# Discussion Paper Series

IZA DP No. 18723

June 2026

# Generative AI and the Redefinition of Entry-Level Software Work

**Samuel Westby**
Northeastern University

**Alicia Modestino**
Northeastern University
and IZA@LISER

**Peiran Cheng**
Northeastern University

IZA DP No. 18723 June 2026

# Generative AI and the Redefinition of Entry-Level Software Work*

## Abstract

Generative AI may change how firms define occupations. We study this process in software development, where large language models overlap with tasks commonly assigned to junior workers. Using the near-universe U.S. online vacancy data from Lightcast, we examine how the public release of ChatGPT changed entry-level software hiring standards. Event-study and difference-in-differences estimates show a 14–15 percent relative decline in junior versus senior software developer vacancies, larger than in related technical occupations and absent in mechanical engineering. A shift-share decomposition shows that rising experience requirements were driven primarily by employers asking for more experience within the same job titles, not by asking for a different composition of titles. Remaining junior vacancies shifted toward problem solving, communication, and attention to detail, not AI-specific skills. The results show how generative AI redefines entry-level work by raising the bar for what counts as a qualified junior hire.

## JEL classification

J23, O33, J24, D83, M51, L86

## Keywords

generative AI, economics of information systems, labor demand, job vacancies, hiring standards, entry-level work

## Corresponding author

Samuel Westby
westby.s@northeastern.edu

* *Acknowledgements*: We gratefully acknowledge funding from the National Science Foundation under grant 2128743. Any views expressed are those of the authors and do not represent the official position of the NSF or National Bureau of Economic Research. Responsibility for any errors or omissions rests solely with the authors.

# 1 Introduction

The broad economic consequences of generative artificial intelligence (AI) remain theoretically and empirically unsettled. Emerging evidence suggests that generative AI can produce heterogeneous productivity gains across workers within the same occupation [16, 17, 31, 43, 44]. Recent information systems research also suggests that generative AI's performance effects depend on worker expertise and the stage of work being performed [29]. Other macroeconomic studies have shown a reduction in opportunities for entry-level employees [15, 28]. Yet it remains unclear how organizations redesign work, redefine expertise, and revise hiring standards.

Software development provides a particularly revealing setting in which to study this question. Large language models can generate code, explain syntax, suggest fixes, and support planning, implementation, and testing [19, 21, 25, 34]. These capabilities overlap closely with many tasks performed by software developers, especially at more junior levels where work is often more narrowly defined, more routine, and more easily reviewed. Early evidence from an online freelancing platform suggests that the demand for certain tasks and/or jobs fell sharply after ChatGPT's release, particularly those in writing and coding categories [19]. Related evidence from an online labor platform similarly shows that AI can shift task demand in ways consistent with both displacement and productivity effects, making labor demand an important setting for observing early AI-driven labor-market adjustment [46]. Moreover, software development is not a homogeneous occupation. Junior and senior developers contribute different bundles of tasks, and generative AI may substitute for some of those bundles while complementing others. This makes software development an ideal context for examining how a general-purpose digital technology reshapes labor demand *within* an occupation.

This distinction matters because firms do not need to respond to new technology only by cutting or expanding headcount. They can also adjust which types of workers they seek and what they require from them. In that context, job vacancies can serve as a leading indicator on the organizational responses of firms, revealing how employers are shifting hiring standards in real-time in anticipation of changing workflows. Prior research shows that posted job requirements are not static features of these roles and can be quite responsive to labor market conditions [26, 39, 40]. If generative AI disproportionately affects the tasks of junior software developers, one might expect to see a decline in the relative number of postings for these workers as well as a change in the experience requirements and/or skill contents of the junior vacancies that remain. Firms can change what they consider an adequate match for a role. This means firms may raise the required experience thresholds for similar jobs and shift toward skill requirements that complement AI-assisted work.

