $x \in (0, 1)$.*

Under cost minimization, there exists a threshold \hat{x} such that juniors perform all tasks $x < \hat{x}$ and seniors perform all tasks $x \geq \hat{x}$.¹²

GenAI affects the production technology through two channels. First, *task displacement*: GenAI automates all tasks $x \in [0, \alpha]$ for some exogenous threshold $\alpha \in (0, \hat{x})$, removing them from human labor; automated tasks remain in the CES aggregator and are performed by GenAI at exogenous unit cost $c_A(x) < w_J/\gamma_J(x)$ for $x \in [0, \alpha]$.¹³ Second,

¹²The threshold \hat{x} satisfies $\gamma_J(\hat{x})/w_J = \gamma_S(\hat{x})/w_S$, where w_J and w_S are wages.

¹³We take α as given rather than modeling the adoption decision. The framework is a model of AI's incidence across worker types, not of firms' adoption choices. The assumption $\alpha < \hat{x}$ is a simplifying benchmark reflecting the empirically relevant case in which the bulk of automated tasks are routine tasks that juniors perform ([Eloundou et al., 2024](#)). If some senior tasks are also automated, seniors face a dis-

within-task augmentation: on tasks that remain human-performed, GenAI augments productivity to $\tilde{\gamma}_k(x) = (1 + \delta_k) \cdot \gamma_k(x)$, where $\delta_J \geq \delta_S \geq 0$ captures the experimental finding that GenAI compresses within-task performance gaps by boosting less-experienced workers more.¹⁴

3.3 Effect of GenAI on Junior Labor Demand

After the firm adopts GenAI, juniors perform tasks (α, \hat{x}') and seniors perform tasks $[\hat{x}', 1]$, where \hat{x}' is the post-AI allocation threshold, determined by the equal-cost condition under augmented productivities (see Supplemental Appendix A.2 for the derivation):

$$\frac{\gamma_S(\hat{x}')}{\gamma_J(\hat{x}')} = \frac{1 + \delta_J}{1 + \delta_S} \cdot \frac{\gamma_S(\hat{x})}{\gamma_J(\hat{x})}. \quad (3)$$

Since we assume that GenAI augments junior productivity at least as much as senior productivity ($\delta_J \geq \delta_S$), the right-hand side weakly exceeds $\gamma_S(\hat{x})/\gamma_J(\hat{x})$, and because $\gamma_S(x)/\gamma_J(x)$ is strictly increasing (Assumption 1), it follows that $\hat{x}' \geq \hat{x}$: the junior task set expands among surviving tasks. When $\delta_J = \delta_S$, the threshold is unchanged exactly.

The change in junior labor demand can be decomposed into three components. We denote by $l_k^{*\text{pre}}(x)$ and $l_k^{*\text{post}}(x)$ the cost-minimizing labor allocations before and after AI adoption. Total junior labor demand is

$$L_J^{*\text{pre}} = \int_0^{\hat{x}} l_J^{*\text{pre}}(x) dx, \quad L_J^{*\text{post}} = \int_{\alpha}^{\hat{x}'} l_J^{*\text{post}}(x) dx, \quad (4)$$

and the change is $\Delta L_J^* \equiv L_J^{*\text{post}} - L_J^{*\text{pre}}$.

Proposition 1 (Decomposition). *The change in junior labor demand following GenAI adoption*

placement channel too, but the asymmetry remains as long as a larger share of automated tasks falls in the junior range. The condition $c_A(x) < w_J/\gamma_J(x)$ ensures that automation is cost-reducing: only tasks for which GenAI is cheaper than human labor are automated.

¹⁴The key results extend to task-specific augmentation $\delta_k(x)$ as long as the post-AI productivity ratio $\tilde{\gamma}_S(x)/\tilde{\gamma}_J(x)$ remains strictly monotone increasing, so that a unique threshold continues to exist.

is:

$$\Delta L_J^* = \underbrace{-\int_0^\alpha l_J^{*\text{pre}}(x) dx}_{\text{displacement } (<0)} + \underbrace{\int_\alpha^{\hat{x}} [l_J^{*\text{post}}(x) - l_J^{*\text{pre}}(x)] dx}_{\text{productivity } (\leq 0)} + \underbrace{\int_{\hat{x}}^{\hat{x}'} l_J^{*\text{post}}(x) dx}_{\text{reallocation } (\geq 0)}. \quad (5)$$

Figure 1 illustrates the task partition before and after GenAI adoption. The three components of the decomposition have distinct determinants and signs. First, the *displacement effect* is the direct loss of automated junior tasks. Its magnitude is increasing in α —the extent of automation—and it is unambiguously negative.

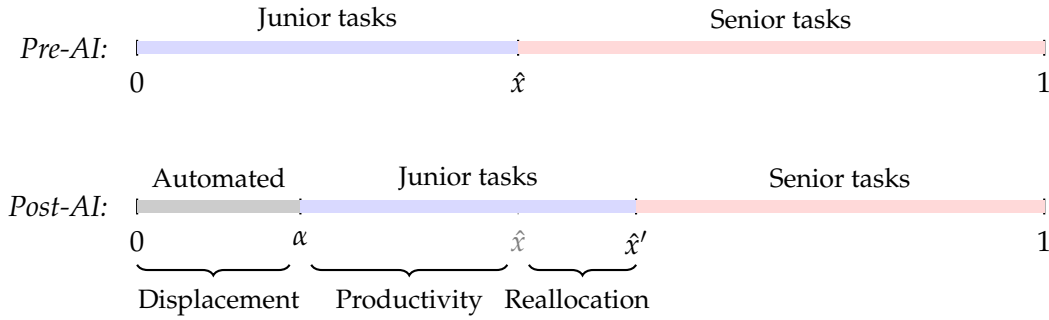


Figure 1: Task Partition Before and After GenAI Adoption

Notes: The figure is schematic and drawn for the maintained case $\alpha < \hat{x}$ (AI automates only junior tasks). The top line shows the pre-AI task partition: juniors perform tasks $x < \hat{x}$ and seniors perform tasks $x \geq \hat{x}$. The bottom line shows the post-AI partition: tasks $x \leq \alpha$ are automated, juniors perform (α, \hat{x}') , and seniors perform $[\hat{x}', 1]$. The three components of Proposition 1 correspond to: *displacement*—junior tasks lost to automation $[0, \alpha]$; *productivity*—changes in labor demand on surviving junior tasks (α, \hat{x}) ; and *reallocation*—tasks shifted from seniors to juniors $[\hat{x}, \hat{x}']$.

Second, the *productivity effect* captures the change in labor needed per surviving task. Under the partial-equilibrium assumption that aggregate output Y and wages are fixed, it reflects two channels (see Supplemental Appendix A.2 for the formal derivation). First, a *cost-index channel*: GenAI reduces the CES cost index from P^{pre} to $P^{\text{post}} \leq P^{\text{pre}}$, which lowers labor demand on all surviving tasks regardless of σ . This channel is larger the greater are α (scope of displacement) and δ_J, δ_S (augmentation). Second, a *within-task augmentation channel*, whose sign depends on the elasticity of substitution σ . When $\sigma < 1$, tasks are complementary in production, so the firm requires a roughly fixed mix of task outputs; making juniors more productive per task means fewer are needed, reinforcing the cost-index channel. When $\sigma = 1$ (Cobb-Douglas), only the cost-index channel operates and the productivity effect is still negative. Thus, if $\sigma \leq 1$, the within-task augmentation

channel is non-positive and the cost-index channel is negative, making the productivity effect unambiguously negative.¹⁵ When $\sigma > 1$, within-task augmentation tends to raise labor demand and can offset the cost-index channel, making the net productivity effect ambiguous.

Finally, the *reallocation effect* captures juniors' expanded task set due to differential augmentation. The number of tasks reallocated to juniors depends on two forces (see Supplemental Appendix A.2 for the local approximation). It is increasing in the differential augmentation $\delta_J - \delta_S$: the larger the gap in productivity gains across worker types, the more tasks shift from seniors to juniors. It is decreasing in the steepness of comparative advantage at the boundary—i.e., how rapidly the productivity ratio $\gamma_S(x)/\gamma_J(x)$ rises at \hat{x} . When this slope is steep, even a large augmentation advantage shifts few tasks. When $\delta_J = \delta_S$, the threshold does not shift and the reallocation effect vanishes entirely.

When tasks are complementary or Cobb-Douglas ($\sigma \leq 1$), two of the three forces reduce junior labor demand: displacement is unambiguously negative and the productivity effect is negative under any nondegenerate GenAI adoption ($\alpha > 0$ or $\delta_J > 0$). The only countervailing force—reallocation—depends on the differential augmentation $\delta_J - \delta_S$ and vanishes entirely when GenAI augments both worker types equally. Junior labor demand therefore declines whenever the reallocation effect is sufficiently small relative to the combined displacement and productivity effects—for example, when comparative advantage is steep near \hat{x} or when the differential augmentation $\delta_J - \delta_S$ is small.

An analogous decomposition holds for seniors. Under the maintained assumption $\alpha < \hat{x}$, no senior tasks are automated, so the displacement effect is zero. The reallocation effect now runs in the opposite direction—seniors lose tasks near \hat{x} to juniors—and the productivity effect operates on surviving senior tasks $[\hat{x}', 1]$. Formally,

$$\Delta L_S^* = \underbrace{0}_{\text{displacement}} + \underbrace{-\int_{\hat{x}}^{\hat{x}'} l_S^{*\text{pre}}(x) dx}_{\text{reallocation } (\leq 0)} + \underbrace{\int_{\hat{x}'}^1 [l_S^{*\text{post}}(x) - l_S^{*\text{pre}}(x)] dx}_{\text{productivity } (\leq 0)}. \quad (6)$$

The reallocation effect is weakly negative: the magnitude of senior task losses depends on

¹⁵The formal expression, derived in Supplemental Appendix A.2, shows that the ratio of post- to pre-AI labor demand on a surviving task is $\frac{l_S^{*\text{post}}(x)}{l_S^{*\text{pre}}(x)} = \rho \cdot (1 + \delta_J)^{\sigma-1}$, where $\rho \equiv (P^{\text{post}}/P^{\text{pre}})^{\sigma} \leq 1$ captures the cost-index channel.

$\delta_J - \delta_S$ and the steepness of comparative advantage near \hat{x} , and vanishes when $\delta_J = \delta_S$. The productivity effect is non-positive when $\sigma \leq 1$, and smaller in magnitude than for juniors because $\delta_S \leq \delta_J$. Supplemental Appendix A.2 provides the formal derivation.

The two decompositions reveal a structural asymmetry: juniors uniquely face displacement, but they are also the only group with an offsetting positive force (reallocation). Seniors have neither. Whether juniors decline by more than seniors therefore depends on three primitives. First, the *scope of automation* α : a larger automation share strengthens the junior-specific displacement channel without affecting seniors directly, tilting the asymmetry toward a sharper junior decline. Second, the *differential augmentation* $\delta_J - \delta_S$: a larger gap raises both the reallocation gain for juniors and the reallocation loss for seniors, working *against* a relatively sharper junior decline. Third, the *steepness of comparative advantage* at \hat{x} : the flatter the productivity ratio $\gamma_S(x)/\gamma_J(x)$, the larger the reallocation effect on both sides. Junior labor demand thus declines by more than senior labor demand when α is sizable, $\delta_J - \delta_S$ is moderate, and comparative advantage is steep enough at the boundary that reallocation is small relative to displacement. The flat or rising senior employment we document below is also consistent with general-equilibrium responses—output expansion or wage adjustment—that lie outside the scope of this partial-equilibrium framework.

3.4 Discussion

The framework yields two sets of testable implications. The first concerns the *channels* themselves: displacement should remove AI-exposed tasks from junior bundles, and reallocation should shift senior-intensive tasks into junior bundles. Section 5.1 brings direct evidence on these two channels, complementing the experimental literature that already provides strong evidence on the productivity channel and the assumption $\delta_J \geq \delta_S$. The second is the *aggregate* implication: junior labor demand at GenAI-adopting firms should fall relative to non-adopters, with a muted or absent senior response. Sections 5.2–5.6 take this implication to the data.

Importantly, the framework abstracts from several additional channels that may further shape the demand for junior and senior workers. First, if GenAI eliminates the routine tasks that juniors effectively provide as “payment” for on-the-job training, the im-

PLICIT apprenticeship contract may break down, reducing firms’ incentives to hire junior workers (Garicano and Rayo, 2025; Ide, 2025). Second, heterogeneous adoption across workers by age or experience may directly shift relative labor demand. Third, by automating and simplifying parts of an occupation’s task bundle, GenAI can alter the expertise required to perform that occupation, thereby changing the pool of qualified workers and, in turn, the effective labor supply for junior and senior roles across occupations (Autor and Thompson, 2025; Hosseini and Lichtinger, 2026; Althoff and Reichardt, 2026). Finally, the framework holds firm output fixed and abstracts from a *scale effect*: productivity gains from GenAI may expand adopting firms, raising labor demand on surviving tasks and attenuating the absolute decline in junior employment.

4 Data and Descriptive Patterns

4.1 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.¹⁶

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).¹⁷ Supplemental Appendices A.5 and A.6 provide

¹⁶The initial data download was performed in May 2025. In April 2026, we re-downloaded the positions data to capture profile updates that users added with a lag, ensuring more complete coverage of recent employment spells.

¹⁷More details on Revelio’s seniority classification methodology are available at <https://www.data-dictionary.reveliolabs.com/methodology.html#seniority>.

comprehensive data validation; Section 4.4 below summarizes their main findings.

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 4.3.

In addition, we 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, for example, [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).

We also construct a representative sample of *tasks* from Revelio’s job postings database. From the Revelio job-postings corpus we draw a representative sample of 1,000 firms containing 355,013 postings, and map each posting’s free-text description to a subset of the 19,265 standardized O*NET tasks using a two-stage pipeline. The first stage is a TF-IDF retriever that returns the top 50 candidate O*NET tasks for each posting; the second stage is a large language model (OpenAI GPT-4o-mini) that reads the posting text together with the candidate list and selects the tasks it judges to be genuinely required. The procedure yields a panel of 2,748,186 posting-task pairs.¹⁸

Our final sample merges all U.S. positions in the Revelio Labs dataset with job postings at the firm level. The resulting dataset covers 281,111 firms that were successfully matched to both employee position data and job postings, and that were actively hiring between January 2021 and March 2025.¹⁹ For these firms, we observe 161,192,969 positions dating back to 2015 and 198,773,384 job postings since 2021, all with usable raw text descriptions.²⁰

¹⁸The extraction pipeline is motivated by [Hampole et al. \(2025\)](#), who construct firm-level AI exposure by mapping résumé text to O*NET tasks. Appendix A.12 describes the pipeline in detail, reports the exact prompt and JSON schema, and presents sanity checks on the output.

¹⁹A firm is considered active if it had at least 20 new positions starting during this period.

²⁰We 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. In addition, we exclude the approximately 800 largest firms in the Revelio data (ranked by post-2021 new hires), representing about 0.2 percent of the firms in the sample, to avoid the influence of extreme outliers on our estimates. Supplemental Appendix A.15 shows that our main DiD results remain robust when these firms are added back.

4.2 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 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.²¹

Figure 2 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. In line with these findings, Supplemental Appendix A.7 shows that the unemployment rate among recent college graduates has risen even as overall unemployment has remained largely stable.

²¹Roughly 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.

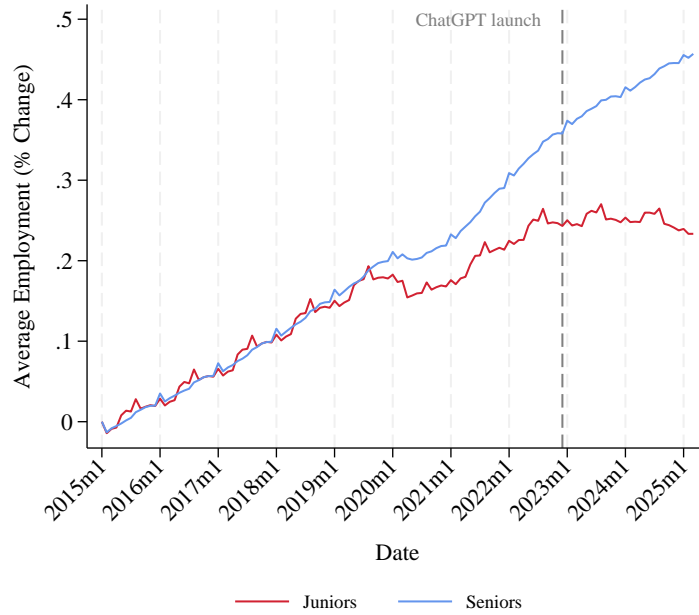


Figure 2: Log Average Employment of Juniors and Seniors in Sample Firms

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-levels, while “Senior” refers to Associate level and above (see Section 4.1 for details).

4.3 GenAI Adoption

4.3.1 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.²² Out of 198.8 million postings, 603,152 (0.30 percent) include at least one keyword. Second, we apply a large language model (LLM) classifier to this subset to 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 generic GenAI mentions in job

²²Keywords 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.

postings. (Supplemental Appendix A.9.1 provides the exact prompt used). This process yields 131,845 postings, 0.066 percent of the full corpus, classified as GenAI integrator roles. Supplemental Appendix A.10 provides a graphical overview of the procedure, and Supplemental Appendix A.11 presents illustrative examples of postings classified as integrators and non-integrators.²³

4.3.2 GenAI-Adopting Firms

We define a firm as a GenAI adopter if it has posted at least one GenAI integrator vacancy. By this criterion, 10,433 firms qualify as adopters. While they make up only 3.71 percent of the 281,111 firms in our sample, adopters are disproportionately large (see descriptive statistics below) and account for 16 percent of the employment (positions) in our dataset.

Figure 3 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.

Table 1 presents descriptive statistics for the full sample, adopters, and non-adopters. Several systematic differences stand out. Adopting firms are much larger, averaging about 546 employees compared to roughly 108 for non-adopters. Their workforces are more senior-heavy, with juniors comprising only 41 percent of employment versus 55 percent in non-adopting firms. Consistent with this pattern, 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: 47 percent of their juniors are in high-exposure occupations (vs. 22 percent at non-adopters), and 33 percent of their seniors (vs. 12 percent at non-adopters).

While non-adopters are widely dispersed across sectors, GenAI adopters are heav-

²³As shown in Supplemental Appendix A.9.1, 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.

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)	123.6 (437.2)	107.6 (368.1)	546.4 (1215.8)
Share junior employees (Entry/Junior)	0.548 (0.227)	0.553 (0.226)	0.414 (0.209)
Share senior employees (Associate+)	0.454 (0.227)	0.449 (0.226)	0.588 (0.209)
Average number of new hires (per quarter)	7.0 (25.0)	6.1 (21.6)	29.9 (65.8)
Share of new hires junior	0.667 (0.351)	0.674 (0.352)	0.496 (0.305)
Average number of separations (per quarter)	5.6 (22.2)	4.9 (19.4)	23.7 (57.0)
Share of separations junior	0.680 (0.353)	0.687 (0.353)	0.517 (0.309)
Average number of promotions for juniors (per quarter)	0.429 (1.9)	0.354 (1.5)	2.4 (6.2)
Juniors in high-exposure jobs (% of all juniors)	0.233 (0.235)	0.224 (0.229)	0.467 (0.281)
Juniors in low-exposure jobs (% of all juniors)	0.264 (0.272)	0.270 (0.273)	0.102 (0.162)
Seniors in high-exposure jobs (% of all seniors)	0.131 (0.176)	0.124 (0.170)	0.332 (0.226)
Seniors in low-exposure jobs (% of all seniors)	0.128 (0.179)	0.131 (0.180)	0.044 (0.088)
Panel B. Industry and Headquarters Location			
Share in NAICS sector 51 (Information)	0.071	0.064	0.243
Share in NAICS sector 52 (Finance and Insurance)	0.067	0.067	0.086
Share in NAICS sector 54 (Professional Services)	0.156	0.152	0.277
Share in NAICS sector 5 (Other)	0.066	0.065	0.093
Share in non-NAICS 5 sectors	0.640	0.652	0.301
HQ in California	0.138	0.135	0.203
HQ in Texas	0.075	0.075	0.060
HQ in New York	0.080	0.080	0.091
HQ in Other States	0.714	0.717	0.651
Observations	11,751,820	11,323,931	427,889
Number of firms	281,111	270,678	10,433

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, and automation exposure. Panel B reports industry and headquarters state distributions.

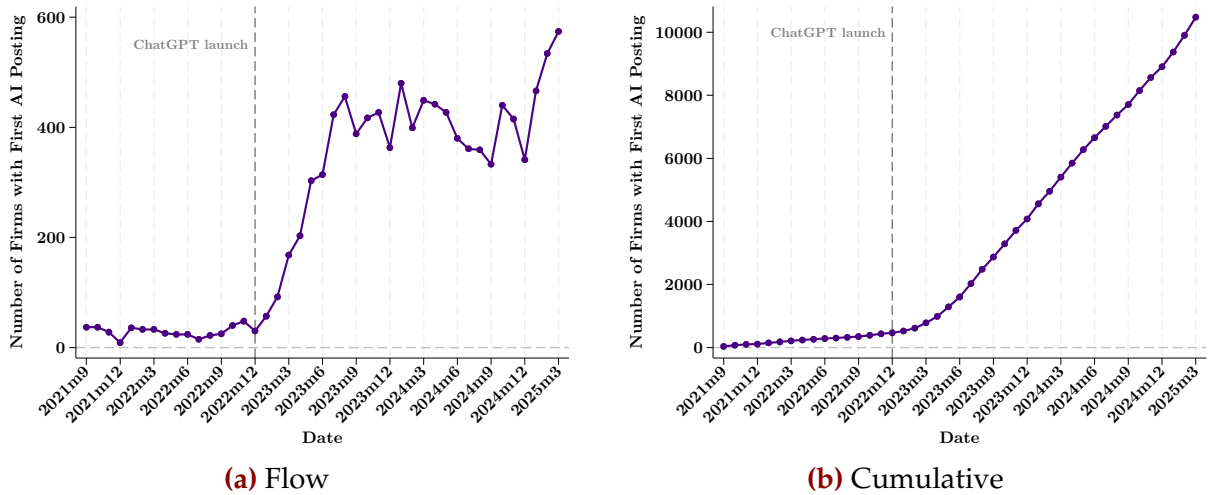


Figure 3: Number of Adopters Over Time

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.

ily concentrated in professional services (28 percent) and information (24 percent). Geographically, adopters are disproportionately headquartered in California, while being slightly less represented in Texas compared to non-adopters (Supplemental Appendix A.13 provides more details on the sectoral and geographic distribution of adopters versus non-adopters).

Taken together, the statistics depict GenAI adopters as larger, more senior-oriented firms, with stronger recruitment from exposed occupations, and greater presence in technology-intensive sectors and states.

4.3.3 Validation: GenAI Mentions in Worker Position Descriptions

As external validation of the adoption measure, we examine whether workers at adopting firms are more likely to report AI-related activities in their own LinkedIn position descriptions—an outcome that reflects actual on-the-job use rather than aspirational hiring language.²⁴ Using a list of GenAI keywords (the full list is reported in Supplemental

²⁴Position descriptions on LinkedIn are considerably less reliable than job posting texts: only about 42% of U.S. positions in our 2021–2025 window have any raw description text at all, and those that exist are often short or formulaic. For this reason, we rely on the richer and more complete raw posting data to construct the firm-level adoption measure throughout the paper. The position-description analysis presented here

Appendix A.4) applied to 18.3 million position descriptions (2021–2025), we compute for each quarter the share of descriptions mentioning GenAI keywords, separately for adopting and non-adopting firms.

Figure 4 plots these shares, normalized to zero in 2021 Q1. Prior to the launch of ChatGPT in late 2022, both series fluctuate near zero with no visible divergence. Beginning in early 2023, the adopter series rises sharply—reaching approximately 3.5 percentage points above baseline by late 2024—while the non-adopter series increases only modestly (roughly 0.7 percentage points). A difference-in-differences regression with firm and month fixed effects formalizes this pattern and yields an estimate of 1.35 percentage points (s.e. = 0.05, $p < 0.001$), and the coefficient remains essentially unchanged (1.40 percentage points) when additionally controlling for description length (full regression results are reported in Supplemental Appendix A.4). Relative to the pre-ChatGPT baseline of 3.1% at adopters, this represents a roughly 43 percent increase. This worker-side evidence supports the interpretation that our firm-level adoption classification captures genuine changes in workplace technology use.

4.4 Data Validation

We subject our data and key variables to a series of validation exercises, reported in full in the Supplemental Appendix.

First, we assess the *representativeness* of the LinkedIn/Revelio data by comparing its industry, occupation, and age distributions against official employment statistics from the QCEW, OEWS, and QWI. The data display the expected tilt of LinkedIn toward managerial, professional, and high-information occupations and industries. However, importantly, this tilt is essentially unchanged from 2021Q1 through 2025Q1, with no discernible break around the ChatGPT launch. Because the level differences are absorbed by firm and sector fixed effects in our main specifications, the differential identification of treatment effects is unaffected (see Appendix A.5 for details).

Second, we verify the *stability of the seniority classification* through two complementary exercises (Appendix A.6). We first analyze the frequency of job-title keywords (e.g., “as-

serves as a complementary validation exercise; in Supplemental Appendix A.8 we additionally use these descriptions to construct an augmented adoption measure and show that our main results are robust.

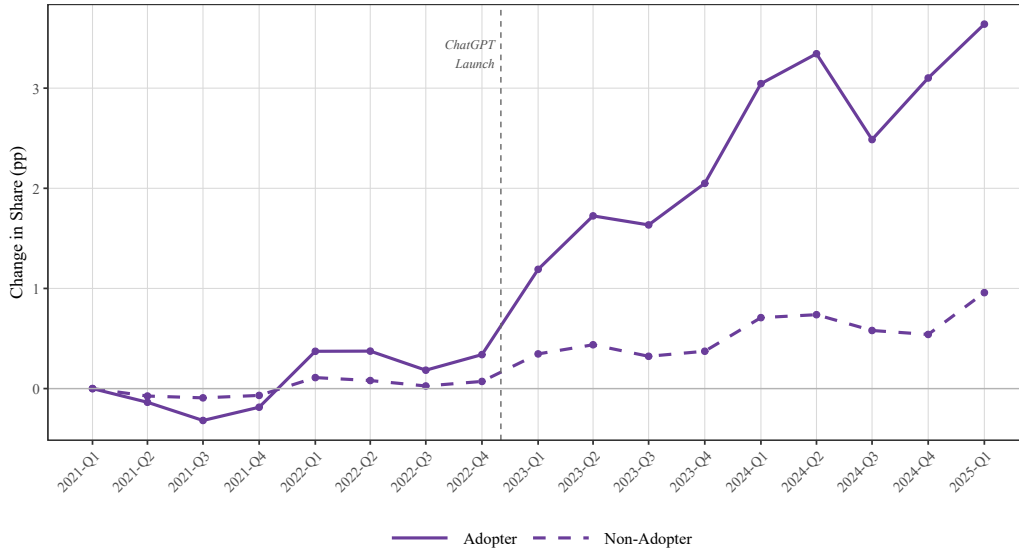


Figure 4: Change in Share of Position Descriptions Mentioning GenAI Keywords

Notes: This figure plots the share of LinkedIn position descriptions containing at least one GenAI keyword, by quarter, separately for GenAI-adopting and non-adopting firms. The sample includes 18.3 million U.S. positions from 2021 Q1 through 2025 Q1 held at firms in the main analysis sample, restricted to positions with a raw description of at least 20 characters. Both series are normalized to zero in 2021 Q1. The vertical dashed line marks the launch of ChatGPT in November 2022. Baseline levels (2021 Q1): Adopters = 2.9%, Non-Adopters = 0.8%.

sistant,” “manager,” “director”) separately for junior and senior workers, and confirm that the classification aligns well with conventional titling. Both pooled pre/post-2023 comparisons and a formal event study at adopting versus non-adopting firms show no meaningful differential shift in keyword distributions around the ChatGPT launch. We then examine the experience composition of junior and senior workers over time, separately at adopting and non-adopting firms. We find that the average labor market experience of junior and senior workers evolves in parallel at adopting and non-adopting firms; a statistically detectable but quantitatively small differential emerges for seniors in the event-study specification, but its magnitude is unlikely to meaningfully affect our employment results. Together, these exercises address the concern that our results could be driven by reclassification rather than by real changes in labor demand.

5 Results

The main empirical exercise of the paper is to test the framework’s *aggregate* implication: that GenAI adoption by firms reduces junior labor demand relative to senior labor demand. We pursue this in Sections 5.2–5.6 using the firm-level résumé and posting data, with a difference-in-differences design, a triple-difference, a staggered event study, and a flow decomposition. Before turning to the aggregate response, Section 5.1 brings direct empirical evidence on the framework’s underlying channels.

5.1 Empirical Evidence on the Framework’s Channels

Proposition 1 decomposes the change in junior labor demand into three forces: a *displacement* effect ($\alpha > 0$, junior tasks lost to automation), a within-task *productivity* effect (on the workers-per-task margin), and a *reallocation* effect ($\hat{x}' > \hat{x}$, juniors expanding into formerly-senior tasks because $\delta_j > \delta_s$). For the productivity channel, and especially for the assumption $\delta_j \geq \delta_s$, a now-large experimental literature provides strong evidence (Brynjolfsson et al., 2025b; Noy and Zhang, 2023; Dell’Acqua et al., 2023; Cui et al., 2025; Cruces, 2026). This subsection tests the remaining two channels using evidence from two sources: task-level information extracted from our own postings data, and external task-level data from the Anthropic Economic Index.

5.1.1 Task-Level Evidence: Testing the Framework’s Mechanisms in Our Data

We begin with direct evidence on the displacement and reallocation channels using task-level information extracted from our own job-postings data (Section 4.1). For each O*NET task j , firm i , 6-digit occupation c , and seniority group $s \in \{\text{junior}, \text{senior}\}$, let $\text{Prev}_{jics\tau}$ be the share of (i, c, s) postings in period τ whose extracted bundle contains task j . Splitting the sample at the ChatGPT launch ($\tau=\text{pre}$: 2021Q3–2022Q4; $\tau=\text{post}$: 2023Q1 through early April 2025, the latest postings in our extract), the dependent variable is the Davis et al. (1998) (DHS) growth rate:

$$\Delta^{\text{DHS}}\text{Prev}_{jics} = 2 \frac{\text{Prev}_{jics}^{\text{post}} - \text{Prev}_{jics}^{\text{pre}}}{\text{Prev}_{jics}^{\text{post}} + \text{Prev}_{jics}^{\text{pre}}} \times 100. \quad (7)$$

This is similar to the task-level DHS measure used by [Hampole et al. \(2025\)](#). Observations are weighted by the harmonic mean of pre- and post-period cell sizes (the standard DHS weight), and standard errors are clustered at the firm level throughout.

Test 1: Displacement. Under the maintained assumption $\alpha < \hat{x}$, automated tasks lie in the junior range, so displacement predicts that AI-exposed tasks are removed from junior bundles at adopters post-ChatGPT and not from senior ones. We operationalize a task's *displaceability* with the standardized [Eloundou et al. \(2024\)](#) GPT-4 exposure score Exp_j , and estimate, separately for juniors and seniors,

$$\Delta^{\text{DHS}}\text{Prev}_{jics} = \beta^S \text{Exp}_j \times \text{Adopt}_i + \alpha_j + \alpha_{ic} + \varepsilon_{jics}, \quad (8)$$

and a pooled triple-difference that adds the Senior_s dimension:

$$\Delta^{\text{DHS}}\text{Prev}_{jics} = \beta^A \text{Exp}_j \times \text{Adopt}_i + \beta^{AS} \text{Exp}_j \times \text{Adopt}_i \times \text{Senior}_s + \mathbf{\Gamma}'\mathbf{Z}_{jics} + \alpha_j + \alpha_{ic} + \varepsilon_{jics}, \quad (9)$$

where Adopt_i is the firm-level adopter indicator (Section 4.3), $\text{Senior}_s = 1$ for seniors, \mathbf{Z}_{jics} collects the lower-order interactions, and α_j and α_{ic} are task and firm-by-occupation fixed effects (firm only in the less demanding spec). Identification comes from within-cell, cross-task variation in exposure interacted with adoption. Under displacement, we expect $\beta^J < 0$, $\beta^S \geq 0$, and $\beta^{AS} > 0$. Table 2 reports the results.

For juniors, a one SD increase in a task's exposure is associated with a roughly four percentage-point larger DHS contraction in junior bundles at adopters relative to non-adopters, marginally significant in both specifications (columns 1–2). Scaling to the gap between the least- and most-exposed tasks (about four SDs) yields a 15–17 percentage-point differential. The coefficient for seniors (columns 3–4) is positive and insignificant: within the same firm-occupation cells, exposed tasks leave junior bundles but not senior ones. The pooled triple-difference makes the contrast precise: $\beta^A = -4.15$ ($p = 0.024$) and $\beta^{AS} = +6.87$ ($p = 0.049$) under firm fixed effects; the saturated specification delivers $\beta^{AS} = +5.83$ ($p = 0.072$). The contraction is therefore a within-occupation removal of AI-exposed tasks from junior bundles, not a compositional shift across occupations.

Table 2: Task-Level DHS Demand by GenAI Exposure and Seniority (Automation)

	Juniors only		Seniors only		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
$Exp_j \times Adopt_i$	-3.923*	-4.140*	+2.408	+1.890	-4.152**	-3.667
	(2.065)	(2.347)	(3.203)	(3.128)	(1.842)	(2.394)
$Exp_j \times Adopt_i \times Senior_s$					+6.868**	+5.826*
					(3.490)	(3.237)
Task FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes		Yes		Yes	
Firm \times Occ FE		Yes		Yes		Yes
N	208,984	208,984	402,363	402,363	611,347	611,347
Firms	923	923	969	969	997	997

Notes: Unit of observation is (task \times firm \times 6-digit SOC occupation \times seniority group). The dependent variable is $100 \times$ the Davis et al. (1998) (DHS) growth rate in the task’s share of the cell’s posting bundles, comparing post-ChatGPT (2023Q1 through early April 2025) to pre-ChatGPT (2021Q3–2022Q4). Exp_j is the standardized Eloundou et al. (2024) GPT-4 exposure of task j ; $Adopt_i$ is the firm-level GenAI-adopter indicator (Section 4.3); $Senior_s$ equals one for seniors. Columns (5)–(6) include the lower-order terms $Senior_s$, $Exp_j \times Senior_s$, and $Adopt_i \times Senior_s$ (coefficients not reported). Observations are weighted by the harmonic mean of pre- and post-period cell sizes (the DHS weight). Standard errors in parentheses are clustered at the firm level. *, **, *** denote significance at 10%, 5%, and 1%.

Test 2: Reallocation. In the framework, the reallocation channel works in the opposite direction: because GenAI augments juniors more than seniors ($\delta_j \geq \delta_s$), the post-AI allocation threshold satisfies $\hat{x}' \geq \hat{x}$, so tasks previously performed by seniors near the comparative-advantage boundary can now be absorbed by juniors. A direct test of this channel asks whether tasks that were ex-ante *senior-intensive* become more prevalent in junior bundles at adopting firms post-ChatGPT.

We operationalize the model’s notion of an “ex-ante senior task” with a task-level senior-intensity score. For each O*NET task j , define

$$SenInt_j = \frac{\#\{(\text{posting}, j) \text{ pairs in pre-period at senior level}\}}{\#\{(\text{posting}, j) \text{ pairs in pre-period}\}}, \quad (10)$$

restricted to tasks appearing in at least 20 pre-period posting-task pairs (6,985 tasks), and standardized to unit variance. $SenInt_j$ measures where, on the pre-AI seniority gradient, task j predominantly sat: tasks close to the boundary \hat{x} on the senior side score high.

We re-estimate equations (8) and (9) with $SenInt_j$ in place of Exp_j , and denote the

resulting coefficients with tildes. Under reallocation, the model predicts $\tilde{\beta}^J > 0$ (senior-intensive tasks shift *into* junior bundles at adopters), $\tilde{\beta}^S \leq 0$ (the mirror image: those same tasks leave senior bundles), and in the pooled specification $\tilde{\beta}^A > 0$ with $\tilde{\beta}^{AS} < 0$. Table 3 reports the results.

Table 3: Task-Level DHS Demand and Senior-Intensity and Seniority (Reallocation)

	Juniors only		Seniors only		Pooled	
	(1)	(2)	(3)	(4)	(5)	(6)
SenInt _j × Adopt _i	−1.114 (2.995)	−2.097 (3.384)	−0.495 (3.284)	−1.657 (3.034)	−0.012 (3.401)	−1.783 (3.407)
SenInt _j × Adopt _i × Senior _s					−0.835 (5.031)	−0.468 (4.354)
Task FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes		Yes		Yes	
Firm × Occ FE		Yes		Yes		Yes
N	189,495	189,495	377,846	377,846	567,341	567,341
Firms	923	923	969	969	997	997

Notes: Unit of observation is (task × firm × 6-digit SOC occupation × seniority group). The dependent variable is 100 × the Davis et al. (1998) (DHS) growth rate in the task’s share of the cell’s posting bundles, comparing post-ChatGPT (2023Q1 through early April 2025) to pre-ChatGPT (2021Q3–2022Q4). SenInt_j is the task-level senior-intensity score defined in equation (10), standardized to unit variance and restricted to tasks with at least 20 pre-period (posting, task) pairs (6,985 tasks). Columns (5)–(6) include the lower-order terms Senior_s, SenInt_j × Senior_s, and Adopt_i × Senior_s (coefficients not reported). Observations are weighted by the harmonic mean of pre- and post-period cell sizes (the DHS weight). Standard errors in parentheses are clustered at the firm level. *, **, *** denote significance at 10%, 5%, and 1%.

The pattern is uniformly null. None of the six coefficients on SenInt_j × Adopt_i is statistically distinguishable from zero, and all point estimates are small and—if anything—slightly *negative*, the wrong sign for the reallocation prediction. The senior–junior differential $\tilde{\beta}^{AS}$ in columns (5)–(6) is similarly close to zero (−0.84 and −0.47, $p > 0.86$). To rule out the possibility that the null is masked by mechanical correlation between SenInt_j and Exp_j, we residualize SenInt_j on Exp_j at the task level—isolating senior-intensity orthogonal to AI-displaceability—and re-estimate the most demanding firm-by-occupation specification. The coefficients are again indistinguishable from zero.

The two tests use only weakly correlated regressors ($\text{corr}(\text{Exp}_j, \text{SenInt}_j) = +0.16$) and examine opposite directions of task movement. The results support the displacement channel, which reduces demand for junior labor in the framework, but we find no support for the reallocation channel, which would raise demand for juniors. Together, these

findings support the prediction that GenAI adoption reduces demand for junior labor, consistent with the firm-level employment results we document in Sections 5.2–5.6.

5.1.2 Evidence from the Anthropic Economic Index

We complement the tests above with external evidence on the framework’s maintained assumption that the bulk of automated tasks lie in the junior range ($\alpha < \hat{x}$).²⁵ We use the Anthropic Economic Index (Appel et al., 2026), which classifies real-world interactions with a frontier large language model by the O*NET task they address and the *collaboration mode* of the interaction. Following the AEI’s methodology, we treat directive interactions and feedback loops as *displacement-type* usage, and learning, task iteration, and validation as *augmentation-type* usage. We merge these AEI measures with occupation-level seniority constructed from the Revelio data, yielding a sample of 387 occupations with sufficient AEI coverage.

We find that *displacement intensity*—the share of classified interactions accounted for by displacement-type modes—exceeds 50% for junior-type occupations but falls to 40% for senior-type occupations, a gap driven by rising task iteration and validation with seniority. The pattern is robust to controlling for broad occupation groups: with two-digit SOC fixed effects, a one-standard-deviation increase in seniority is associated with a 3.68 percentage point decline in displacement intensity ($p = 0.002$). These patterns are consistent with the framework’s assumption that junior-type tasks are disproportionately the ones GenAI can automate. Supplemental Appendix A.3 provides full details on the data, methodology, and the complete set of results.

5.2 Employment Dynamics by Adoption and Seniority

5.2.1 Descriptive Time Trends:

We begin by comparing the evolution of junior and senior employment over time in GenAI-adopting versus non-adopting firms. Figure 5a plots average junior employment

²⁵This concerns the *mode* of GenAI engagement across tasks, not the size of the within-task productivity gain: the model assumes $\delta_j \geq \delta_s$ (juniors gain at least as much productivity conditional on task survival), but implies that a larger *share* of senior-type tasks will involve augmentation-type interaction because fewer of those tasks are amenable to outright automation.

across the two groups. From 2018 through the end of 2022, junior employment in adopting and non-adopting firms 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 5b 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 2023.²⁶

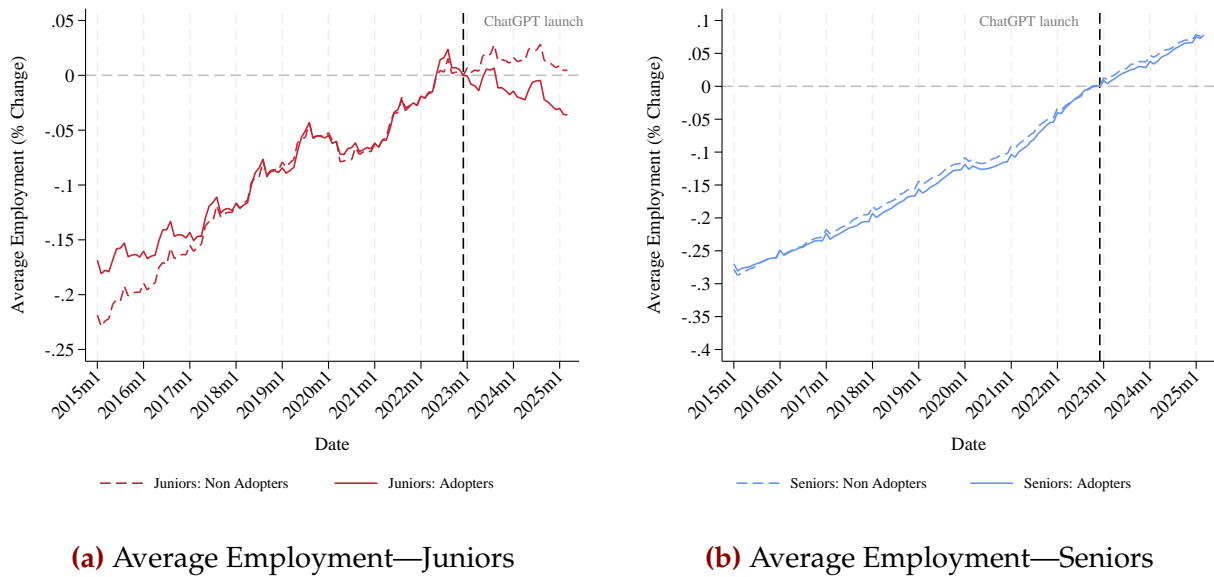


Figure 5: Average Junior and Senior Employment by Adoption Status

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). 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 5.4.