In this paper, we study the public release of ChatGPT in November 2022 as an empirical setting to examine how generative AI reshaped hiring standards within the software development occupation. Using a comprehensive dataset of U.S. online job vacancies from Lightcast, we compare changes in the number of openings and the skill requirements for junior software developer postings (requiring three or fewer years of experience) to more senior postings (requiring four or more years experience), before and after ChatGPT's public release. This setting works well for three reasons. First, ChatGPT initiated a sudden and widely adopted change in software development. Second, the capabilities of large language models overlap with the tasks of junior software developers. And third, online job postings allow us to observe how firms updated their hiring behavior and job requirements in near real-time.

We document three main findings. First, we align with [15, 28] and show that the public release of ChatGPT coincided with a sharp decline in demand for junior software developers relative to senior developers. Our difference-in-differences estimates show that, over the 12 months following November 2022, software developer postings requiring fewer than four years of experience fell by roughly 15 percent relative to postings requiring four or more years of experience. This shift is visible almost immediately in the monthly ratio of senior to junior postings and is comparable in magnitude to the sudden disruption that was observed previously during the onset of the COVID-19 labor-market shock. We also show that this pattern is not explained by broader trends affecting all computer and mathematical occupations. Junior software developers were more adversely affected than junior workers in other computer and mathematical occupations. Moreover, there was no impact on junior workers in other STEM occupations that have far less exposure to generative AI such as mechanical engineering.

Second, we show that employers redefined junior-level software work. In addition to employers posting fewer junior-style jobs after the popularization of large language models, they sought more experienced candidates to fill the remaining entry-level roles. We also find that the junior vacancies that remained shifted toward requiring complementary higher-order skills, like problem solving, interpersonal communication, and attention to detail. Notably, this evidence does not support the contention that AI-specific skills (like PyTorch, CNNs, etc.) made junior workers broadly more employable. Instead, generative AI appears to have increased the value of skills that *complement* AI-assisted work among junior software engineers. Prior research shows that a similar pattern emerged among online freelancing markets, where the jobs that remain after generative AI diffusion tend to be more complex [19].

Third, these effects were not uniform across organizational contexts. The decline in junior relative to senior hiring was concentrated among larger firms and larger labor markets. This is consistent with the theory that organizations with more resources and flexibility as well as those located in larger labor markets that are more competitive, are more likely to adopt and reorganize around new technologies. Moreover, the impact on junior software developers followed a U-shaped pattern such that industries with either limited software demand or software-dense sectors showed larger shifts, while industries with moderate software demand did not. This heterogeneity suggests that generative AI does not simply reduce hiring in proportion to software usage intensity. Instead, its effects depend on how software tasks are integrated into production and how easily firms can reallocate those tasks.

This paper contributes to the literature by showing how the diffusion of generative AI reshapes organizational demand for human capital through the redefinition of entry-level work. We examine how firms revise the roles, requirements, and hiring standards attached to software work. This connects emerging evidence on the heterogeneous productivity effects of generative AI to employer-side labor demand. Existing research shows that generative AI can improve software development performance, often with large gains for less-experienced workers. We show that these productivity gains can coincide with reduced demand for junior workers when firms raise the experience and skill thresholds associated with entry-level roles. The paper also contributes to the literature on technological change and skill-biased labor demand by showing that adjustment can occur *within* occupations, not only between them. In our setting, firms reduce the number of junior software vacancies and also redefine what counts as a qualified junior by increasing experience requirements within job titles and shifting the remaining junior vacancies toward skills that complement AI-assisted work. This connects to emerging work showing that generative AI is reshaping both the supply of human output and the demand for human labor [47].