5.2.2 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

²⁶These 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.

workers:

$$\log(\text{Employment}_{it}) = \alpha + \sum_{j=2015Q2}^{2025Q1} \beta_j \mathbf{1}\{t = j\} \times \text{Adopt}_i + \delta_t + \text{Adopt}_i + \varepsilon_{it}, \quad (11)$$

where the dependent variable $\log(\text{Employment}_{it})$ denotes the log employment of junior (or senior) workers at firm i in quarter t .²⁷ 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.²⁸ 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 6 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.

5.2.3 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:

²⁷For computational efficiency we aggregate our panel data to quarterly frequency in all analyses.

²⁸Note 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 4.3 for more details). In Section 5.4, we relax this definition by exploiting variation in the timing of adoption across firms using a staggered event-study design.

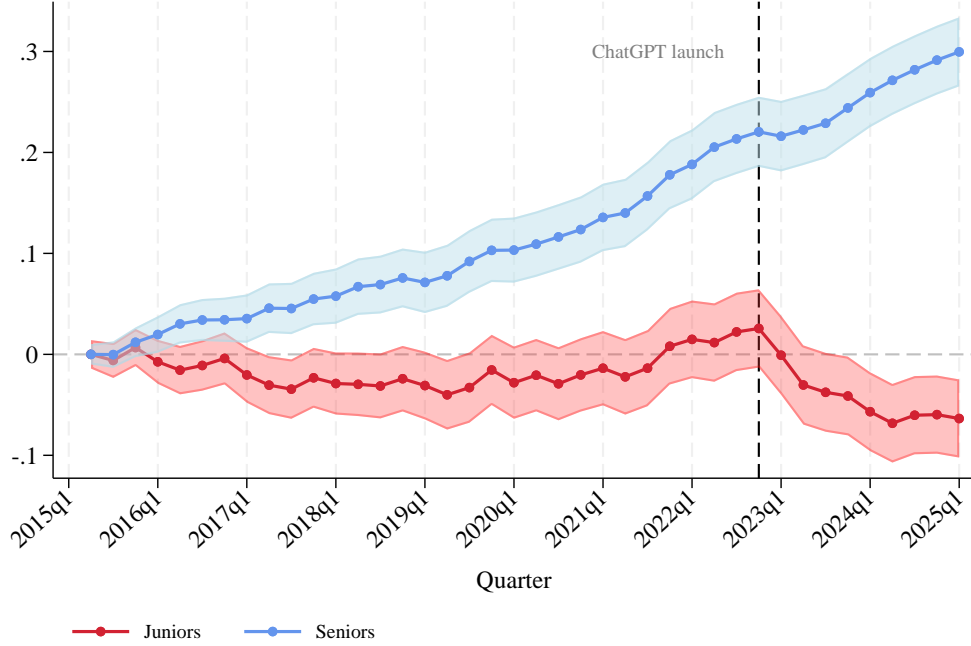


Figure 6: Difference-in-Differences Estimates by Seniority

Notes: The figure presents the estimated DiD coefficients β_j from Equation 11, estimated separately for juniors and seniors. The adoption measure used in this analysis is time-invariant; hence, the figure compares adopters and non-adopters. For results using a time-varying adoption measure, see Section 5.4. Standard errors are clustered at the firm level.

$$\begin{aligned}
\log(\text{Employment}_{ist}) = & \alpha + \sum_{j=2015Q2}^{2025Q1} \beta_j \mathbf{1}\{t = j\} \times \text{Adopt}_i \times \text{Junior}_s \\
& + \sum_{j=2015Q2}^{2025Q1} \pi_j \mathbf{1}\{t = j\} \times \text{Adopt}_i + \sum_{j=2015Q2}^{2025Q1} \rho_j \mathbf{1}\{t = j\} \times \text{Junior}_s \\
& + \kappa (\text{Adopt}_i \times \text{Junior}_s) + \gamma_{it} + \xi_{p(i)st} + \varepsilon_{ist}.
\end{aligned} \tag{12}$$

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 7 presents the estimated coefficients β_j from Equation 12. 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 5a and 5b 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.

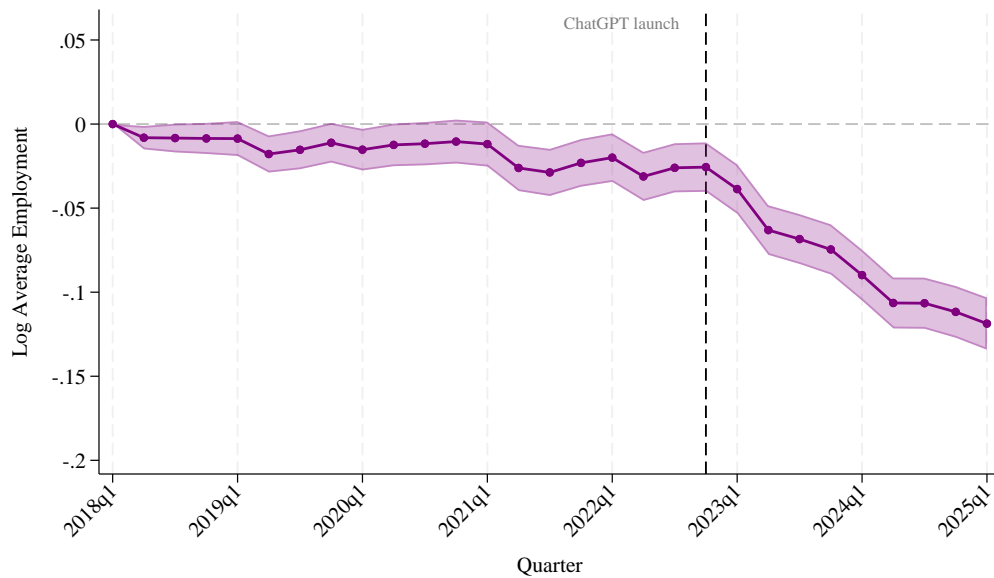


Figure 7: Triple-Difference Estimates: Juniors vs. Seniors

Notes: The figure presents the estimated triple-difference coefficients β_j from Equation 12, capturing how junior employment evolves relative to senior employment within the same firm and period, comparing adopters to non-adopters. Standard errors are clustered at the firm level.

Beginning in 2023Q1, the coefficients drop sharply, reaching roughly a 9 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.²⁹

The β_j coefficients can be interpreted as reflecting the differential impact of GenAI 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 5.3) and a staggered event-study design (Section 5.4). Section 5.5 further discusses possible forces that could challenge this interpretation and how we address them.

5.2.4 The Timing of the Junior Employment Decline:

The early and pronounced decline in the DiD and triple-difference coefficients—beginning in 2023Q1, shortly after the release of ChatGPT—may appear surprisingly abrupt, given that automation effects typically unfold more gradually. One plausible explanation is that the rapid diffusion of GenAI tools at the end of 2022 constituted a discrete shock to firms’ expectations about future automation. In response, firms may have adjusted in a forward-looking manner by scaling back hiring for 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.³⁰ Supplemental Appendix A.1 formalizes this mechanism in a simple dynamic model in which expectations of *future* automation—combined with labor-adjustment costs—lead firms to reduce hiring *today*.

One might still question whether firms were sufficiently aware of GenAI tools by early 2023 for such expectation shifts to occur. However, the sharp surge in adoption from 2023Q1 onward—as shown in Figure 3 and documented by studies such as [Bick et al. \(2024\)](#)—suggests that firms’ behavior changed almost immediately following the

²⁹In Supplemental Appendix A.14, we present triple-difference estimates excluding the industry-by-seniority-by-time fixed effects. The results are very similar to those in Figure 7.

³⁰Indeed, in Section 5.6 we show that the decline in junior employment is driven by reduced hiring rather than increased separations.

release of ChatGPT. Consistent with this, mentions of “AI” in firms’ earnings calls began rising steeply as early as 2022Q4, tripling by 2023Q2 (Supplemental Appendix A.16, Figure A.18). Finally, Supplemental Appendix A.17 compiles anecdotal evidence indicating that several prominent U.S. companies publicly adjusted their hiring practices and internal workflows in early 2023, explicitly attributing these changes to the adoption or anticipated impact of GenAI.

5.3 Which Junior Jobs Decline the Most Among Adopters

The previous section showed that GenAI adoption is associated with a decline in junior employment relative to senior employment. We now examine whether this seniority-biased effect is concentrated in occupations most exposed to GenAI. To do so, we extend the triple-difference specification of Equation 12 by replacing the Junior_s dimension with an indicator for *high-exposure* occupations, and by estimating the specification separately for juniors and seniors. This yields, for each seniority group, a comparison of how employment in high-exposure occupations evolves relative to low-exposure occupations, between adopters and non-adopters.³¹ Specifically, for each seniority group $s \in \{\text{junior}, \text{senior}\}$ we estimate:

$$\begin{aligned} \log(\text{Employment}_{ict}^s) = & \alpha + \sum_{j=2015Q2}^{2025Q1} \beta_j^s \mathbf{1}\{t = j\} \times \text{Adopt}_i \times \text{HighExp}_c \\ & + \sum_{j=2015Q2}^{2025Q1} \pi_j \mathbf{1}\{t = j\} \times \text{Adopt}_i + \sum_{j=2015Q2}^{2025Q1} \rho_j \mathbf{1}\{t = j\} \times \text{HighExp}_c \\ & + \kappa (\text{Adopt}_i \times \text{HighExp}_c) + \gamma_{it} + \zeta_{p(i)ct} + \varepsilon_{ict}, \end{aligned} \quad (13)$$

where $\log(\text{Employment}_{ict}^s)$ denotes log employment of seniority group s in exposure cell $c \in \{\text{high}, \text{low}\}$ at firm i in quarter t ; HighExp_c equals one for high-exposure occupations and zero for low-exposure occupations; γ_{it} are firm-by-time fixed effects; and $\zeta_{p(i)ct}$ are industry-by-exposure-by-time fixed effects. Equation 13 has the same structure as the seniority-based triple-difference in Equation 12, but here the third dimension is *occupational exposure* (high vs. low) rather than seniority. The coefficients β_j^s trace, separately

³¹Supplemental Appendix A.18 shows the most common high- and low-exposure occupations by industry, ranked by number of positions.

for juniors and seniors, how employment in high-exposure occupations evolves relative to low-exposure occupations within the same firm and quarter, comparing adopters to non-adopters.

Figure 8 plots the estimated β_j^s coefficients for juniors and seniors, with the coefficients normalized to zero in 2018Q1.³² Two patterns stand out. First, for both seniority groups, high-exposure occupations expand faster than low-exposure occupations at adopters relative to non-adopters throughout the pre-ChatGPT period (2018–2022), consistent with adopters tilting their workforce toward exposed occupations over time. The two profiles track each other closely through 2022, with seniors drifting somewhat above juniors. Second, the two series diverge sharply after the launch of ChatGPT in 2022Q4. The junior coefficient peaks at roughly 0.14 log points just before ChatGPT and then declines steadily to about 0.07 by 2025Q1—a relative contraction of roughly 7 log points in high-exposure versus low-exposure junior employment at adopters. In contrast, the senior coefficient continues to rise after 2022Q4, reaching about 0.22 by 2025Q1, with no visible break around the ChatGPT launch.

Taken together, these results provide two insights. First, the post-2022 contraction in junior employment at adopters is concentrated in high-exposure occupations rather than reflecting a broad-based decline across all junior roles. Second, the design helps rule out a purely compositional explanation—namely, that high-exposure occupations declined economy-wide after 2023 and adopters simply employed a larger share of such roles. By comparing high- and low-exposure occupations *within* the same firm-by-quarter, and by running the same specification for seniors as a placebo-like comparison, we find that high-exposure occupations contract relative to low-exposure occupations only for juniors and only at adopters after the ChatGPT launch. The senior profile continues along its pre-existing upward trend, consistent with the interpretation that GenAI adoption is seniority-biased: it shifts demand away from junior workers in the very occupations where the technology is most applicable, while leaving senior employment in those same occupations essentially unaffected.

Supplemental Appendix A.19 reports a complementary heterogeneity exercise that re-estimates the DiD specification within tiers of juniors' educational background, measured by the prestige of their alma mater. The relative decline in junior employment is U-shaped

³²Estimation begins in 2018Q1 for consistency with the other specifications in this section.




Figure 8: Triple-Difference by Occupational Exposure—Juniors vs. Seniors

Notes: This figure reports the estimated coefficients β_j^s from Equation 13, estimated separately for juniors (red) and seniors (blue). Each series captures the evolution of log employment in high-exposure occupations relative to low-exposure occupations at adopters relative to non-adopters, conditional on firm-by-time and industry-by-exposure-by-time fixed effects. Details on the exposure measure are provided in Section 4.1. Shaded bands are 95% confidence intervals. Standard errors are clustered at the firm level.

across tiers: largest for juniors from mid-tier institutions and smallest for juniors from the most and least selective institutions.

5.4 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:

$$\log(\text{JuniorEmployment}_{it}) = \alpha + \sum_{j=2}^J \beta_j (\text{Lag}_j)_{it} + \sum_{k=1}^K \gamma_k (\text{Lead}_k)_{it} + \mu_i + \lambda_t + \varepsilon_{it}, \quad (14)$$

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 9: Event Study

Notes: The graph presents the estimated coefficients β_j from Equation 14 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.

Figure 9 reports the results. The coefficients remain flat during the pre-adoption period, supporting the validity of the parallel-trends assumption. Approximately two 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. Appendix A.22 repeats this exercise for senior employment and shows no comparable post-adoption decline, consistent with the seniority-biased pattern documented in the DiD and triple-difference specifications.

5.5 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.

5.5.1 Confounding Labor Demand Shocks

A key concern is that other contemporaneous shocks may have disproportionately affected labor demand for junior workers in GenAI-adopting firms relative to non-adopters. We focus on two commonly discussed shocks. 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 9, 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 5.2.

Third, we directly test for this channel following [Ottonello and Winberry \(2020\)](#). We use the high-frequency identified monetary-policy surprises of [Jarocinski and Karadi \(2020\)](#), aggregated to the quarterly level as the in-quarter sum of monthly surprises, and estimate the differential-sensitivity regression

$$\log L_{i,t+1} - \log L_{i,t} = \alpha_i + \gamma_t + \beta(\text{Adopter}_i \times \varepsilon_t^m) + \eta_{i,t}, \quad (15)$$

on the pre-2022 firm-quarter panel ($t \in 2015\text{Q1}–2021\text{Q3}$), to see if our adopter firms were disproportionately responsive to previous monetary policy shocks. The outcome is the forward one-quarter log change in total or junior employment, α_i and γ_t are firm and year-quarter fixed effects. Under the monetary-policy confounder story, $\beta < 0$: adopters should contract *more* than non-adopters in the quarter following a tightening shock.

Table 4 reports the estimates. For total employment, the interaction is small and of the

Table 4: Sensitivity of adopter employment to identified monetary-policy shocks

	$\log L_{i,t+1} - \log L_{i,t}$	
	Total employment	Junior employment
Adopter _{<i>i</i>} × ε_t^m	0.046** (0.020)	0.018 (0.025)
Firm FE, Year-quarter FE	Yes	Yes
Firms	275,809	272,549
Observations	5,919,703	5,692,315
Quarters	23	23

Notes. LinkedIn/Revelio firm-quarter panel; shock quarters $t \in 2015Q1$ – $2021Q3$ (sample stops before 2022 so the forward one-quarter outcome window closes no later than 2021Q4). The shock ε_t^m is the in-quarter sum of monthly pure monetary-policy surprises from [Jarocinski and Karadi \(2020\)](#). Each column is one regression of the forward one-quarter log change in total (column 1) or juniors (column 2) employment on $\text{Adopter}_i \times \varepsilon_t^m$, with firm and year-quarter fixed effects. Outcomes are winsorised at 0.5/99.5%. Standard errors in parentheses are Cameron–Gelbach–Miller two-way cluster-robust at the firm and year-quarter level; t -statistics use $\min(G_{\text{firm}}, G_{\text{quarter}}) - 1$ degrees of freedom. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

opposite sign to what the confounder story would predict ($\hat{\beta} = 0.046$, s.e. 0.020): adopters expand slightly more, not less, than non-adopters in the quarter following a tightening. For junior employment, the interaction is small and statistically insignificant ($\hat{\beta} = 0.018$, s.e. 0.025). Either way, there is no evidence that adopters contract differentially after identified monetary-policy shocks. Appendix [A.20](#) reports lag-augmented local-projection impulse responses ([Jordà, 2005](#); [Montiel Olea and Plagborg-Møller, 2021](#)) across horizons $h = 0, \dots, 8$ quarters, which deliver the same null at every horizon. This finding is also consistent with evidence in the literature (see [Chodorow-Reich, 2014](#); [Gertler and Gilchrist, 1994](#)) showing that larger firms, like adopters in this case, are less likely to be sensitive to monetary-policy shocks.

An additional 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 [4.3](#)). However, these dynamics are also likely to be 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 significant relative increase in junior employment among adopters (relative to non-adopters) prior to GenAI adoption, which we do not (see Figures [5](#) and [9](#)).

5.5.2 Labor 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 11 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.

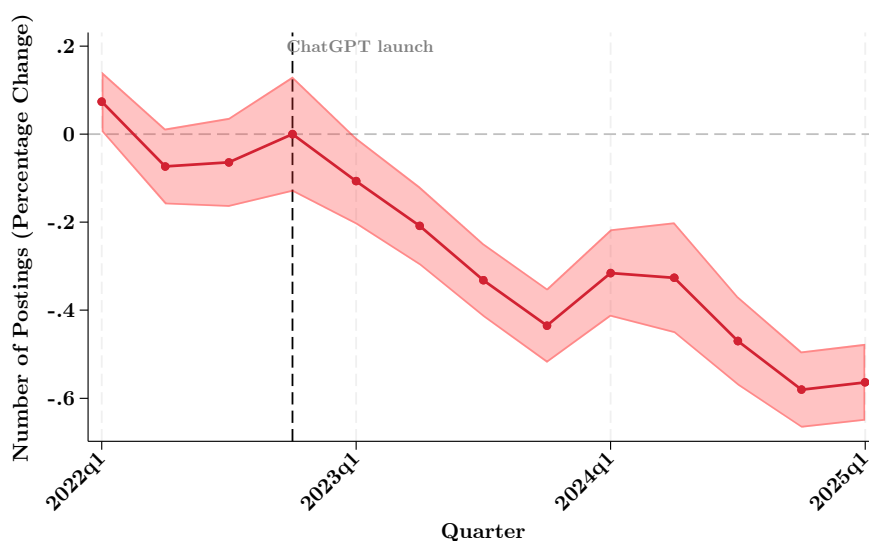


Figure 10: DiD—Job Postings

Notes: This figure reports the difference-in-differences estimates (β_j from Equation 11) 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.

Figure 10 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.

5.6 Decomposing Decline in Junior Employment Into Flows

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:

$$y_{it} = \alpha + \beta (\text{Adopt}_i \times \text{Post}_t) + \delta_t + \gamma_i + \xi_{pt} + \varepsilon_{it}, \quad (16)$$

where y_{it} denotes the number of junior hires, number of junior separations, or the junior promotion rate³³ in firm i at quarter t . γ_i denotes firm fixed effects, δ_t are time fixed effects and ξ_{pt} are industry-by-time fixed effects.

Table 5 reports the estimated β from Equation 16. The results indicate that the sharp contraction in junior employment among adopters is driven primarily by a slowdown in hiring. Specifically, the coefficient on *Hiring* implies that, relative to non-adopters, GenAI-adopting firms hired on average 4.0 fewer junior workers per quarter after 2023Q1—about an 80 percent reduction relative to the pre-period mean of 5.06 hires per firm-quarter. Junior separations also declined among adopters, by roughly 1.1 workers per quarter, partially offsetting the hiring contraction; but this reduction is about one-fourth the magnitude of the hiring decline, so net junior employment still contracts sharply.³⁴ Junior promotion rates rose by a small but statistically significant 0.033 percentage points (relative to a pre-period mean of 0.82 percent). Because promotions move workers out of the junior pool, this rise further reduces the junior stock, though its magnitude is very small relative to the hiring channel. Appendix A.23 explores sectoral heterogeneity in the decline in junior hiring, showing that the contraction is broad-based across industries and not driven disproportionately by any single sector.

³³The junior promotion rate is defined as $\frac{\text{Promotions}_t}{\text{Junior Employment}_{t-1}}$, expressed in percentage points.

³⁴See Supplemental Appendix A.21 for the time-series DiD results for junior hires and exits.

Table 5: Hiring, Separations, and Promotions

	Hiring	Separation	Promotion (rate)
Treat × Post	−4.006*** (0.221)	−1.089*** (0.163)	0.033** (0.014)
Pre-period mean	5.059	4.154	0.822
Observations	7,941,179	7,941,179	7,690,581
Clusters (firms)		280,866	

Notes: This table reports the estimated β from Equation 16. Estimation begins in 2018 to exclude early-sample trends. Standard errors clustered by firm in parentheses.

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

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 roughly 281,000 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 trends remain unchanged. Difference-in-differences, triple-difference, and staggered event-study estimates consistently point to this pattern. Comparing high- versus low-exposure occupations, we find that within adopting firms junior employment in high-exposure occupations contracts sharply relative to low-exposure occupations after 2022, while the corresponding senior series continues to rise—consistent with GenAI substituting for exposed junior tasks while complementing senior work within the same occupations. Finally, decomposing employment flows reveals that this effect stems primarily from a slowdown in hiring rather than increased separations (which in fact decline), with promotions playing only a minor role.

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. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.
- 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.
- Althoff, L. and Reichardt, H. (2026). Task-specific technical change and comparative advantage.
- Appel, R., Massenkoff, M., McCrory, P., McCain, M., Heller, R., Neylon, T., and Tamkin, A. (2026). The anthropic economic index report: Economic primitives.
- 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. and Thompson, N. (2025). Expertise. NBER Working Paper 33941, National Bureau of Economic Research.
- Autor, D. H. and Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market. *American Economic Review*, 103(5):1553–1597.
- Azar, J., Gaggl, P., and Taska, B. (2025). The eroding foundation: How AI is reshaping entry-level

- labor markets. Technical report, SSRN. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5842084.
- 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.
- Bessen, J. E. (2018). AI and jobs: The role of demand. *NBER Working Paper No. 24235*. National Bureau of Economic Research.
- 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.
- Crane, L. D. and Soto, P. E. (2026). AI and Coder Employment: Compiling the Evidence.
- Cruces, J. (2026). Does AI substitute for or augment human expertise? evidence from legal analysis. Technical report, Working Paper.
- 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>.
- Davis, S. J., Haltiwanger, J. C., and Schuh, S. (1998). Job creation and destruction. *MIT Press Books*, 1.
- Dell'Acqua, F., Ayoubi, C., Lifshitz, H., Sadun, R., Mollick, E., Mollick, L., Han, Y., Goldman, J., Nair, H., Taub, S., et al. (2025). The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise. Technical report, National Bureau of Economic Research.
- Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F., and Lakhani, K. R. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. *Harvard Business School Technology & Operations Management Unit Working Paper*, (24-013).
- Deming, D. J. (2023). Why Do Wages Grow Faster for Educated Workers? Technical report, National Bureau of Economic Research.
- Dominski, J. and Lee, Y. S. (2025). Advancing AI Capabilities and Evolving Labor Outcomes. *arXiv Preprint arXiv:2507.08244*.

- Eckhardt, S. and Goldschlag, N. (2025). AI and Jobs: The Final Word (Until the Next One). Economic Innovation Group. Accessed 2025-08-30.
- Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs Are GPTs: Labor Market Impact Potential of LLMs. *Science*, 384(6702):1306–1308.
- FactSet Research Systems Inc. (2025). Events and transcripts (callstreet). Earnings-call transcripts and events data (CallStreet).
- 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.
- Friebel, G., Huang, Y., Li, J., Shukla, S., and Zhang, A. (2026). Pyramids, Diamonds, and Oscillations: AI and the Structure of Internal Labor Markets.
- 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.
- Garicano, L. and Rossi-Hansberg, E. (2006). Organization and inequality in a knowledge economy. *Quarterly Journal of Economics*, 121(4):1383–1435.
- 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.
- Hershbein, B. and Kahn, L. B. (2018). Do Recessions Accelerate Routine-Biased Technological Change? Evidence From Vacancy Postings. *American Economic Review*, 108(7):1737–1772.
- Hosseini, S. M. and Lichtinger, G. (2026). Generative AI and occupational entry barriers: The labor-supply channel of technological change. Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6059674.
- Houston, D. (2023). A message from our CEO. *Dropbox Blog*. Accessed 2025-10-29.
- Humlum, A. and Vestergaard, E. (2025). Large Language Models, Small Labor Market Effects. Technical report, National Bureau of Economic Research.
- Ide, E. (2025). Automation, AI, and the intergenerational transmission of knowledge. *arXiv preprint arXiv:2507.16078*.
- Ide, E. and Talamas, E. (2025). Artificial Intelligence in the Knowledge Economy. *Journal of Political Economy*. Forthcoming.
- Ilut, C., Kehrig, M., and Schneider, M. (2018). Slow to hire, quick to fire: Employment dynamics with asymmetric responses to news. *Journal of Political Economy*, 126(5):2011–2071.

- 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.
- Jarocinski, M. and Karadi, P. (2020). Deconstructing Monetary Policy Surprises—The Role of Information Shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.
- Jordà, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.
- 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.
- Klein Teeselink, B., Meloso, D., Checkoway, H., and Tyson, L. (2025). Generative AI and labor market outcomes. Technical report, SSRN. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5516798.
- Lodefalk, M., Löthman, L., Koch, M., and Engberg, E. (2026). Same storm, different boats: Generative AI and the age gradient in hiring. Technical report, Working Paper.
- Medical Design & Development Staff (2023). Nuance, Microsoft Develop AI-Automated Clinical Documentation Application. Accessed 2025-10-30.
- Mia Sato (2023). CNET is Doing Big Layoffs Just Weeks After AI-generated Stories Came to Light. Accessed 2025-10-30.
- 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.
- Montiel Olea, J. L. and Plagborg-Møller, M. (2021). Local Projection Inference Is Simpler and More Robust Than You Think. *Econometrica*, 89(4):1789–1823.
- Morgan Stanley (2023). Morgan Stanley Wealth Management Announces Key Milestone in Innovation Journey with OpenAI. *Press Release*. Accessed 2025-10-29.
- Noy, S. and Zhang, W. (2023). Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Science*, 381(6654):187–192.
- Ottonello, P. and Winberry, T. (2020). Financial Heterogeneity and the Investment Channel of Monetary Policy. *Econometrica*, 88(6):2473–2502.
- Reuters (2023a). IBM to pause hiring as it plans to replace 7,800 jobs with AI. *Reuters*. Accessed 2025-10-29.
- Reuters (2023b). Meta creates Top-Level Team to Focus on AI. *Reuters*. Accessed 2025-10-29.
- Reuters (2023c). Meta to Cut 10,000 Jobs, Eliminate 5,000 Open Roles in Efficiency Push. *Reuters*. Accessed 2025-10-29.
- Revelio Labs (2025). Revelio Labs Workforce Data. Accessed via Wharton Research Data Services (WRDS).
- Sherwood News (2024). Companies are talking about ai considerably less than before. Article discussing trends in AI mentions; cites FactSet CallStreet transcripts.
- 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.
- The Wendy’s Company (2023). Wendy’s Taps Google Cloud to Revolutionize the Drive-Thru Experience with Artificial Intelligence. *PR Newswire*. Accessed 2025-10-29.
- TMCnet (2023). Presto and CKE Restaurants Announce Drive-Thru Voice AI Automation Partnership. Press release via TMCnet, accessed 2025-10-30.
- 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 Supplemental Appendix

A.1 Theoretical Framework: 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 \(2019\)](#).

A.1.1 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

$$\max_{n_0, n_1} \Pi = A_0 f(n_0) - (w_0 + c_h)n_0 + \beta \mathbb{E}_0[A_1 f(n_1) - w_1 n_1 - \phi(n_0 - n_1)_+]. \quad (17)$$

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 tomor-

row. The FOC is

$$A_0 f'(n_0) = \begin{cases} w_0 + c_h, & \text{if no firing is expected,} \\ w_0 + c_h + \beta\phi, & \text{if firing is anticipated,} \end{cases} \quad (18)$$

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

$$n_0^C = \left(\frac{\alpha A_0}{w_0 + c_h + \beta\phi} \right)^{\frac{1}{1-\alpha}}. \quad (19)$$

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

$$n_1^- = \left(\frac{\alpha A_1}{w_1 - \phi} \right)^{\frac{1}{1-\alpha}} \quad (\text{assuming } w_1 > \phi). \quad (20)$$

The contraction regime is relevant whenever

$$\frac{A_1}{w_1 - \phi} \leq \frac{A_0}{w_0 + c_h + \beta\phi}. \quad (21)$$

Differentiating (19) yields

$$\frac{\partial n_0^C}{\partial \phi} = -\frac{\beta}{1-\alpha} \cdot \frac{n_0^C}{w_0 + c_h + \beta\phi} < 0, \quad (22)$$

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

A.1.2 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 $\sigma > 1$,

$$Y_t = A_t \left[\Gamma_t^{1/\sigma} n_t^{(\sigma-1)/\sigma} + (1 - \Gamma_t)^{1/\sigma} (A^K K)^{(\sigma-1)/\sigma} \right]^{\sigma/(\sigma-1)}, \quad (23)$$

where capital K is fixed and the constants are absorbed into Ω_t below to simplify notation. The MPL implied by (23) is

$$\frac{\partial Y_t}{\partial n_t} = A_t \left[\Gamma_t^{1/\sigma} n_t^{(\sigma-1)/\sigma} + \Omega_t \right]^{1/(\sigma-1)} \Gamma_t^{1/\sigma} n_t^{-1/\sigma}. \quad (24)$$

The firm now maximizes (17) with $A_t f(n_t)$ replaced by Y_t . The optimality conditions take exactly the same form as in the baseline:

$$\frac{\partial Y_1}{\partial n_1} = \begin{cases} w_1, & n_1 \geq n_0, \\ w_1 - \phi, & n_1 < n_0, \end{cases} \quad \frac{\partial Y_0}{\partial n_0} = \begin{cases} w_0 + c_h, & \text{no firing expected,} \\ w_0 + c_h + \beta\phi, & \text{firing anticipated.} \end{cases} \quad (25)$$

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 (25) 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 $\partial n_0 / \partial (\mathbb{E}_0 I_1) < 0$.

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.

A.2 Conceptual Framework: Formal Details

This appendix provides the formal derivations and proofs for the conceptual framework in Section 3.

A.2.1 Threshold Shift

After GenAI adoption with augmentation (δ_J, δ_S) , the post-AI allocation threshold \hat{x}' is determined by the condition $\tilde{\gamma}_J(\hat{x}')/w_J = \tilde{\gamma}_S(\hat{x}')/w_S$, or equivalently

$$\frac{\gamma_S(\hat{x}')}{\gamma_J(\hat{x}')} = \frac{1 + \delta_J}{1 + \delta_S} \cdot \frac{w_S}{w_J} = \frac{1 + \delta_J}{1 + \delta_S} \cdot \frac{\gamma_S(\hat{x})}{\gamma_J(\hat{x})}. \quad (26)$$

When $\delta_J = \delta_S$, the right-hand side equals $\gamma_S(\hat{x})/\gamma_J(\hat{x})$, so $\hat{x}' = \hat{x}$ exactly. For small differential augmentation, a first-order Taylor expansion of $\ln(\gamma_S(x)/\gamma_J(x))$ around \hat{x} (holding wages fixed) yields

$$\hat{x}' - \hat{x} \approx \frac{\ln\left(\frac{1+\delta_J}{1+\delta_S}\right)}{\left.\frac{d}{dx} \ln\left(\frac{\gamma_S(x)}{\gamma_J(x)}\right)\right|_{x=\hat{x}}} \geq 0. \quad (27)$$

The numerator is increasing in $\delta_J - \delta_S$; the denominator—the steepness of comparative advantage at the boundary—is positive and finite by Assumption 1. By the implicit function theorem applied to (26), \hat{x}' is a continuous function of (δ_J, δ_S) .

A.2.2 Senior Employment Decomposition

By symmetric reasoning, the change in senior labor demand under $\alpha < \hat{x}$ is

$$\Delta L_S^* = \underbrace{0}_{\text{displacement}} + \underbrace{-\int_{\hat{x}}^{\hat{x}'} l_S^{*\text{pre}}(x) dx}_{\text{reallocation } (\leq 0)} + \underbrace{\int_{\hat{x}'}^1 [l_S^{*\text{post}}(x) - l_S^{*\text{pre}}(x)] dx}_{\text{productivity } (\leq 0)}. \quad (28)$$

Seniors face no displacement (no senior tasks are automated). The reallocation effect is weakly negative—seniors lose tasks near the boundary to newly-productive juniors. The productivity effect operates in the same direction as for juniors (see Appendix A.2.4), but

the within-task augmentation component is smaller when $\delta_S < \delta_J$; the cost-index channel, which is common to both worker types, is unchanged. The absence of displacement—the dominant channel for the junior decline—yields a muted net effect on senior labor demand.

A.2.3 Proof of Proposition 1

Proof. The change in junior labor demand is

$$\Delta L_J^* = L_J^{*\text{post}} - L_J^{*\text{pre}} = \int_{\alpha}^{\hat{x}'} l_J^{*\text{post}}(x) dx - \int_0^{\hat{x}} l_J^{*\text{pre}}(x) dx. \quad (29)$$

Splitting each integral at the shared endpoints α and \hat{x} :

$$\begin{aligned} \Delta L_J^* &= \left[\int_{\alpha}^{\hat{x}} l_J^{*\text{post}}(x) dx - \int_{\alpha}^{\hat{x}} l_J^{*\text{pre}}(x) dx \right] + \int_{\hat{x}}^{\hat{x}'} l_J^{*\text{post}}(x) dx - \int_0^{\alpha} l_J^{*\text{pre}}(x) dx \\ &= \underbrace{- \int_0^{\alpha} l_J^{*\text{pre}}(x) dx}_{\text{displacement}} + \underbrace{\int_{\alpha}^{\hat{x}} [l_J^{*\text{post}}(x) - l_J^{*\text{pre}}(x)] dx}_{\text{productivity}} + \underbrace{\int_{\hat{x}}^{\hat{x}'} l_J^{*\text{post}}(x) dx}_{\text{reallocation}}. \quad \square \end{aligned}$$

A.2.4 Sign of the Productivity Effect

We maintain the partial-equilibrium assumption that Y and wages (w_J, w_S) are fixed.

Under cost minimization with the CES aggregator (1), conditional demand for task x is $y^*(x) = Y \cdot (p(x)/P)^{-\sigma}$, where $p(x)$ is the unit cost of task x and

$$P = \left(\int_0^1 p(z)^{1-\sigma} dz \right)^{\frac{1}{1-\sigma}} \quad (30)$$

is the CES cost index (with the Cobb-Douglas limit $P = \exp\left(\int_0^1 \ln p(z) dz\right)$ when $\sigma = 1$). For a junior task $x \in (\alpha, \hat{x})$, the unit cost is $p(x) = w_J/\gamma_J(x)$ and output is $y(x) = \gamma_J(x) l_J(x)$, so the cost-minimizing labor allocation is

$$l_J^*(x) = Y \cdot P^{\sigma} \cdot w_J^{-\sigma} \cdot \gamma_J(x)^{\sigma-1}. \quad (31)$$

After GenAI adoption, the cost index changes from P^{pre} to P^{post} . On surviving junior

tasks $x \in (\alpha, \hat{x})$, augmented productivity is $\tilde{\gamma}_J(x) = (1 + \delta_J)\gamma_J(x)$, so the ratio of post- to pre-AI labor demand is

$$\frac{l_J^{*\text{post}}(x)}{l_J^{*\text{pre}}(x)} = \left(\frac{P^{\text{post}}}{P^{\text{pre}}} \right)^\sigma (1 + \delta_J)^{\sigma-1}. \quad (32)$$

Since GenAI weakly reduces unit costs on every task—automated tasks satisfy $c_A(x) < w_J/\gamma_J(x)$ by assumption, augmented tasks have unit cost $w_k/((1 + \delta_k)\gamma_k(x)) \leq w_k/\gamma_k(x)$, and the threshold reallocation only occurs when it is cost-reducing—the post-AI cost index satisfies $P^{\text{post}} \leq P^{\text{pre}}$, with strict inequality whenever $\alpha > 0$ or $\delta_k > 0$ for some k .³⁵ Write $\rho \equiv (P^{\text{post}}/P^{\text{pre}})^\sigma \in (0, 1]$.

The ratio (32) is uniform across common junior tasks, so the sign of the productivity effect is determined by whether $\rho \cdot (1 + \delta_J)^{\sigma-1}$ exceeds, equals, or falls below one:

- $\sigma < 1$: $(1 + \delta_J)^{\sigma-1} \leq 1$ (with equality iff $\delta_J = 0$) and $\rho \leq 1$. The product is ≤ 1 , with strict inequality whenever $\delta_J > 0$ or $\alpha > 0$. The productivity effect is non-positive.
- $\sigma > 1$: $(1 + \delta_J)^{\sigma-1} \geq 1$ but $\rho \leq 1$. Within-task augmentation raises labor demand per task, while the reduction in the cost index lowers it. The sign is ambiguous in general.
- $\sigma = 1$: $(1 + \delta_J)^0 = 1$, so the ratio reduces to $\rho = P^{\text{post}}/P^{\text{pre}} \leq 1$. The within-task augmentation effect vanishes and only the cost-index effect remains. The productivity effect is non-positive, and strictly negative whenever $P^{\text{post}} < P^{\text{pre}}$.³⁶

Sufficient condition for net decline when $\sigma \leq 1$.

When $\sigma \leq 1$, the productivity effect is non-positive, so

$$\Delta L_J^* \leq - \int_0^\alpha l_J^{*\text{pre}}(x) dx + \int_{\hat{x}}^{\hat{x}'} l_J^{*\text{post}}(x) dx. \quad (33)$$

³⁵The CES cost index P is increasing in each task's unit cost. Since cost minimization assigns each task to the cheapest available input, and GenAI either automates tasks at lower cost or augments human productivity, every task's effective unit cost weakly decreases after GenAI adoption.

³⁶With $\sigma = 1$, the production function is $Y = \exp\left(\int_0^1 \ln y(x) dx\right)$ and cost minimization equalizes expenditure across tasks: $p(x)y(x) = P \cdot Y$ for all x . Since $p(x) = w_J/\gamma_J(x)$ and $y(x) = \gamma_J(x)l_J(x)$, we obtain $w_J l_J^*(x) = P \cdot Y$, i.e., $l_J^*(x) = P Y/w_J$. After GenAI adoption, $l_J^{*\text{post}}(x) = P^{\text{post}} Y/w_J < P^{\text{pre}} Y/w_J = l_J^{*\text{pre}}(x)$ whenever $P^{\text{post}} < P^{\text{pre}}$.

The reallocation effect is bounded above by $(\hat{x}' - \hat{x}) \cdot \bar{l}$, where $\bar{l} = \sup_{x \in [\alpha, \hat{x}']} l_J^{*\text{post}}(x)$. The displacement effect $\int_0^\alpha l_J^{*\text{pre}}(x) dx > 0$ for any $\alpha > 0$. Therefore $\Delta L_J^* < 0$ whenever

$$(\hat{x}' - \hat{x}) \cdot \bar{l} < \int_0^\alpha l_J^{*\text{pre}}(x) dx. \quad (34)$$

Note that (34) is only a sufficient condition for a net decline in junior labor demand, not a characterization of all parameter configurations under which the decline occurs; the net effect can be negative even when this bound is not satisfied, because the bound replaces the productivity effect with zero.

When $\delta_J = \delta_S$, the augmented productivity ratio $\tilde{\gamma}_S(x)/\tilde{\gamma}_J(x) = \gamma_S(x)/\gamma_J(x)$ is unchanged, so $\hat{x}' = \hat{x}$ exactly. The reallocation effect is zero and $\Delta L_J^* < 0$ follows immediately. More generally, by the implicit function theorem (applicable since $\frac{d}{dx} \ln(\gamma_S/\gamma_J)|_{\hat{x}} > 0$ by Assumption 1), \hat{x}' is continuous in (δ_J, δ_S) . Since $\hat{x}' - \hat{x}$ vanishes at $\delta_J = \delta_S$ while the displacement effect is bounded away from zero, there exists a neighborhood around $\delta_J = \delta_S$ in which (34) holds.