Our findings also have broader implications for talent pipelines in the software industry. If the rapid adoption of generative AI significantly disrupts or even replaces the initial rung on the career ladder for software engineers, this may have long-run consequences for both entry-level software engineers and the firms that typically hire them. This is because entry-level roles typically allow workers to gain experience, learn new systems, and move into more advanced roles over time. For recent college graduates who majored in software engineering, the initial labor market disruption associated with the introduction of ChatGPT may lead to under- or unemployment in the short-run and lower earnings due to labor market scarring in the long-run. For software firms, the efficiency improvements from AI adoption may increase productivity in the short-run but disrupt the long-term supply of senior software engineers with more experienced technical expertise in the long-run. This can occur even without reducing the number of entry-level job openings if firms increase the skill requirements for entry-level positions. Finally, this within-occupation upskilling creates a moving bar for entry-level workers who are no longer qualified for these positions, and if widespread, can potentially give rise to persistent labor market mismatch that requires short-term re-training for the current cohort of software engineering graduates and/or longer-term changes to curriculum for future cohorts of computer science majors [41].

The remainder of this paper is structured as follows. Section 2 develops the theoretical argument linking generative AI to a change in relative demand *within* software development. Section 3 describes the data and empirical strategy. Section 4 presents the main results, including evidence on heterogeneity and changing experience and skill requirements. Section 5 discusses the implications of this work for theory, organizations, and policy.

# 2 Theoretical Development

While many studies on generative AI have focused on whether it will change aggregate employment, we choose to answer a more well-defined question to reduce the amount of heterogeneity and better understand the nuances of adopting this new technology. Specifically, we examine how organizations revise work requirements after the diffusion of a new information system. Generative AI may change how firms define the experience and skills required for a role. We argue that the public release of ChatGPT created a task shock in software development that led firms to adjust both the relative demand for juniors versus seniors and also the hiring standards for entry-level roles.

## 2.1 Generative AI as a Task Shock in Software Development

Software development is a particularly salient occupation for studying the labor-market effects of generative AI. This is because many basic or core software development tasks in an organization overlap with the capabilities of large language models (LLMs). These LLM systems can generate code, suggest fixes, explain unfamiliar syntax, and support brainstorming across multiple stages of the development process, including planning, implementation, and testing [19, 21, 34]. However, in the early years following the dissemination of large language models, software developers did not initially use these systems as fully autonomous substitutes. Instead, developers often used them for examples, direction, and preliminary solution paths, while continuing to evaluate outputs for correctness and trustworthiness [50]. This combination of high task overlap and semi-automation sets the stage for software development to be an ideal case-study for studying how new technologies may affect the relative productivity of different types of workers rather than entirely replacing jobs within an occupation.

Yet it was not until the public release of ChatGPT in November 2022 that generative AI became an economically meaningful task shock. Although the underlying technology had been advancing for years, ChatGPT dramatically lowered access costs by making powerful LLM capabilities available through a simple interface. Early survey evidence suggests that computer and mathematical occupations had especially high rates of generative AI use during the initial diffusion period of ChatGPT, particularly among software-related jobs, making these positions extremely vulnerable to potential labor-market adjustments [14]. For employers, these adjustments may have initially taken the form of changes in hiring practices instead of layoffs or other formal personnel reorganizations. Changing posted job requirements for new hires is likely quicker, less costly, and less disruptive to established workflows compared to redesigning teams, changing reporting structures, or reducing headcount. For that reason, vacancy postings can provide a meaningful early signal of how firms are responding to the widespread adoption of this new general-purpose technology [20, 26, 40].

## 2.2 Uneven Productivity Gains and Within-Occupation Substitution

Our theoretical argument builds on the premise that generative AI does not equally raise the productivity of all software developers – at least not initially during its introduction.

This is because junior developers are more likely to be assigned relatively narrower, more routine, and easily reviewable tasks that are exactly the kinds of tasks where generative AI appeared to be especially useful early on [32]. Recent evidence suggests that AI assistance provided a sizable productivity gain for software work and these gains were often large for the type of tasks done by less-experienced workers [16, 17, 44]. As a result, generative AI should therefore substitute more strongly for junior task bundles than for senior ones [1]. This interpretation is consistent with trends from online freelancing markets showing larger post-ChatGPT declines in automation-prone writing and coding jobs than those requiring more manual tasks or more complex work [19, 46].