A.3 AI Usage Patterns Across the Seniority Spectrum

This appendix provides external descriptive evidence on how GenAI is used for tasks at different levels of the occupational seniority hierarchy, using data on GenAI usage patterns that are independent of the firm-level employment data in the main analysis. The analysis draws on the Anthropic Economic Index (AEI), a public dataset that classifies real-world interactions with a large language model by the occupational tasks they serve and the nature of human–AI collaboration. We merge these GenAI usage measures with occupation-level seniority indicators from Revelio Labs/LinkedIn to document systematic differences in the *mode* of GenAI engagement across junior-type and senior-type tasks. The goal is not to establish causal effects but to provide theory-motivated descriptive evidence that complements the conceptual framework in Section 3: the framework predicts that junior-type tasks are disproportionately exposed to automation, while senior-type tasks—conditional on remaining human-performed—are more likely to involve augmentation-type GenAI usage. Note that while the main analysis studies within-firm junior versus senior employment, Appendix A.3 uses occupation-level seniority as a proxy for the type of tasks typically performed by more junior or more senior workers—it captures cross-occupation variation in task complexity, not the within-firm hierarchy itself. The data allow us to examine whether observed GenAI usage patterns are consistent with the framework’s prediction.

A.3.1 Data Sources

Anthropic Economic Index (AEI). The AEI dataset (Appel et al., 2026) is constructed from a large sample of anonymized conversations with Claude, Anthropic’s large language model, during the week of November 13–20, 2025. Each conversation is algorithmically matched to the O*NET task it most closely corresponds to and classified into one of five *collaboration modes* that characterize how the user interacts with the AI:

- **Directive:** the user instructs the GenAI to produce an output with minimal back-and-forth. The GenAI executes a well-specified request—analogueous to delegating a task to a subordinate.
- **Learning:** the user seeks to acquire knowledge or understanding from the AI, using

it as an informational resource or tutor.

- **Task iteration:** the user and GenAI collaboratively refine an output through repeated exchange, with the user providing feedback and the GenAI revising its work.
- **Validation:** the user presents work or reasoning to the GenAI for checking, verification, or quality assurance.
- **Feedback loop:** a hybrid mode combining elements of directive and iterative interaction.

We use these collaboration modes as an interpretive proxy for the framework’s conceptual distinction between displacement and augmentation, recognizing that this mapping is a classification choice rather than a direct measurement of economic substitution. Specifically, we classify directive use and feedback loops as *displacement-type* interactions, in which GenAI usage tilts toward substitution for human task execution, and learning, task iteration, and validation as *augmentation-type* interactions, in which GenAI complements human judgment and expertise. We use data from both the Claude.ai consumer platform and the first-party API, though we focus primarily on Claude.ai usage in the reported results.

Revelio Labs/LinkedIn workforce composition. Our primary seniority measure is based on Revelio Labs data derived from public LinkedIn profiles. For each occupation, we compute the average seniority level of active workers as of January 2023 (pre-ChatGPT) using Revelio’s 1–7 seniority scale, and standardize to a z-score. The Revelio measure captures the *actual* seniority composition of the workforce—whether an occupation is predominantly staffed by junior or senior workers—which directly corresponds to the junior–senior distinction emphasized in the conceptual framework. We restrict the sample to occupations with at least 50 observed workers and partition occupations into terciles based on this measure (“low/junior-type,” “medium,” “high/senior-type”).

O*NET experience requirements. As a complementary validation, we also report the average years of related work experience required for each occupation from the O*NET Education, Training, and Experience database (version 30.0). This normative measure

captures how much experience an occupation *requires* rather than the actual seniority of its workforce, providing an independent check that the Revelio-based classification aligns with established occupational characteristics.

O*NET task-to-occupation mapping. The AEI matches conversations to O*NET task descriptions, and the O*NET task statements database provides the mapping from tasks to detailed occupations (8-digit O*NET-SOC codes). Each task may appear in multiple occupations, and each occupation contains multiple tasks. We merge the AEI usage data with this mapping to create a dataset at the task–occupation level, where each observation represents a specific O*NET task within a specific occupation, with associated GenAI usage statistics. After merging, the analysis sample contains 4,494 task–occupation pairs spanning 753 occupations across 22 two-digit SOC groups.

A.3.2 Descriptive Patterns

Table A.1 reports average GenAI usage patterns by Revelio seniority tercile. To summarize the balance between automation and augmentation in GenAI usage, we construct two measures at the task level: *displacement intensity*, defined as the share of classified GenAI interactions (displacement + augmentation) accounted for by displacement-type modes (directive and feedback loop), and *augmentation intensity*, defined analogously for augmentation-type modes (learning, validation, and task iteration). The two sum to 100% by construction for each task.

The aggregation from raw conversations to the tercile-level statistics in Table A.1 proceeds as follows. First, the AEI assigns each conversation a collaboration mode and maps it to an O*NET task; we compute mode shares at the task–occupation level (the share of conversations for a given task in a given occupation that fall into each mode). Second, we aggregate to the occupation level by taking weighted averages of task-level mode shares, with weights proportional to AEI task usage counts. Third, we classify occupations into terciles based on Revelio seniority and report unweighted averages across occupations within each tercile. The regression analyses in Table A.2 use the occupation-level measures as dependent variables. The sample shrinks from 753 occupations (full merged sample) to 559 in Table A.1 (requiring ≥ 50 Revelio workers for tercile assignment) and to 387 in Table A.2 (additionally requiring ≥ 3 AI-used tasks and non-missing seniority).

The central finding is that displacement intensity exceeds 50% only for junior-type occupations (52.0%), meaning that when GenAI is used for tasks characteristic of these occupations, the typical task tilts toward automation. For senior-type occupations, augmentation intensity dominates at 59.6% (displacement intensity: 40.4%). In other words, tasks in junior-intensive occupations are more likely to engage GenAI as a substitute, while tasks in senior-intensive occupations are more likely to engage GenAI as a complement.

This pattern is driven by the augmentation side. The individual collaboration mode shares show that task iteration rises monotonically with seniority (12.5% to 22.8%) and validation increases from 1.1% to 3.7%, while the directive share remains roughly flat across terciles. We also report the average O*NET experience requirements for each tercile as a complementary validation: occupations classified as junior-type by Revelio indeed require substantially less experience (1.4 years vs. 3.6 years for senior-type), confirming that the two seniority measures capture similar variation.

A.3.3 Regression Evidence

To assess whether the descriptive patterns hold within broad occupation groups, we regress occupation-level GenAI usage shares on the seniority measures, controlling for two-digit SOC fixed effects with heteroskedasticity-robust standard errors. This comparison exploits variation in seniority *within* occupation families—for instance, comparing secretaries to managers within “Office and Administrative Support,” or paralegals to lawyers within “Legal” occupations.

Table A.2 reports the results. The key finding is that displacement intensity *declines* significantly with seniority, confirming the descriptive pattern within narrow occupation groups. A one-standard-deviation increase in Revelio seniority is associated with a 3.68 percentage point decline in displacement intensity with 2-digit SOC fixed effects ($p = 0.002$). The individual mode shares tell a consistent story: the augmentation share increases significantly with seniority, driven primarily by task iteration (1.74 pp, $p = 0.029$) and validation (0.82 pp, $p = 0.005$), while the displacement share declines (−1.91 pp, $p = 0.018$). These results are consistent with the framework’s prediction that GenAI serves as a substitute for routine task execution in junior-intensive occupations and as a

Table A.1: AI Usage Patterns by Revelio Seniority Tercile

	Low (Junior-type)	Medium	High (Senior-type)
Number of occupations	151	202	206
Avg. Revelio seniority	1.36	2.03	3.11
Avg. experience (years, O*NET)	1.39	2.34	3.57
<i>Collaboration mode shares (%)</i>			
Directive	24.4	23.9	26.1
Feedback loop	3.4	2.2	3.0
Learning	13.0	16.7	15.5
Validation	1.1	2.6	3.7
Task iteration	12.5	17.9	22.8
Displacement intensity (%)	52.0	42.0	40.4
Augmentation intensity (%)	48.0	58.0	59.6

Notes: Each observation is an occupation. Occupations are divided into terciles based on the average seniority level of workers observed in Revelio Labs/LinkedIn data as of January 2023, restricted to occupations with ≥ 50 observed workers. Collaboration mode shares are weighted averages across an occupation’s tasks, with weights proportional to AEI task usage counts. Displacement and augmentation intensities are computed at the task level as displacement/(displacement + augmentation) and augmentation/(displacement + augmentation), respectively, then averaged across tasks within each occupation and across occupations within each tercile. Displacement-type interactions include directive and feedback loop modes; augmentation-type interactions include learning, validation, and task iteration modes. The five collaboration mode shares do not sum to 100% because a substantial fraction of task-matched conversations are not confidently assigned to any mode by the AEI classification algorithm; displacement and augmentation intensities are computed using only the classified subset and sum to 100% by construction. Values above 50% for displacement intensity indicate that GenAI usage for a given task tilts toward automation rather than augmentation. The AEI data cover Claude interactions during November 13–20, 2025.

complement to expert judgment in senior-intensive occupations.

A.3.4 Discussion

The AEI analysis provides descriptive evidence consistent with the core mechanism of the conceptual framework. The framework assumes that GenAI is more likely to automate the routine, codifiable tasks that constitute the bulk of junior work, while the complex, judgment-intensive tasks that characterize senior work are more likely to involve augmentation-type GenAI usage. The data show that this asymmetry manifests

Table A.2: Revelio Seniority and GenAI Usage Mode: Occupation-Level Regressions

Dependent variable	No fixed effects			2-digit SOC FE		
	β	SE	p	β	SE	p
Directive	-1.00	0.74	0.18	-1.15	0.80	0.15
Learning	-0.18	0.71	0.80	0.22	0.89	0.81
Validation	0.83	0.26	0.002	0.82	0.29	0.005
Task iteration	2.55	0.71	<0.001	1.74	0.79	0.029
Displacement	-1.25	0.77	0.11	-1.91	0.80	0.018
Augmentation	3.20	0.92	<0.001	2.78	1.05	0.008
Displacement intensity	-3.26	1.02	0.002	-3.68	1.15	0.002

Notes: $N = 387$ occupations (with ≥ 3 AI-used tasks, non-missing seniority, and ≥ 50 Revelio workers). The independent variable is the Revelio seniority z-score (standardized average seniority from LinkedIn profiles, January 2023). Each row is a separate regression. Dependent variables are occupation-level weighted averages of GenAI collaboration mode shares, with weights proportional to task usage counts. Displacement = directive + feedback loop. Augmentation = learning + validation + task iteration. Displacement intensity = displacement / (displacement + augmentation) $\times 100$, computed at the task level and then averaged to the occupation level. Heteroskedasticity-robust standard errors.

not merely in which tasks are exposed to AI, but in *how* AI is used: junior-type tasks are engaged through directive interactions that substitute for human execution, while senior-type tasks are engaged through iterative and learning-oriented interactions that complement human expertise. This pattern—a qualitative difference in the mode of human–AI collaboration across the seniority spectrum—provides descriptive support for the aggregate employment patterns documented in the main analysis.

Several caveats apply. First, the AEI data classify conversations by the O*NET task they address, not by the seniority of the individual user. We observe how AI is used *for* tasks that are characteristic of junior or senior occupations, but we cannot determine whether a junior or senior worker initiated the interaction. The relevant variation is therefore at the task level, not the worker level. Second, the data cover a single week in November 2025 on a single AI platform, and usage patterns may evolve as the technology and its user base mature. Third, the collaboration mode classification is algorithmic and imperfect; some interactions may be misclassified. Despite these limitations, the consistency of the patterns across occupation-level and task-level specifications and across a range of professional domains suggests that the qualitative pattern is stable across specifications.

A.4 GenAI Keywords in Worker Position Descriptions

This appendix reports the full regression results underlying the validation exercise in Section 4.3.3. The main text presents the time-series pattern visually (Figure 4); here we describe the keyword list and formalize the pattern in a difference-in-differences framework.

We apply a curated list of GenAI keywords to 18.3 million U.S. position descriptions spanning 2021 Q1 through 2025 Q1, drawn from the Revelio Labs individual-positions data. The keyword list builds on the list used to identify GenAI integrator postings in Section 4.3, augmented with additional terms that workers are more likely to use when describing their own adoption of GenAI tools on their résumés. Matching is case-insensitive; ambiguous short forms (e.g., LLM, GPT, NLP, Bard) are matched with word-boundary regex to avoid false positives. Positions with descriptions shorter than 20 characters are excluded as likely placeholders.

The full keyword list is as follows:

- **Original list (job-postings analysis).** *Models and products:* ChatGPT, Chat GPT, GPT, GPT-4, GPT-3, GPT-4o, GPT integration, Claude, Copilot, Gemini. *Concepts:* large language model (LLM), generative AI, Gen AI, transformer-based model, prompt engineering, prompt design, vector embeddings, vector database, retrieval-augmented generation (RAG). *Frameworks, APIs, and infrastructure:* LangChain, LlamaIndex, RetrievalQA, Pinecone, Weaviate, Milvus, OpenAI API, Anthropic Claude API, Azure OpenAI, Google Vertex AI Generative, HuggingFace Transformers.
- **Additional résumé-specific terms.** *Worker self-description of AI use:* AI tools, AI-powered, AI-driven, AI-assisted, AI-based, AI automation. *Adjacent ML/DL vocabulary:* machine learning, ML model, deep learning, neural network, natural language processing (NLP), computer vision. *Consumer-facing GenAI tools:* GitHub Copilot, Midjourney, DALL-E, Stable Diffusion, Bard, Perplexity. *Workflow terms:* fine-tuning, text generation, image generation, AI chatbot, conversational AI.

The additional terms are designed to capture how workers themselves describe GenAI use on their résumés, which differs in tone and vocabulary from how firms advertise integrator roles in vacancy postings: workers are more likely to name the specific consumer-

facing tools they use (e.g., Midjourney, GitHub Copilot) and to describe their work in adjacent ML/DL or workflow terms, while job postings emphasize technical infrastructure (vector databases, frameworks, APIs).

Table A.3 reports the difference-in-differences estimates. The dependent variable is an indicator for whether the position description contains at least one GenAI keyword. Column (1) reports the basic OLS estimate; column (2) adds firm and month fixed effects, yielding a DiD coefficient of 1.35 percentage points (s.e. = 0.05, $p < 0.001$); and column (3) additionally controls for description length.

Table A.3: Difference-in-Differences: GenAI Keywords in Position Descriptions

Dependent Variable:	has_ai		
Model:	OLS (1)	FE (2)	FE + Length (3)
Adopter \times Post	0.0176*** (0.0015)	0.0135*** (0.0005)	0.0140*** (0.0005)
<i>Fixed effects</i>			
Firm		Yes	Yes
Month		Yes	Yes
Observations	18,345,320	18,345,320	18,345,320
R ²	0.010	0.117	0.121

Notes: Standard errors clustered by firm in parentheses. *** $p < 0.01$.

A.5 Representativeness of the LinkedIn/Revelio Data

This appendix evaluates how the LinkedIn/Revelio sample compares to official benchmarks along three dimensions—industry, occupation, and an age proxy based on seniority—both in 2024 levels and in stability over the GenAI diffusion period. Two findings emerge. First, the data display the well-documented compositional tilt of LinkedIn toward managerial, professional, and high-information occupations and industries. Second, this tilt is essentially unchanged from 2021Q1 through 2025Q1, so the level differences are absorbed by the firm and sector fixed effects in our main specifications and the differential identification of treatment effects is unaffected.

We use three official benchmarks. The industry benchmark is constructed from BLS Quarterly Census of Employment and Wages (QCEW) national-area files, retaining all-size observations, averaging employment across the three months within each quarter, and aggregating to 3-digit NAICS. The occupation benchmark uses annual national OEWS employment shares aggregated to 2-digit SOC. After aggregating LinkedIn occupations to 2-digit SOC, the common support covers 100% of 2024 official OEWS employment. The age benchmark is constructed from quarterly QWI age-by-industry employment aggregated from 3-digit industries to broad sectors. To match LinkedIn’s junior/senior split, we divide the QWI 25–34 age bin evenly between 25–29 and 30–34 and compare LinkedIn juniors with the under-30 workforce and LinkedIn seniors with the 30+ workforce.

Industry representation. Figure A.1 compares LinkedIn and QCEW employment shares for the twenty broad sectors in 2024. Manufacturing, Educational Services, Professional Services, Information, Finance and Insurance, and Public Administration are over-represented in LinkedIn, while Retail, Accommodation and Food, Administrative and Waste, and Transportation are under-represented. These patterns are consistent with prior work using large-scale online employment data, which finds that coverage is concentrated in industries with a high share of white-collar and college-educated workers (e.g., [Hershbein and Kahn, 2018](#)). Figure A.2 shows the analogous comparison at the more granular 2-digit NAICS level: despite the level differences, the points line up closely along the 45-degree line, with a 2024 cross-sectional correlation of 0.70.




Figure A.1: Broad Sector Shares in LinkedIn/Revelio and Official Employment Data

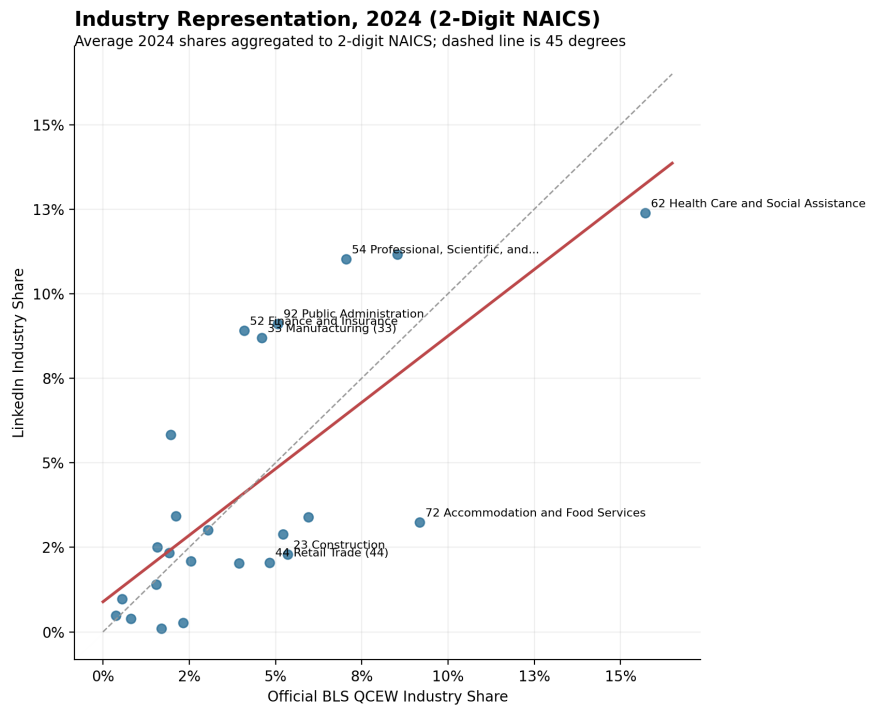


Figure A.2: 2-Digit Industry Shares in LinkedIn/Revelio and Official Employment Data

Occupation representation. Figure A.3 performs the same comparison at the 2-digit SOC level using the OEWS benchmark. The compositional tilt is sharper at the occupation level: Management, Business and Finance, Computer and Mathematical, and Arts/Media occupations are over-represented, while Office and Administrative, Transportation, Food Service, Production, and lower-wage service occupations are under-represented. The 2024 cross-sectional correlation is 0.37, reflecting the well documented over-sampling of higher-skill occupations within sectors. Crucially, because these compositional gaps are stable over time (see below), they are absorbed by the firm and sector fixed effects in our main specifications and do not contaminate the differential estimates of interest.

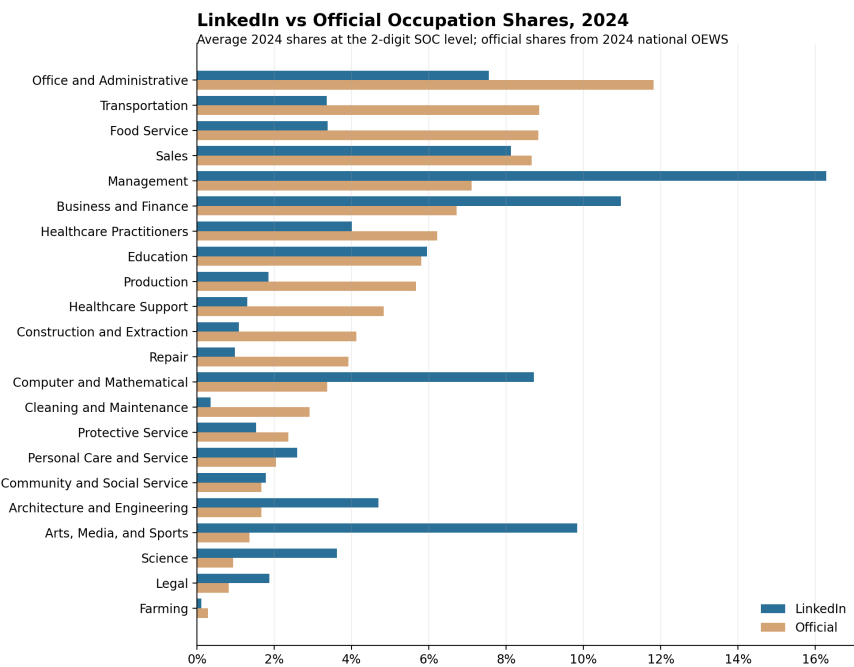


Figure A.3: 2-Digit Occupation Shares in Revelio and Official Employment Data

Stability over time. Because our identification rests on the *differential* evolution of employment around the diffusion of GenAI, the relevant question is not whether LinkedIn matches the official distribution in levels, but whether its compositional structure is stable across the treatment window. Figure A.4 plots the quarterly cross-sectional correlation between LinkedIn and the official benchmarks from 2021Q1 through 2025Q1, separately for industries (top panel) and occupations (bottom panel). The industry correlation stays in a tight 0.77–0.82 band; the occupation correlation stays in a 0.37–0.41 band. Neither

series shows a discernible break around the launch of ChatGPT (vertical dashed line at 2023Q1), indicating that the compositional structure of the LinkedIn sample is essentially unchanged across the pre- and post-GenAI periods.

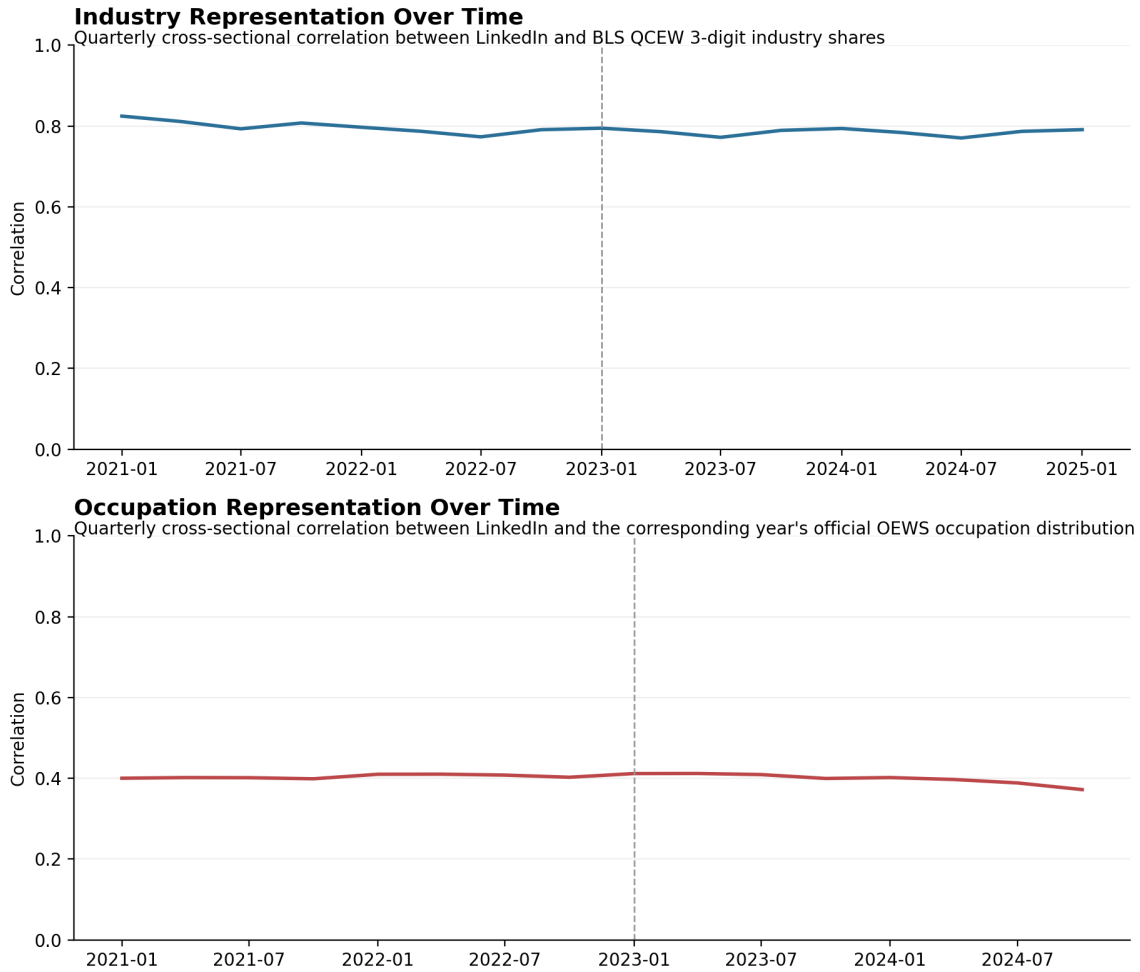


Figure A.4: Stability of Industry and Occupation Representation Over Time

Figure A.5 repeats the exercise for the age dimension. The top panel plots the quarterly cross-sectional correlation between LinkedIn junior sector shares and the corresponding QWI under-30 sector shares; the bottom panel does the same for senior versus 30+ shares. The junior correlation is stable around 0.44 throughout the sample, and the senior correlation is stable around 0.60, with neither series showing a break around the GenAI launch. The relative representation of junior and senior workers across sectors in the LinkedIn data therefore evolves in line with the official QWI age structure, supporting the use of LinkedIn-based seniority composition to identify changes in junior and senior employ-

ment around GenAI adoption.




Figure A.5: Stability of Age-Proxy Representation Over Time

A.6 Validation of the Seniority Classification

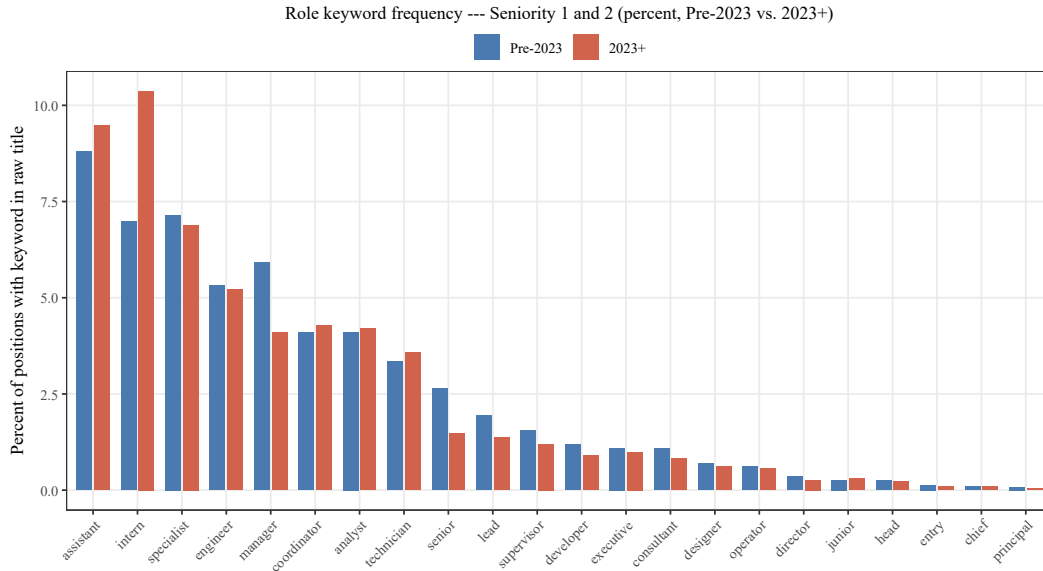
We conduct two complementary exercises to evaluate the seniority classification produced by Revelio and to verify its stability over the GenAI diffusion period. First, we analyze job title keywords for junior and senior roles. Second, we examine the average experience levels within each seniority group. In both exercises, we study patterns over time and across adoption status.

A.6.1 Job Title Keywords by Seniority

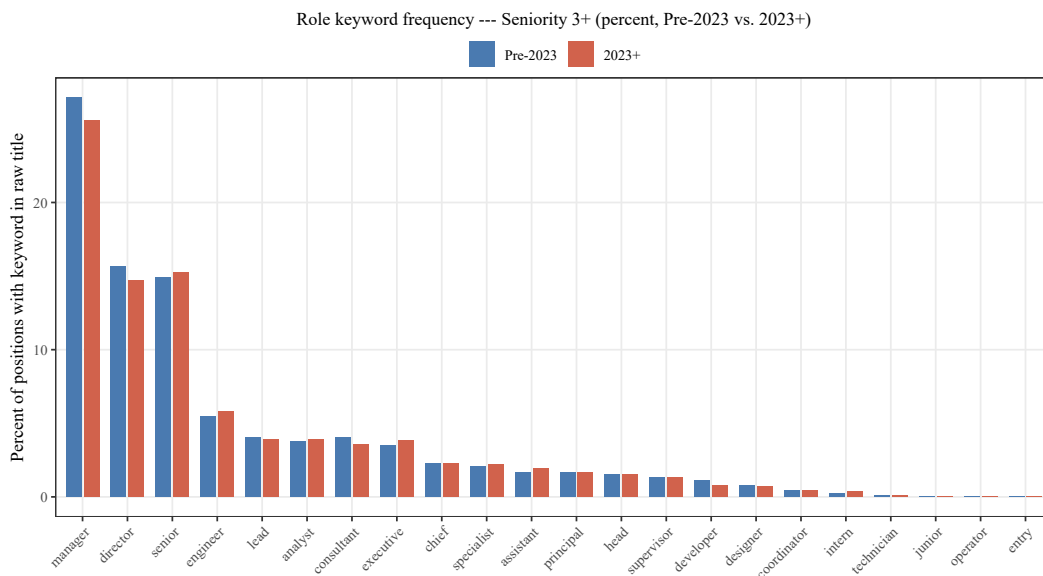
What junior and senior positions look like. We first examine the distribution of keywords in raw job titles, separately for junior and senior workers. Figure A.6 presents word clouds summarizing the most frequently occurring keywords in each seniority group across the full sample. The classification is meaningful: junior positions are dominated by terms such as *assistant*, *intern*, *specialist*, *associate*, *technician*, and *analyst*, while senior positions are concentrated around *manager*, *director*, *senior*, *consultant*, *president* and *lead*. These patterns confirm that Revelio’s seniority variable aligns well with conventional titling and motivate the keyword lists used in the exercises below.

Stability of keyword distributions across 2023. We next compare the distribution of keyword frequencies in positions that started before 2023 to those that started in 2023 or later, pooled across all firms. Figure A.7 shows that the relative prominence of keywords is stable over time within each seniority group. Junior positions continue to be dominated by the same keywords before and after 2023, and senior positions by theirs. The main exception is *intern*, whose share increases after 2023 relative to the earlier period. This stability argues against a systematic shift in titling conventions or classification methodology around the diffusion of GenAI.

Differential evolution at adopting vs. non-adopting firms. The previous exercises pool all firms together. To directly address the concern that *adopting* firms may be reclassifying positions differentially, we track the evolution of keyword shares over time, separately at adopting and non-adopting firms. For each firm-quarter cell, we compute two shares: the share of junior-classified positions whose raw title contains any junior-leaning



(a) Junior positions (Seniority 1–2)



(b) Senior positions (Seniority 3+)

Figure A.7: Keyword Frequency by Seniority Group, Pre-2023 vs. 2023+

where $\text{Share}_{f,t}^{\text{match}}$ is the share of positions at firm f in quarter t whose raw title contains a matching keyword (junior-leaning for juniors, senior-leaning for seniors). The regression is weighted by the number of positions per firm-quarter cell, and standard errors are clustered at the firm level. Figure A.9 reports the estimated coefficients. For juniors, the




Figure A.8: Share of Positions with Matching Keywords, by Seniority and Adoption

Notes: This figure plots, for each quarter, the share of positions whose raw title contains a “matching” keyword: junior-leaning keywords among junior-classified positions (left panel) and senior-leaning keywords among senior-classified positions (right panel), at GenAI-adopting (solid) and non-adopting (dashed) firms. Shares are position-weighted within each firm-quarter cell, then aggregated across firms by adoption status. The vertical dashed line marks the launch of ChatGPT in November 2022. Sample: quarters 2020Q1–2025Q1.

coefficients hover around zero throughout the sample, with a visible quarterly saw-tooth that reflects the seasonal pattern of summer hiring and no discernible break at the ChatGPT launch. For seniors, the coefficients drift mildly upward in the post-period, reaching approximately +0.3 percentage points by 2025Q1—a quantitatively small movement relative to a baseline share of roughly 70%. Both patterns suggest that the seniority classification remains essentially stable across the diffusion of GenAI.

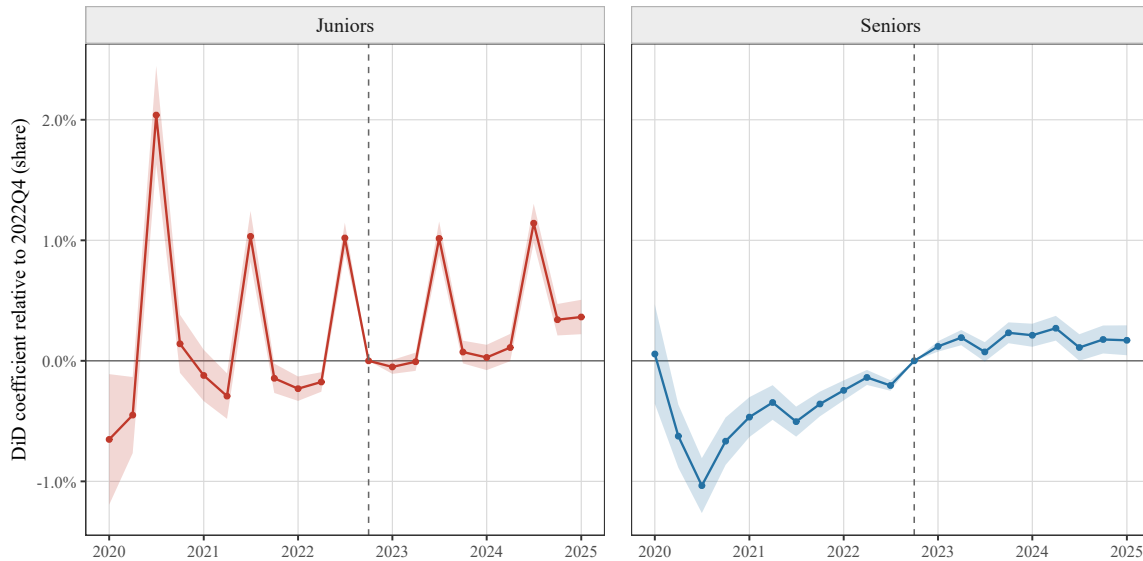


Figure A.9: Event Study: Differential Evolution of Matching-Keyword Shares

Notes: This figure plots quarterly event-study coefficients β_k from the specification $\text{Share}_{f,t}^{\text{match}} = \alpha_f + \alpha_t + \sum_{k \neq 2022\text{Q4}} \beta_k \cdot \mathbf{1}\{t = k\} \cdot \text{Adopter}_f + \varepsilon_{f,t}$, estimated separately for juniors (left panel: share of positions with any junior-leaning keyword in title) and seniors (right panel: share with any senior-leaning keyword). Regressions are weighted by the number of positions per firm-quarter cell, and standard errors are clustered at the firm level. Shaded areas denote 95% confidence intervals. The vertical dashed line marks the reference quarter 2022Q4 (the last pre-period quarter). Sample: quarters 2020Q1–2025Q1.

A.6.2 Worker Experience by Seniority and Adoption Status

A potential concern with our seniority-based analysis is that the classification of positions may itself shift with GenAI diffusion—for instance, if adopting firms differentially re-label roles as more senior after 2022. If so, the decline in junior employment at adopting firms could partly reflect mechanical reclassification rather than genuine changes in labor demand. To assess this possibility, we examine the evolution of actual worker experience among juniors and seniors, separately at adopting and non-adopting firms. We measure experience as *Labor Market Experience* (LME): the number of years since the worker’s first observed position. For each firm-quarter-seniority cell, we compute the worker-weighted average LME and then aggregate across firms by adoption status.

Figure A.10 plots average LME over time, separately for juniors (left panel) and seniors (right panel). For each seniority group, the adopter and non-adopter series evolve in close parallel. The pre-existing level gap between adopters and non-adopters—about

0.56 years for juniors and 0.28 years for seniors at 2022Q4—remains roughly constant throughout the sample, with no visible break around the ChatGPT launch. The upward trend in each series reflects the gradual aging of the observed workforce: average LME rises by approximately 2.8 years in all four series over the five-year window.




Figure A.10: Average Labor Market Experience by Seniority and GenAI Adoption Status

Notes: This figure plots average Labor Market Experience (LME), defined as years since the worker’s first observed position, for juniors (left panel) and seniors (right panel) at GenAI-adopting (solid) and non-adopting (dashed) firms over time. Averages are worker-weighted within each firm-quarter cell and then aggregated across firms by adoption status. The vertical dashed line marks the launch of ChatGPT in November 2022. Sample: quarters 2020Q1–2025Q1.

To isolate differential dynamics from the common time trend, we estimate, separately for juniors and seniors, the specification

$$\overline{\text{LME}}_{f,t} = \alpha_f + \alpha_t + \sum_{k \neq 2022\text{Q4}} \beta_k \cdot \mathbf{1}\{t = k\} \cdot \text{Adopter}_f + \varepsilon_{f,t},$$

where f indexes firms and t indexes quarters. The regression is weighted by the number of workers in each firm-quarter cell, and standard errors are clustered at the firm level. Each coefficient β_k captures the adopter–non-adopter difference in average LME at quarter k , relative to the baseline quarter 2022Q4.

Figure A.11 reports the results. For juniors, the coefficients are tightly centered around zero throughout the sample, with no discernible break at the ChatGPT launch. For se-

niors, the pre-period coefficients are slightly positive and drift gently toward zero by late 2022; post-2022Q4 they turn mildly negative, reaching approximately -0.04 years by 2025Q1. This indicates that, relative to 2022Q4, the average LME of seniors at adopting firms declined by about 0.04 years compared to non-adopting firms by 2025Q1. The magnitude is economically small, corresponding to less than three weeks of experience, or roughly 0.3% of the average post-period LME for seniors (approximately 14 years).

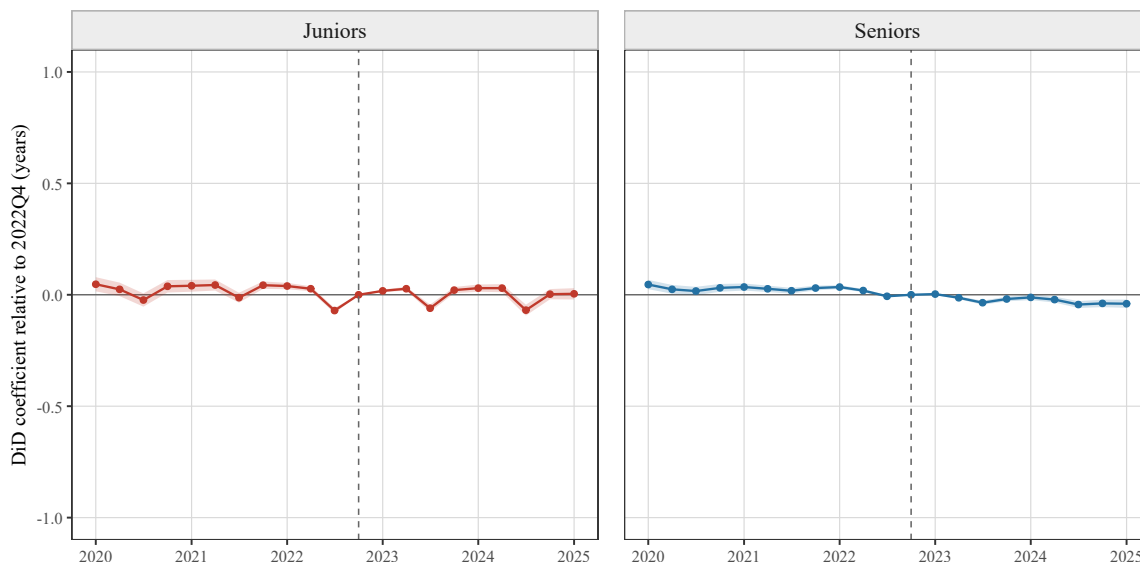


Figure A.11: Event Study: Differential Evolution of Labor Market Experience

Notes: This figure plots quarterly event-study coefficients β_k from the specification $\overline{\text{LME}}_{f,t} = \alpha_f + \alpha_t + \sum_{k \neq 2022\text{Q4}} \beta_k \cdot \mathbf{1}\{t = k\} \cdot \text{Adopter}_f + \varepsilon_{f,t}$, estimated separately for juniors (left panel) and seniors (right panel). The outcome is the firm-quarter worker-weighted average LME. Regressions are weighted by the number of workers per firm-quarter cell, and standard errors are clustered at the firm level. Shaded areas denote 95% confidence intervals. The vertical dashed line marks the reference quarter 2022Q4 (the last pre-period quarter). Sample: quarters 2020Q1–2025Q1.