In contrast, senior developer tasks appear to be more complementary to AI tools, at least in the short run, because their work is less centered on raw production of code and more focused on system-level decision-making, architecture, review, and coordination. This reflects a form of cognitive reapportionment: AI changes the location of production, judgment, and decision rights between humans and intelligent systems, making some tasks easier to delegate while increasing the value of human oversight and integration [35]. Recent field evidence even suggests that applying AI tools to more complex tasks may even slow experienced developers in realistic open-source settings, showing that AI's productivity effects are not solely positive for software development [11].

In addition, senior software developers spend substantial time on non-coding activities, such as information seeking, helping coworkers, meetings, and other collaborative tasks that are harder to automate [38]. This does not imply that senior developer tasks are completely insulated from the effects of generative AI nor are they immune to future labor market impacts arising from later iterations of this technology. We argue that during the initial public adoption period of ChatGPT, the relative value of more experienced senior workers who can supervise, evaluate, and integrate AI-assisted work is likely to increase compared to less experienced junior workers who are typically assigned more routine entry-level tasks.

This creates an important distinction between worker productivity and labor demand. A technology can increase the productivity of junior workers while reducing the number of junior workers that firms need to perform those tasks. If generative AI lowers the cost of junior tasks (routine coding, debugging, implementation), firms may require fewer units of junior labor despite the remaining juniors becoming more productive. So, productivity gains do not necessarily imply expanded entry-level hiring in the short run. This paper studies the immediate transition period of the rapid diffusion of generative AI in the years following the sudden public release of ChatGPT in November 2022. We acknowledge that over longer horizons, lower production costs may also expand total software output, create new tasks and roles, and eventually generate new entry points into the software engineering occupation [7, 8, 12]. But during the initial transition period, it is plausible that the substitution effects may dominate such that the public release of ChatGPT would initially reduce employer demand for junior software developers and the tasks that they perform relative to senior software developers.

### 2.3 From Task Reallocation to Hiring Standards and Vacancy Requirements

If generative AI changes the relative productivity of junior and senior software developers, then firms may respond by changing vacancies and hiring standards instead of a headcount reduction alone. Prior research shows that employers revise posted education, experience, and skill requirements in response to changing labor-market conditions, effectively redefining who is considered employable for a given role [41]. During weak labor markets, firms raise requirements within occupations, and those increases may appear directly in vacancy postings even before they affect employment totals [26, 40]. When labor markets tighten, posted requirements can fall again, which suggests that hiring standards are adjustable properties of the job itself [39].

Vacancy requirements therefore reveal how firms define both the work to be performed and the type of human capital considered appropriate for that work. When a new information system changes the value of tasks, firms may first respond by editing hiring requirements before making deeper changes to teams or organizational structure. Evidence from prior waves of AI adoption is consistent with this possibility. For example, a study of 1,300 publicly traded firms in China from 2007 to 2018 finds that AI adoption was associated with increased employment of nonacademically trained workers and decreased employment of academically trained workers, suggesting that AI can alter the types of workers firms demand rather than simply reducing labor demand overall [51].

Thus, the relevant outcome we focus on is not simply whether firms post relatively fewer junior vacancies, but also if they rewrite junior job requirements in ways that affect the employability of junior workers. If routine coding becomes easier to automate, firms may not entirely eliminate junior roles. Instead, they may raise the bar for who qualifies as "junior" by placing greater weight on tasks that are harder to automate, such as communication, problem solving, and attention to detail. Junior work need not disappear entirely, but its task content may shift upward, such that recent college graduates who are seeking to enter the profession are no longer qualified for these roles.