A.7 College Graduates Unemployment

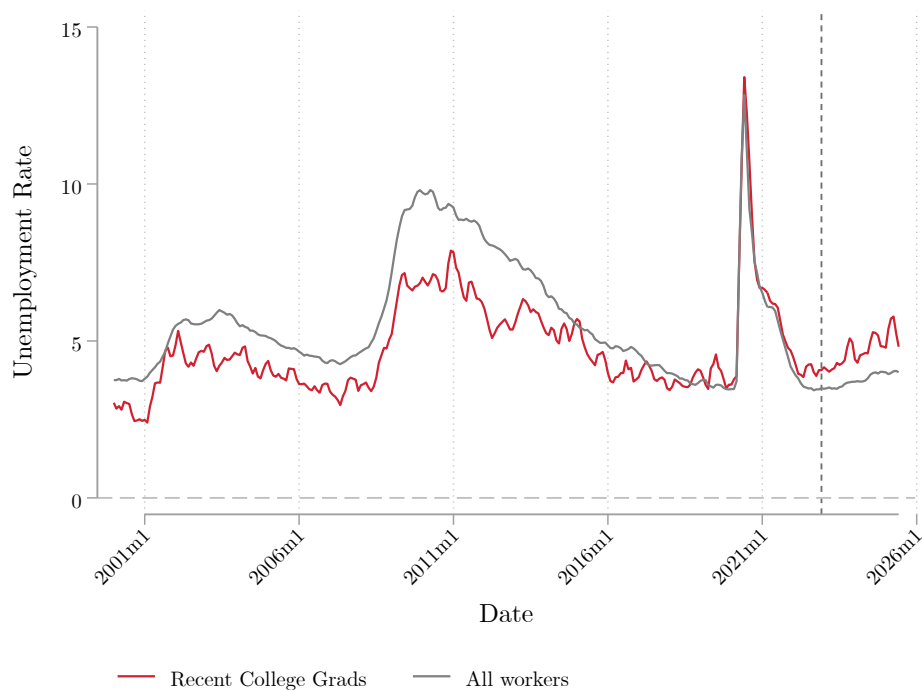


Figure A.12: Unemployment Rates for Recent College Graduates vs. All Workers

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

A.8 Robustness of Adoption Measure: Postings and Worker Position Descriptions

The baseline adoption indicator Adopt_i in the main text is identified from explicit “GenAI integrator” job postings (Section 4.3). This appendix re-estimates the headline triple-difference of Section 5.2.3 using an *augmented* adoption indicator, $\text{Adopt}_i^{\text{aug}}$, that flags firm i as an adopter if it either (a) appears as an adopter under the baseline postings-based definition, or (b) has at least one worker whose own LinkedIn *position* description is classified as an integrator role. The augmented measure helps address two residual concerns about the baseline definition. First, a fraction of GenAI integrator postings may be “ghost” or signaling postings that never translate into on-the-job integration work. Second, “silent adopters”—firms that integrate GenAI through internal reskilling or in-house tooling rather than through external hiring—can be missed by a postings-only definition.

The position-description component is constructed in close parallel to the postings-based measure (see Section 4.3 for the full pipeline): we first filter the 2021–2025 worker position spells to those whose free-text description contains at least one GenAI-related keyword, and we then ask an LLM (Llama-3.3-70B-Instruct, run with temperature = 0 and a fixed seed) to keep only those whose description indicates that the worker integrates GenAI into a product, workflow, or pipeline. A firm is flagged by component (b) above if it has at least one such position.

Across the 281,898 firms in the analysis sample, the postings-based definition flags 10,433 adopters (3.71%). Adding the position-description component flags an additional 34,387 firms not reached by the postings data, for a union of 38,441 augmented adopters (13.64%). The augmented measure thus significantly expands the treated group relative to the baseline. Figure A.13 reports the triple-difference event-study profile from Equation 12, replacing Adopt_i with $\text{Adopt}_i^{\text{aug}}$. The pattern is qualitatively unchanged: The magnitude is very similar to the baseline estimate of about -0.09 in Figure 7. Even when the treated group is expanded sevenfold to include firms whose adoption is visible only through workers’ own descriptions of what they do, the seniority-biased contraction in junior employment after the launch of ChatGPT remains intact.

It is worth noting that the position-description data are considerably less complete and less informative than the job-posting texts on which the baseline measure is built. Only




Figure A.13: Triple-Difference Estimates—Augmented Adoption Measure

Notes: This figure reproduces Figure 7, replacing the postings-based adoption indicator $Adopt_i$ with the augmented indicator $Adopt_i^{aug}$, which additionally treats a firm as an adopter if any of its worker position descriptions in the 2021–2025 window is classified as an integrator role by Llama-3.3-70B-Instruct. Specification, fixed effects, and inference follow Equation 12. Standard errors are clustered at the firm level.

about 42% of U.S. position spells in the 2021–2025 window carry a free-text description at all, and those that do are often short or formulaic.

A.9 API Prompts

A.9.1 Prompt: Identify Job Postings for GenAI 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.9.2 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):

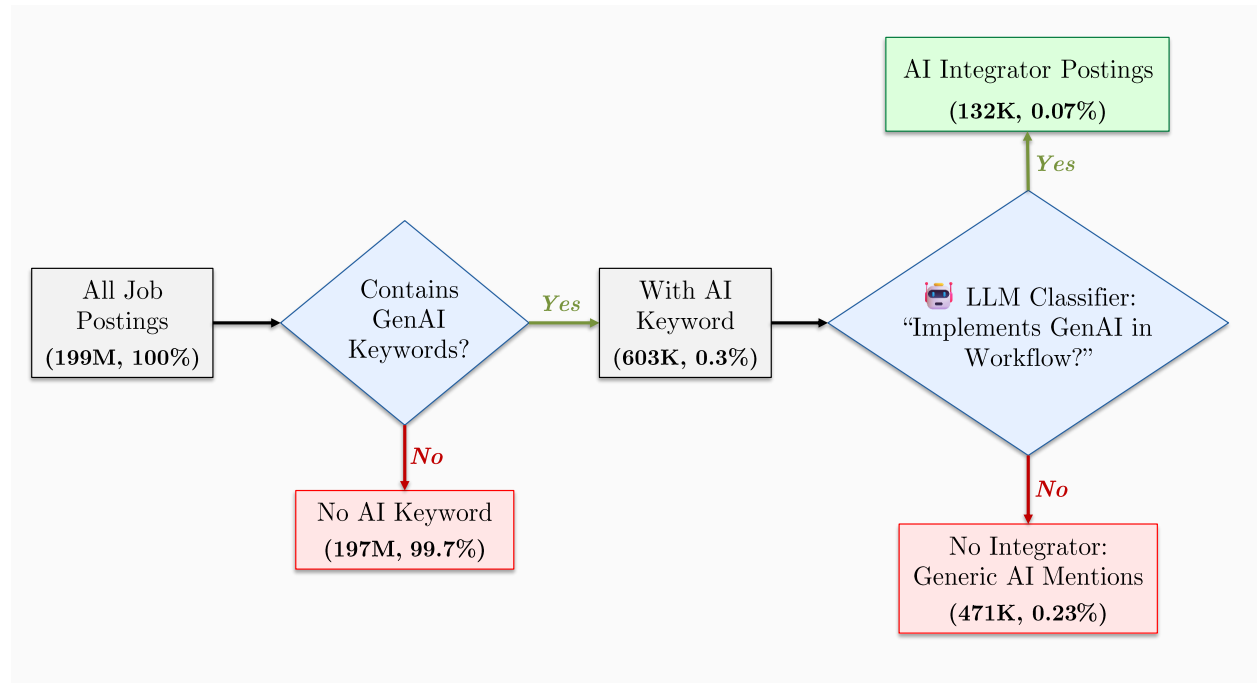
- Return EXACTLY {n} lines.
- Each line contains ONLY one integer in 1..5 for the corresponding line below.
- Do NOT include any keys, bullets, indexes, punctuation, or extra text.
- Do NOT include blank lines.
- STOP OUTPUT immediately after printing the {n}th line.

NAMES (one per line, in order):

{names_block}

"""

A.10 Detecting GenAI Integrator Postings—Graphical Illustration



A.11 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.11 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.12 Task Extraction Pipeline

This appendix documents the pipeline that maps each of the 355,013 raw job-postings in our 1,000-firm sample to a subset of the 19,265 standardized O*NET tasks, producing the 2,748,186 (posting, task) pairs used in Section 5.1.1.

The pipeline has two stages. Stage 1 is a classical information-retrieval step that, given a posting’s raw text, narrows the 19,265-task universe to the $K = 50$ O*NET tasks most textually similar to the posting. Stage 2 is a large language model (LLM) that reads the posting together with these 50 candidates and selects the subset of candidates it judges to be genuinely required by the posting, tagging each selection with a confidence label. The retrieval stage is deliberately occupation-agnostic—we do *not* restrict candidates to tasks listed in O*NET under the posting’s Revelio-assigned 6-digit SOC code—because Revelio’s SOC assignment is itself imperfect and many postings span tasks from adjacent occupations. Stage 2 is responsible for enforcing relevance and is the sole source of the final match decision.

Stage 1: Multi-view TF-IDF retrieval Rather than hashing the posting into a single bag of words, the retriever scores each of the 19,265 O*NET task strings against three views of the posting and combines the three similarities into a weighted score. The three views are:

1. the **full raw text** of the posting, after stripping boilerplate (equal-employment-opportunity statements, benefits enumerations, legal disclaimers);
2. the **duties block**, extracted by locating section headers such as “Responsibilities,” “What you’ll do,” “Key tasks,” or “Duties,” and taking the following paragraph(s) up to the next section header;
3. the **requirements block**, extracted analogously from headers such as “Qualifications,” “Requirements,” or “What we’re looking for.”

When a duties block is detected, the three views are combined with weights (0.40, 0.45, 0.15); otherwise the full-text view carries 0.85 and the requirements view 0.15. Each view is built from TF-IDF over (1, 2)-grams after lowercasing and stop-word removal; the posting’s job title is boosted by a factor of two in every view. Final candidates are the $K = 50$ O*NET tasks with the highest combined similarity, returned without any SOC restriction.

This retrieval step is deterministic, runs locally on a sparse matrix, and serves purely to bound the LLM’s search space so that the Stage 2 prompt remains short enough for batch-API throughput.

Stage 2: LLM selection with constrained JSON output The LLM is OpenAI’s `gpt-4o-mini`, queried through the OpenAI Batch API (24-hour asynchronous window; batches of up to 10,000 requests per JSONL file). For each posting we send one request whose content is the posting’s raw text and the Stage-1 candidate list (fifty pairs of `task_id` and task description). The response must conform to a JSON schema with `"strict": True` and the following structure:

```

{
  "seniority":          "junior" | "senior",
  "years_experience":  integer | null,
  "likely_occupation": string,
  "extracted_tasks":  [
    { "task_id":      <enum of the 50 candidate IDs>,
      "confidence":  "HIGH" | "MEDIUM" | "LOW" }
  ]
}

```

The `task_id` field is constrained at token-sampling time to one of the fifty candidate identifiers, which prevents the model from hallucinating task IDs outside the retrieved set. In our output, *all 2,748,260* returned task IDs satisfy this constraint.

The system prompt, reproduced verbatim, instructs the model to perform three classifications (seniority, years of experience, likely occupation) and a task-selection step:

You are an expert labor economist analyzing job postings. For each posting, you must:

1. SENIORITY: Classify as “junior” or “senior” based on Revelio’s 7-tier ordinal scale. Mentally place the posting into one of these 7 levels, then collapse to junior/senior:

L1 Entry Level (intern, trainee, entry-level)	→ JUNIOR
L2 Junior Level (jr., associate analyst, bookkeeper)	→ JUNIOR
L3 Associate Level (senior individual contributor, attorney)	→ SENIOR
L4 Manager Level (manager, superintendent, lead)	→ SENIOR
L5 Director Level (director, head of, VP of a function)	→ SENIOR
L6 Executive Level (managing director, partner, SVP)	→ SENIOR
L7 Senior Executive (CEO, CFO, COO, president)	→ SENIOR

2. YEARS OF EXPERIENCE: Extract the minimum years of experience required as an integer. Return null if not mentioned.

3. LIKELY OCCUPATION: Provide your best guess of the O*NET occupation title.

4. O*NET TASKS: From the candidate task list below, select which tasks are genuinely required by this posting.

- You MUST only return `task_id` values that appear in the candidate list. Do NOT invent or modify task IDs.
- Be SELECTIVE: only include tasks clearly implied by the posting duties/responsibilities.
- Typical range: 5-20 tasks per posting.

- For each task, rate confidence: HIGH (explicitly described), MEDIUM (strongly implied), LOW (plausible but not stated).

Return ONLY valid JSON matching the schema exactly.

Batches are submitted sequentially so that the in-flight token budget stays below the organisation quota. A manifest file records the `batch_id`, submission time, and completion state of every batch. Failed or partially completed batches are resubmitted; successfully completed batches are skipped on reruns. Stage-1 retrieval artefacts (candidate lists per posting) are cached, so reruns of Stage 2 (e.g., after prompt changes) do not require recomputing TF-IDF.

The pipeline produces two parquet files: a *posting-level* table with one row per posting (`job_id`, `rcid`, `post_date`, `onet_code`, `job_title`, LLM-inferred seniority, `years_experience`, `likely_occupation`, and `n_tasks`), and a *task-level* table with one row per extracted (posting, task) pair (`job_id`, `task_id`, `task_text`, `confidence`, `valid_id`). Table A.4 reports summary statistics.

Table A.4: Task Extraction Output: Summary Statistics

<i>Posting-level</i>	
Postings processed	355,013
classified junior	158,139
classified senior	196,874
Tasks per posting, mean (s.d.)	7.74 (2.65)
Tasks per posting, median	8
<i>Task-level</i>	
Total (posting, task) pairs	2,748,260
Distinct O*NET tasks appearing	16,342 of 19,265 (84.8%)
Task IDs within the candidate enum	100.00%
<i>Confidence distribution</i>	
HIGH (explicitly described)	58.4%
MEDIUM (strongly implied)	40.1%
LOW (plausible but not stated)	1.4%

Three features of the output are worth highlighting. First, the LLM respects the JSON schema perfectly: every returned `task_id` is one of the 50 candidates retrieved in Stage 1, so there are no hallucinated tasks. Second, the output is broad—84.8% of the O*NET task universe appears at least once—but selective within each posting: the median bundle has eight tasks, consistent with the prompt’s “typical range 5–20” instruction. Third, the vast majority of selections (98.5%) are labelled HIGH or MEDIUM confidence; in our main analysis we use all three tiers, but all results are robust to dropping the 1.4% of LOW-confidence matches.

Our pipeline is motivated by the task-extraction step in [Hampole et al. \(2025\)](#), but the two exercises differ in input and target. They extract AI *applications* from résumé text and map each application to the single nearest O*NET task via sentence-embedding cosine similarity, in order to construct a firm-level AI-exposure score. We instead extract the *multi-task bundle* of a job posting—the set of O*NET tasks the posting is asking the worker to perform—because our object of interest is the composition of labour demand and

how it shifts with GenAI adoption. The two designs share the use of the O*NET task taxonomy as a standardized vocabulary and the reliance on an LLM to bridge the gap between free-form text and structured task codes.

A.13 Sectoral and Geographical Patterns by GenAI Adoption

Sectoral Distribution of Adopters: We provide here more detailed evidence on the sectoral distribution of GenAI adopting firms. Figure A.14a documents the share of firms in each major sector that have adopted AI, while Figure A.14b 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.14: Sectoral Distribution of Adopters

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).

Geography of Adopters: Figure A.15 shows more details on the distribution of adopters across U.S. states. As expected, adoption is highly concentrated in California, which alone accounts for about 26% of all adopters. The top five adopter states are California, New York, Texas, Massachusetts, and Virginia, all of which are technology-intensive regions. This pattern highlights that AI adoption is not exclusively a Silicon Valley phenomenon but instead spans several large, tech-heavy states.




Figure A.15: Share of U.S. Adopters by State

A.14 Triple-Diff Without Industry-by-Time-by-Seniority Fixed Effects

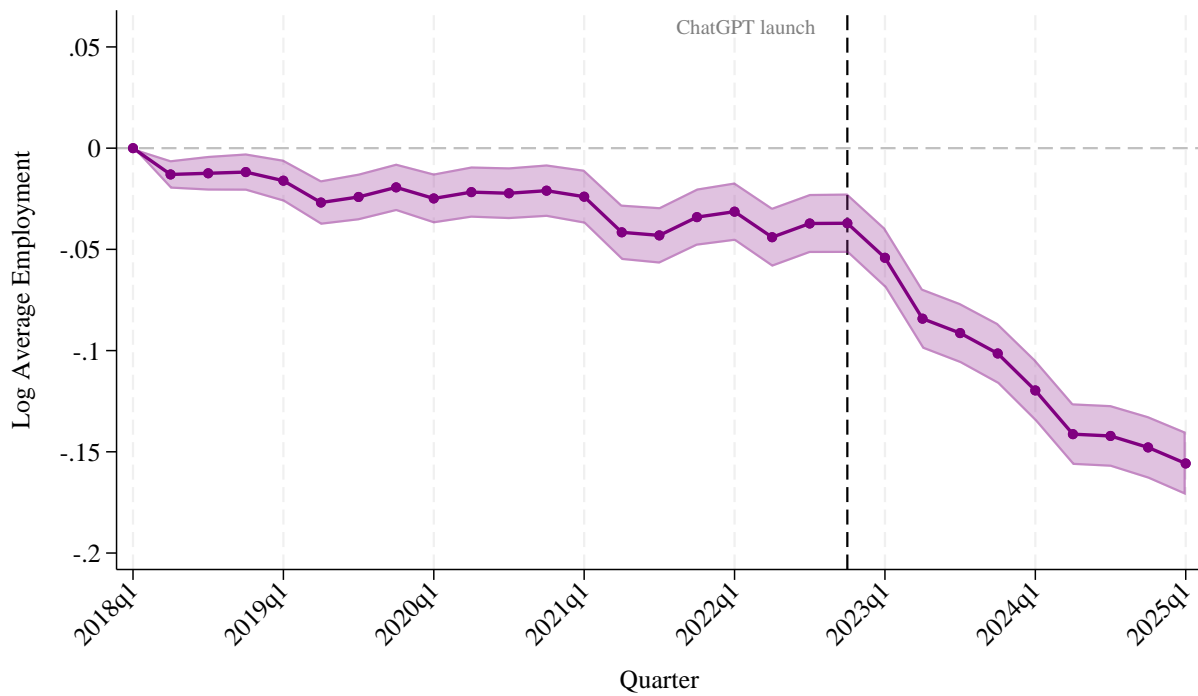


Figure A.16: Triple Differences (Without Industry-by-Time-by-Seniority Fixed Effects)

Notes: This graph shows the results of the same exercise as in Figure 7, excluding the industry-by-time-by-seniority fixed effects.

A.15 Main DiD Including the Largest Firms

As described in Section 4.1, our main sample excludes the approximately 800 largest firms in the Revelio data (ranked by post-2021 new hires) to prevent the estimates from being driven by a handful of extreme outliers. In this appendix, we re-estimate the main DiD specification from Equation 11 on a sample that adds these firms back. Acquiring and processing the full raw job-postings corpus for these firms—needed to classify their postings as GenAI integrator vacancies—is computationally prohibitive given the scale of their posting volume. However, because these firms are very likely to qualify as adopters under our definition given their size and posting activity, we include them in this robustness check as adopters.

Figure A.17 reports the resulting coefficients. The pattern is consistent with Figure 6: the junior coefficients are stable between 2019 and 2023 and then decline sharply, falling by roughly 9–10 additional log points through 2024Q3, while senior employment continues to expand throughout the sample. The sharp break around the diffusion of GenAI—the central pattern documented in the paper—is preserved.

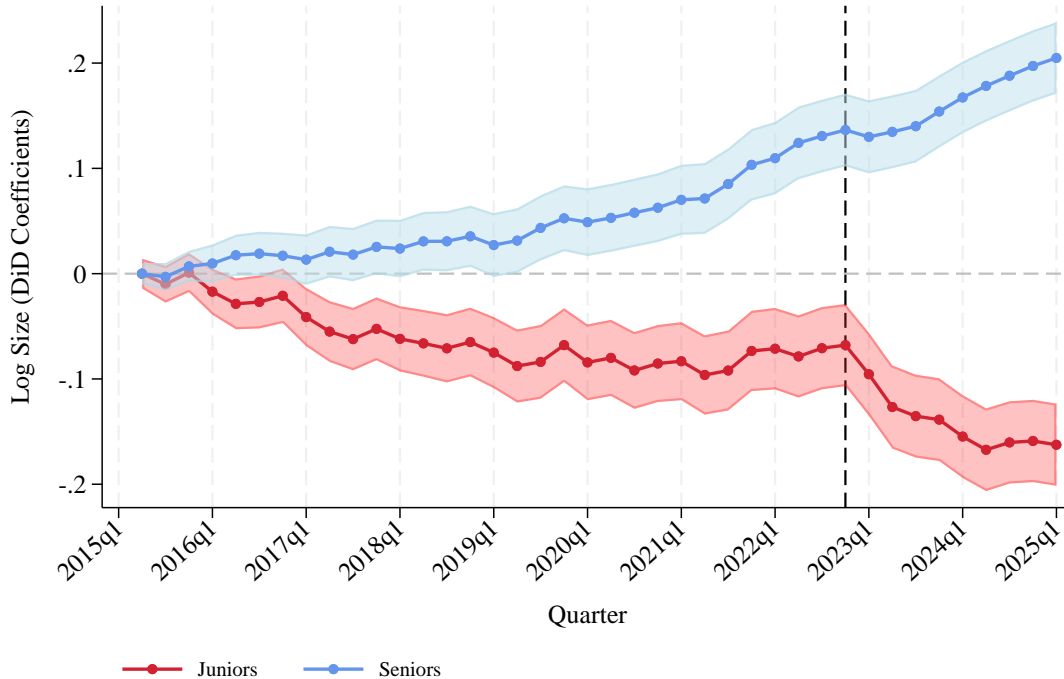


Figure A.17: Main DiD Including the Largest Firms

Notes: This figure re-estimates the DiD specification in Equation 11 on the sample used in Figure 6 augmented with the approximately 800 largest firms (ranked by post-2021 new hires) that are excluded from the main analysis. Because their firm-level adoption date is poorly identified, these firms are treated as GenAI adopters. Standard errors are clustered at the firm level.

A.16 Mentions of “AI” on Earnings Calls

Figure A.18 shows that mentions of “AI” on U.S. firms’ earnings calls began rising steeply as early as 2022Q4, tripling by 2023Q2. This pattern is referenced in Section 5.2 as evidence that firms were aware of and responding to the diffusion of GenAI almost immediately after the release of ChatGPT, consistent with the forward-looking interpretation of the junior employment decline.

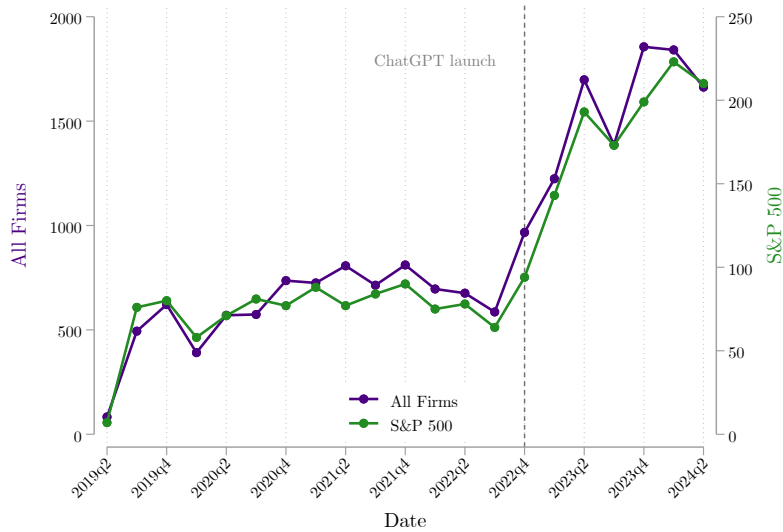


Figure A.18: Mentions of “AI” on Earnings Calls

Notes: Reproduced from Sherwood News (2024), using data from FactSet CallStreet earnings transcripts (FactSet Research Systems Inc., 2025).

A.17 Early 2023 Labor-Market Adjustments to GenAI—Anecdotal Evidence

In this appendix, we compile narrative evidence from early 2023 documenting how major U.S. companies began adjusting their hiring practices and internal workflows, explicitly attributing these changes to the adoption or anticipated impact of generative AI.

IBM: In May 2023, IBM’s CEO told *Bloomberg* that the company “*expects to pause hiring for roles as roughly 7,800 jobs could be replaced by Artificial Intelligence in the coming years*” (Reuters, 2023a).

Dropbox: In April 2023, Dropbox’s CEO announced that the company will reduce its workforce by 16 percent. According to his letter, one of the main explanations is that “*the AI era of computing has finally arrived. We’ve believed for many years that AI will give us new superpowers and completely transform knowledge work.*” (Houston, 2023).

Meta: In February 2023, Meta created a “*top-level product group*” to “*turbocharge*” generative-AI work and then announced 10,000 layoffs and cancellation of 5,000 open roles (March 2023) while emphasizing a pivot to AI (Reuters, 2023b,c).

CNET: In March 2023, shortly after revealing its use of AI-generated articles, CNET laid off about 10 percent of its staff. Parent company Red Ventures cited a refocus on search-optimized content and appointed the editor-in-chief to lead AI strategy, signaling ongoing integration of generative AI (Mia Sato, 2023).

Nuance and Microsoft: In March 2023, Nuance, a Microsoft subsidiary, launched *DAX Express*, a GPT-4–powered tool that automates clinical documentation from patient interactions, aiming to reduce physicians’ administrative workload and burnout (Medical Design & Development Staff, 2023).

Morgan Stanley: In March 2023, the Wealth Management division launched a new GPT-4–powered assistant. According to its co-president, the technology was meant to help “*freeing up valuable time for financial advisors to do what they do best—serve their clients*” (Morgan Stanley, 2023).

Wendy’s: In May 2023, Wendy’s partnered with Google Cloud to pilot *FreshAI*, a generative-AI system automating drive-thru orders using large language models to handle complex, customizable requests and improve accuracy and speed (The Wendy’s Company, 2023).

Presto and CKE Restaurants: Similarly, in May 2023, Presto Automation expanded its partnership with CKE Restaurants to deploy *Presto Voice*, an AI system automating drive-thru orders at Carl’s Jr. and Hardee’s. After pilots showed higher sales and efficiency, CKE began offering the technology to franchisees nationwide (TMCnet, 2023).

A.18 Most Common High- and Low-Exposed Occupations by Industry

Table A.5: 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.19 Heterogeneity by Educational Background

This appendix examines whether the decline in junior employment varies by workers' educational background, as referenced in Section 5.3. We first describe the construction of the school-quality variable, then present the heterogeneity results.

A.19.1 Construction of the School-Quality Variable

To capture educational background, we construct a position-level measure of school quality. For each position, we assign the institution where the worker most recently studied, provided that the education ended no later than one year after the job start. If no such record exists, we use the first recorded institution, provided it began before the job start date. Each institution is then assigned a GPT-4-based quality rating on a five-point scale, produced using OpenAI's GPT-4o-mini model (see Supplemental Appendix A.9.2 for the full prompt). A score of 1 corresponds to Ivy League and other globally elite institutions; 2 to highly respected international universities; 3 to strong national or regional schools; 4 to lower-tier but standard institutions; and 5 to minimally selective or non-accredited institutions.

A.19.2 Results

We re-estimate Equation 11 for junior employment separately within each of the five school-quality tiers. For clarity of presentation, we use a single post-treatment indicator for quarters after 2023Q1 instead of a full set of time dummies.

The results, shown in Figure A.19, reveal a U-shaped pattern. Juniors from tier-3 universities experienced the steepest relative declines in employment, while those from tiers 2 and 4 also saw significant reductions. Juniors from tier-1 and tier-5 universities experienced the smallest decline.³⁷ Figure A.20 reports the time-series version of these estimates, with a full set of time dummies in place of the single post-2023 indicator.

To help interpret this pattern, Figures A.21 and A.22 show average predicted salaries³⁸ and GenAI exposure levels of junior positions across school-quality tiers. As expected, there is a monotonic positive rela-

³⁷In the analysis, we exclude positions where the worker's most recent educational entry is reported as a high school (approximately 5 percent of positions). Including these positions, based on high school quality, has virtually no effect on the results.

³⁸Actual position-level salaries are unavailable. However, Revelio Labs provides a predicted salary variable imputed using position-specific information. For more details, see: <https://www.data-dictionary.reveliolabs.com/methodology.html#salary>.




Figure A.19: Heterogeneity by School Quality

Notes: Each bar reports the estimated coefficient from re-estimating Equation 11 for junior employment separately within each of the five school-quality tiers. The specification replaces time dummies with a single post-treatment indicator for quarters after 2023Q1. Estimation begins in 2018 to exclude early-sample trends. Standard errors are clustered at the firm level.

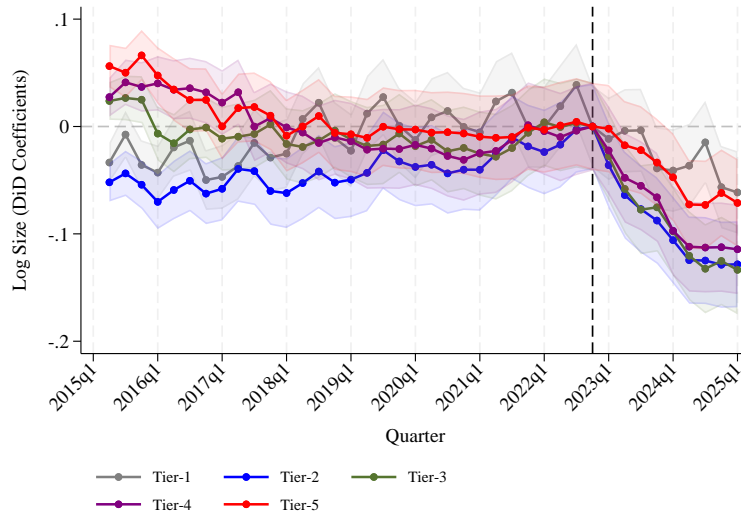


Figure A.20: DiD by Educational Background—Time Series (Juniors)

Notes: This figure plots β_j from Equation 11, estimated separately for each of the five school-quality tiers. Standard errors are clustered at the firm level.

tionship between juniors' salaries and the prestige rank of their alma mater. GenAI exposure also increases monotonically with school quality. This pattern may partly account for the attenuation in the decline from tier 3 to tier 5, but it does not help explain the increase in the magnitude of the effect from tier 1 to tier 3.

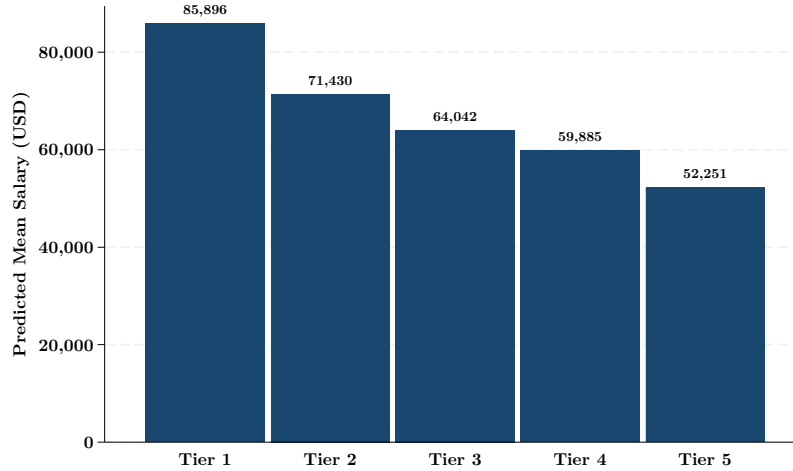


Figure A.21: Predicted Salary by School Quality (Juniors, 2022)

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




Figure A.22: 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.

A.20 Monetary-policy sensitivity: data and impulse responses

Shock series. We use the high-frequency identified monetary-policy surprises of [Jarocinski and Karadi \(2020\)](#), who decompose FOMC announcement-window co-movements of federal funds futures and the S&P 500 into a “pure” monetary-policy component and a central-bank information component. We use the monthly pure-policy median series (`MP_median`) and aggregate to the quarterly frequency as the in-quarter sum. The resulting series is positive on net tightenings and negative on net easings and is, by construction, orthogonal to information shocks. For the main-text regression we use shocks at $t \in 2015\text{Q1}–2021\text{Q3}$ so that the forward one-quarter outcome window closes no later than 2021Q4; for the local projections we restrict each horizon h to $t + h \leq 2021\text{Q4}$.

Sample. The LinkedIn/Revelio Labor firm-quarter panel covers 284,384 firms observed 2015Q1–2024Q4. Firm-level adoption (“`Adopteri`”) is defined as in Section 4.3.2. The main-text outcomes are the forward one-quarter log-change in total headcount ($\log L_{i,t+1} - \log L_{i,t}$) and in junior headcount (seniority levels 1–2; $\log L_{i,t+1}^{1,2} - \log L_{i,t}^{1,2}$), enforced by a strict one-calendar-quarter gap and winsorised per outcome at the 0.5/99.5% level. The local-projection outcomes below use the cumulative log-change $\log L_{i,t+h} - \log L_{i,t-1}$ at each horizon h .

Local-projection impulse responses. To verify that the null result in Table 4 is not horizon-specific, we estimate the lag-augmented Jordà local projection ([Jordà, 2005](#); [Montiel Olea and Plagborg-Møller, 2021](#))

$$\log L_{i,t+h} - \log L_{i,t-1} = \alpha_i^h + \gamma_t^h + \beta_h (\text{Adopter}_i \times \varepsilon_t^m) + \sum_{\ell=1}^{12} \theta_{h,\ell} \Delta \log L_{i,t-\ell} + \eta_{i,t}^h, \quad (35)$$

for horizons $h = 0, \dots, 8$ quarters. The specification includes 12 quarterly lags of log-employment growth as controls, following [Montiel Olea and Plagborg-Møller \(2021\)](#), and firm and year-quarter fixed effects; standard errors are Cameron–Gelbach–Miller two-way cluster-robust at the firm and year-quarter level. We also report the pooled impulse response from the analogous specification without the adopter interaction and with firm fixed effects only (time fixed effects would absorb ε_t^m).

Figure A.23 plots the pooled (top row) and adopter-interaction (bottom row) impulse responses with 95% confidence bands. The interaction β_h is statistically indistinguishable from zero at every horizon for both total and junior employment, consistent with the one-quarter sensitivity estimate in the main text.

Lag-augmented LP-IRFs (MOPM 2021; L=12 qtrly lags of $\Delta \log L$): LinkedIn employment · JK(2020) MP_pure shock · 2015Q1–2021Q4 (pre-2022)

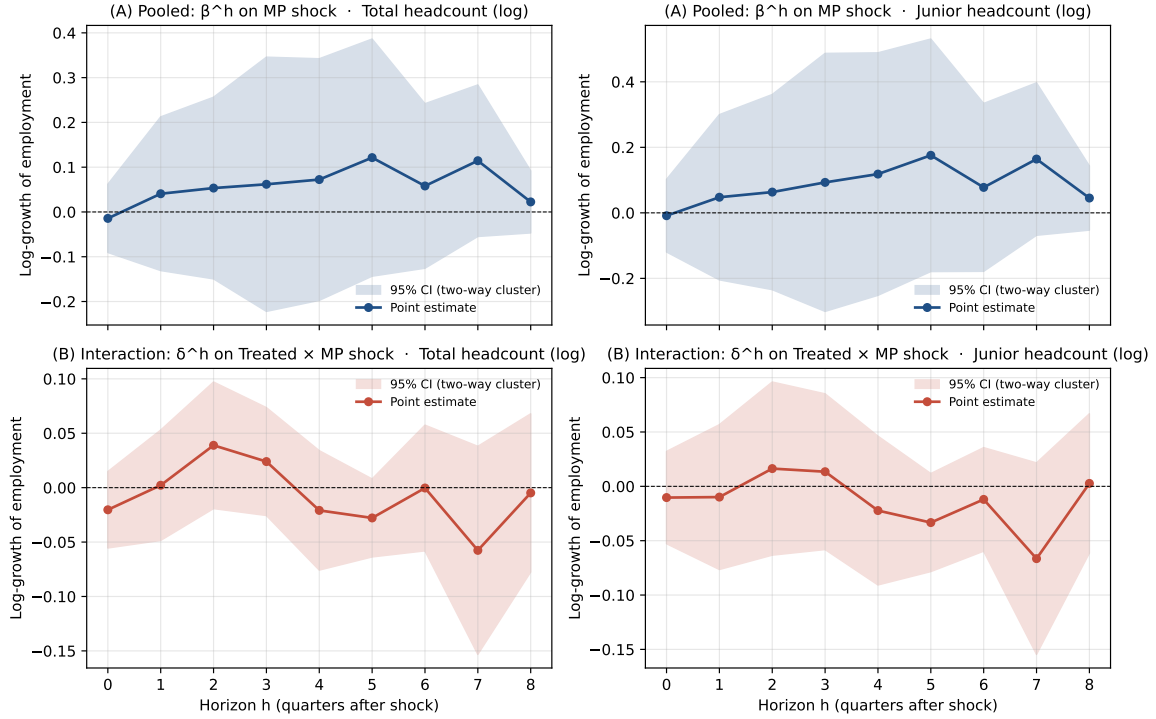
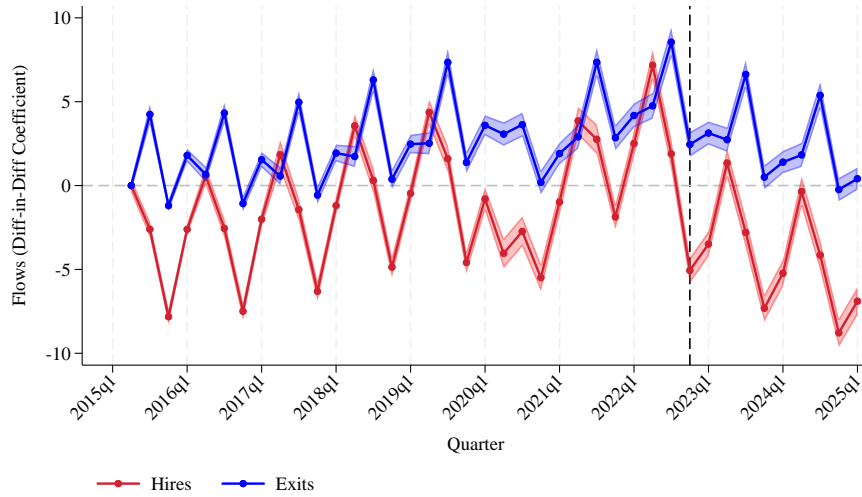



Figure A.23: Impulse responses of LinkedIn employment to Jarocinski and Karadi (2020) pure monetary-policy shocks, lag-augmented Jordà local projection, pre-2022 sample. Top row: pooled response β_h from $\log L_{i,t+h} - \log L_{i,t-1} = \alpha_i^h + \beta_h \varepsilon_t^m + \text{lag controls} + \eta$, firm fixed effects only. Bottom row: adopter-interaction response β_h from equation (35), firm and year-quarter fixed effects. Shaded bands are 95% confidence intervals using Cameron–Gelbach–Miller two-way cluster-robust standard errors at the firm and year-quarter level.

A.21 DiD for Hires and Exits—Time Series (Juniors)



(a) Not Seasonally Adjusted



(b) Seasonally Adjusted (LOWESS)

Figure A.24: DiD for Hires and Exits—Time Series (Juniors)

Notes: This figure plots β_j from Equation $y_{it} = \alpha + \sum_{j=2015Q2}^{2025Q1} \beta_j \mathbf{1}\{t = j\} \times \text{Adopt}_i + \delta_t + \gamma_i + \xi_{pt} + \varepsilon_{it}$, which corresponds to Equation 16 in Section 5.6, but replaces the single post-2023Q1 indicator with a full set of quarter dummies. Panel (a) presents the raw coefficient estimates, while Panel (b) shows estimates adjusted for seasonal variation by smoothing the coefficients using the LOWESS method (bandwidth = 0.5) prior to plotting. In both panels, the coefficients are normalized to zero in 2015Q2.

A.22 Event Study for Seniors

This appendix reports the staggered event-study specification of Equation 14, estimated on log *senior* employment rather than log junior employment, as a placebo-like comparison for the junior event-study results in Figure 9. Sample, timing, fixed effects, and inference are identical to the junior specification; only the outcome variable changes.




Figure A.25: Event Study—Seniors

Notes: The graph presents the estimated coefficients β_j from Equation 14 applied to senior employment, using the method of Callaway and Sant’Anna (2021) for staggered adoption. Firms with zero recorded senior employment in 2021Q1—eight quarters before the diffusion of GenAI—are excluded. Standard errors are clustered at the firm level.

Unlike the junior results, senior employment shows no post-adoption decline. The pre-trend coefficients slope modestly upward—rising from roughly -0.10 eight quarters before adoption to near zero just before the event—reflecting the secular growth of senior employment at adopting firms documented elsewhere in the paper. After adoption, the coefficients remain slightly above zero and do not decline. There is thus no break in senior employment at the event time, in sharp contrast to the clear post-adoption drop for juniors in Figure 9. This mirrors the seniority-biased pattern documented in the DiD and triple-difference specifications: the post-2022 contraction is confined to junior employment, while senior employment at

adopters continues along its pre-existing trajectory.

A.23 The Junior Hiring Decline—Heterogeneity by Sector

This appendix examines heterogeneity in the decline in junior hiring by sector. For this, we re-estimate Equation 16 for junior hires separately by sector. Results are presented in Figure A.26. 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 A.26: 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-adoption group. Standard errors are clustered by firm.