This is consistent with task-based theories of technological change, which argue that new technologies typically substitute for routine tasks while complementing non-routine skills [9]. Generative AI extends this logic to software development by reducing the cost of routine coding, debugging, and implementation tasks while increasing the relative value of judgment, coordination, and problem solving. This shift is a reallocation of cognitive responsibility across humans and AI systems [35]. Research has also shown that the labor market rewards roles that combine quantitative skills with social and higher-order cognitive skills [18]. This prior literature leads to three empirical expectations. First, if generative AI substitutes for tasks commonly assigned to junior developers, junior software vacancies should decline relative to senior vacancies after ChatGPT's public release. Second, if firms respond by redefining entry-level work, required experience should rise *within existing job titles.* Third, if the remaining junior roles are reorganized around tasks that complement AI-assisted production, junior vacancies should shift toward higher-order cognitive and interpersonal skills.

## 2.4 Where Effects Should First Emerge

The impact of generative AI on junior versus senior software hiring is unlikely to be uniform across firm types. Even if the underlying task shock is similar, firms must have the resources, managerial attention, technical infrastructure, and labor-market flexibility to translate a new tool into changed hiring standards. It is likely that any potential decline in junior relative to senior software vacancies will be strongest when adoption is more feasible (e.g., larger firms) and/or in places where firms have greater flexibility to redesign teams around AI-assisted production (e.g., larger labor markets).

One reason larger firms may exhibit changes in hiring dynamics earlier is that they are typically better positioned to adopt new technologies. Prior work shows that the use of advanced technologies like AI is concentrated among larger firms [2, 24, 37]. Recent evidence also suggests that firms may treat AI investment as a strategic response to labor-related pressures in addition to a productivity upgrade [36]. Larger firms may also have more opportunities to redesign entry-level hiring. If generative AI allows teams to produce more routine code with the same number of senior workers, large firms may be especially able to respond by reducing junior hiring while maintaining output. This does not imply that small firms are unaffected, only that the earliest and strongest hiring adjustments may appear where experimentation is easiest.

Geography may also matter for similar reasons. Larger labor markets can facilitate the spread of new technologies and work practices [13]. In software, larger metropolitan areas contain more innovative firms, larger pools of specialized labor, and more opportunities for worker relocation. We therefore expect software-dense geographies to show earlier and/or greater shifts in hiring. However, we acknowledge that not all metropolitan areas will be early AI adopters. For that reason, geography is seen as a facilitator, not a primary mechanism.

Industry context may also shape the adjustment, although the direction is less straightforward. In software-intensive industries, generative AI may directly substitute for routine junior developer tasks, making redefinition more likely. In industries where software teams are smaller or more peripheral, the same technology may either reduce the need for junior specialists or expand the scope of what small technical teams can accomplish. Because of this large scope, we treat industry differences as exploratory.

### 2.5 Summary of Expectations

Taken together, the literature implies that ChatGPT created a task shock in software development that should affect the organizational definition of junior and senior work. Because junior developers are more likely to perform narrower and more explicit tasks, generative AI should reduce relative demand for junior developers in the short run. However, the central takeaway is not only that firms post fewer junior vacancies. Firms may also raise the experience threshold for similar roles and shift remaining junior vacancies toward skills that complement AI-assisted production, such as problem solving, communication, and attention to detail. These changes should be most visible where firms have greater capacity to adopt and reorganize around generative AI, particularly larger firms and larger labor markets.

## 3 Data and Research Design

### 3.1 Real-Time Vacancy Data

Our empirical setting is the U.S. labor market for software developers. We use online job vacancy data from Lightcast (formerly Burning Glass Technologies), which represents the near-universe of online job postings scraped from over 50,000 sources. For each of the 460+ million unique U.S. job postings, Lightcast provides a six-digit Standard Occupation Code

(SOC), North American Industry Classification System (NAICS) industry code, job title, employer name, location, as well as the required level of education, years of experience, and specific skillsets (e.g., Python). Prior research has shown that these data track aggregate labor-demand patterns well and are particularly useful for studying changes in employer requirements both within occupations and at high frequencies (e.g., monthly) [4, 26, 39, 40].

These data are well suited to our research question because job vacancies encode the organizational requirements attached to work. Our goals are to measure if firms posted fewer junior software developer vacancies after the release of ChatGPT and to observe how firms revised the experience thresholds and skill requirements associated with those vacancies. Online job postings are useful because they reveal employer-side adjustments to hiring standards in near real-time, long before those adjustments show up in organizational changes. This makes vacancy data particularly appropriate for studying the redefinition of entry-level work.

The full dataset that we use spans from January 2019 to March 2025 and contains approximately 5.7 million unique U.S. job postings for software developers (SOC 15-1252). We use this longer window for descriptive context to better understand the long-term trends prior to the widespread adoption of generative AI. Our main empirical design relies on a more narrow window to capture labor market conditions during the 12 months before and after the public release of ChatGPT in November 2022.

## 3.2 Sample Characteristics and Key Measures

We restrict our primary analysis to software developer postings within a U.S. Combined Statistical Area (CSA) and exclude vacancies from U.S. territories. We classify postings by required years of experience to differentiate labor demand for more- versus less-experienced workers. *Junior* postings are those requiring three or fewer years of experience and *senior* postings are those requiring four or more years. Postings with missing experience requirements are excluded from the main analysis but we find qualitatively similar results using all job postings and predicting years of experience based on the observable characteristics of each job posting (see the appendix).

This experience-based classification operationalizes the boundary between entry-level and experienced software work. Because our theory concerns whether generative AI shifted this boundary, our analysis examines both the relative number of junior versus senior vacancies and changes in the requirements attached to software roles over time.

We aggregate the raw job posting data to the *experience-group* $\times$ *job title* $\times$ *CSA* $\times$ *month* cell level. Our dependent variable is the natural logarithm of the count of vacancies in each cell. We create a dummy variable to differentiate postings by experience group that $Junior_e$, equals one for junior-level and zero for senior-level. We also construct a post-period indicator, $PostChatGPT_m$, equal to one for months on or after the release of ChatGPT starting in November 2022. We include control variables, $X_{cjm}$, that capture differences across labor markets that may affect hiring such as population quintile rank, firm-size, and industry mix that are measured before the widespread public release of ChatGPT in November 2022 so that they are independent of the widespread generative AI technological shock.

We measure three items. First, we measure whether the relative availability of junior software vacancies changed after ChatGPT's public release. Second, we measure whether required experience increased within job titles, which would indicate that firms raised the bar for nominally similar roles rather than simply changing the mix of titles posted. Third, we examine whether the skill content of junior vacancies shifted toward capabilities that complement AI-assisted production.

The choice of November 2022 as the inflection point captures the public release of ChatGPT. Our design does not claim developers had no exposure to these tools prior to ChatGPT. For example, GitHub Copilot released a technical preview in June 2021 that acted more as a virtual "pair-programmer" than a labor substitute. ChatGPT, on the other hand, reached 100 million users within the first two months of its introduction. It became the fastest growing consumer application ever [30] with wide adoption by workers in computer and mathematical occupations [14], making it ideal as a discrete and measurable event for studying changes in employer hiring behavior for software engineers.

## 3.3 Research Design

Our empirical strategy takes four steps. First, we estimate event-study and difference-in-differences models to test whether the public release of ChatGPT coincided with a relative decline in junior software developer vacancies. This establishes whether firms shifted away from entry-level software opportunities after generative AI became widely available. Second, we compare software developers to other computer and mathematical occupations, and Mechanical Engineers (a similar technical occupation less exposed to AI), to assess whether the pattern reflects a software-specific adjustment or a broad decline in junior technical hiring. Third, we use a shift-share decomposition to determine whether rising experience requirements reflect a change in the mix of job titles posted or higher experience requirements within the same titles. Fourth, we analyze changes in required skills to assess whether remaining junior vacancies shifted toward capabilities that complement AI-assisted work.

### 3.3.1 Event Study

Our event study tests if the boundary between junior and senior software hiring shifted after the public release of ChatGPT. The key identifying assumption is that, absent ChatGPT's release, the relative gap between junior and senior software vacancies would have evolved smoothly through November 2022.

Specifically, we estimate:

$$\begin{aligned} \ln Vacancies_{cjm} &= \alpha + \delta\, Junior_{cjm} \\ &\quad + \sum_{k \neq -1} \beta_k \Big( \mathbf{1}_{\{m-m_0=k\}} \times Junior_{cjm} \Big) + \mu_m + \gamma_c + X_{cjm} + \varepsilon_{cjm}. \end{aligned} \tag{1}$$

where $Vacancies_{cjm}$ denotes the number of software vacancies in CSA $c$, job title $j$, and month $m$. $Junior_{cjm}$ is an indicator equal to one for junior vacancies. $m_0$ denotes November 2022, and $k$ indexes months relative to $m_0$, such that September 2022 corresponds to $k = -2$ and December 2022 to $k = +1$. We omit $k = -1$ (October 2022) as the reference period, so the coefficients $\beta_k$ are interpreted relative to the month immediately preceding November 2022. Further, $\mu_m$ and $\gamma_c$ denote month and CSA fixed effects, respectively, $X_{cjm}$ is a vector of controls, and $\varepsilon_{cjm}$ is a stochastic error term and $\alpha$ is the regression constant.

The coefficients $\beta_k$ trace the change in junior software vacancies relative to senior software vacancies over time. Pre-event coefficients provide a test of the parallel-trends assumption, while post-event coefficients show how quickly the relative demand for junior developers changed relative to senior developers after November 2022.

We also estimate a triple-difference event-study specification with Equation 2 to account for confounding changes in the labor market that might affect all mathematical and computer science occupations, such as the slowdown in hiring after the initial COVID-19 recovery. Specifically, we examine the trend in software job vacancies with varying experience levels (Junior vs Senior), over time (before versus after the release of ChatGPT), and relative to other computer science occupations (all other vacancies in SOC 15 except software engineers).

$$\begin{aligned}\ln(\text{Vacancies}_{cjm}) = \theta_1 \text{Junior}_{cjm} + \theta_2 \text{Software}_{cjm} + \theta_3 \big(\text{Junior}_{cjm} \times \text{Software}_{cjm}\big) \\ + \sum_{k \neq -1} \beta_k \Big[\mathbf{1}_{\{m-m_0=k\}} \times \text{Junior}_{cjm} \times \text{Software}_{cjm}\Big] \\ + \alpha + \mu_m + \gamma_c + X_{cjm} + \varepsilon_{cjm}.\end{aligned} \tag{2}$$

where the $Junior_{cjm}$ is an indicator equal to one for junior positions in a particular $CSA \times JobTitle \times Month$ cell. We also include fixed effects to absorb local labor market shocks associated with all jobs within a CSA-month, persistent differences across job titles within cities, and national time-varying shocks specific to each job title. The event-study coefficients $\beta_k$ trace the dynamic path of junior software vacancies relative to November 2022.[1]

### 3.3.2 Difference-in-Differences Specification

We next estimate the average post-ChatGPT change in junior software developer vacancies relative to senior software developer vacancies using the following difference-in-differences specification:

$$\begin{aligned} \ln Vacancies_{cjm} = \; & \alpha + \delta \, Junior_{cjm} \\ & + \theta \Big( PostChatGPT_m \times Junior_{cjm} \Big) + \mu_m + \gamma_c + X_{cjm} + \varepsilon_{cjm}. \end{aligned} \tag{3}$$

where the variables are defined the same as in Equation 1 with the addition of $PostChatGPT_m$, which is a binary indicator that equals one when $m$ is on or after November 2022. In this equation, $\theta$ is the coefficient of interest that captures the average change in junior relative to senior software vacancies after November 2022. A negative $\theta$ value indicates that entry-level software vacancies became less available relative to senior vacancies in the post-ChatGPT period.

The main assumption is that the relative difference between junior and senior software vacancies would have evolved similarly across the pre- and post-periods without the release of ChatGPT. We estimate the model over the one-year period before and the one-year period after the public release of ChatGPT 3.5 in November 2022. Summary statistics for the estimation sample are reported in Appendix Table A1, and summary statistics for the control variables are reported in Appendix Table B1. Standard errors are clustered at the job title level to allow for within-category correlation in outcomes over time.

[1] Our specification follows the broader event-study literature by embedding our event-time indicator in a high-dimensional fixed-effects framework, allowing for flexible control of local and occupation-specific shocks without requiring a large number of explicit lower-order interactions [23, 45, 49].

### 3.3.3 Synthetic Control

A central alternative explanation is that junior hiring declined across computer and mathematical occupations during the broader post-pandemic cooling of the technology labor market. To address this concern, we supplement the triple-difference comparison with a synthetic comparison group constructed from other SOC 15 occupations. This approach re-weights related occupations to better match the pre-ChatGPT junior-to-senior posting ratio and overall posting level for software developers.

We construct the synthetic SOC 15 comparison group using the 12 month pre-treatment window from November 2021 through October 2022. We restrict the weights to be nonnegative and contribute no more than 20 percent of the variation to find the combination of weights that best approximates the SOC 15-1252's pre-period junior-to-senior posting ratio and overall posting level. The resulting synthetic panel is combined with the treated SOC 15-1252 panel and used as an alternative to the aggregate SOC-15 in the difference-in-difference-in-differences analysis. A summary of occupational weights and descriptive statistics for the synthetics dataset can be found in Table A2.

### 3.3.4 Shift-Share Analysis of Experience Requirements

The event-study and difference-in-differences estimates show whether firms shifted away from junior software developer vacancies. They do not reveal whether firms also redefined the requirements attached to software roles. A decline in junior vacancies could occur because firms post fewer explicitly junior titles, because they shift toward more senior titles, or because they keep posting similar titles while asking for more experienced candidates. The last possibility is central to our argument: firms may redefine entry-level work by raising experience thresholds within similar roles. To distinguish among these possibilities, we conduct a shift-share decomposition of average required experience across job titles within the software development occupation.

Formally, let $n$ be the number of distinct job titles within SOC 15-1252. For each job title $i \in \{1, 2, ..., n\}$, let $s_{i,pre}$ and $s_{i,post}$ be the share of job title $i$ in the total number of job postings *pre-* and *post-ChatGPT*, respectively. Similarly, let $e_{i,pre}$ and $e_{i,post}$ be the average required years of experience for job title $i$ *pre-* and *post-ChatGPT*, respectively. Then, the overall average required years of experience in each period is given by:

$$E_{pre} = \sum_{i=1}^{n} s_{i,pre} \cdot e_{i,pre}; \quad E_{post} = \sum_{i=1}^{n} s_{i,post} \cdot e_{i,post}$$

To decompose the change in the overall average required years of experience ($\Delta E = E_{post} - E_{pre}$), we construct two counterfactual scenarios:

1. **Counterfactual 1: Holding Constant the Composition of Job Titles.** We calculate the overall average required experience across all postings if the distribution of job titles remained at its *pre-ChatGPT* level, but the experience requirements for each title changed to their *post-ChatGPT* values. This counterfactual average is:

   $$E_{cf1} = \sum_{i=1}^{n} s_{i,pre} \cdot e_{i,post}$$

   The difference between this counterfactual and the *pre-ChatGPT* average ($E_{cf1} - E_{pre}$) indicates the change attributable to new experience demands within the same job titles.

2. **Counterfactual 2: Holding Constant Experience Requirements.** We calculate the overall average required experience across postings if the experience requirements for each job title remained at their *pre-ChatGPT* level, but the distribution of job titles changed to its *post-ChatGPT* values. This counterfactual average is:

   $$E_{cf2} = \sum_{i=1}^{n} s_{i,post} \cdot e_{i,pre}$$

The difference between this counterfactual and the *pre-ChatGPT* average ($E_{cf2} - E_{pre}$) indicates the change attributable to a shift in the mix of job titles being posted.

By comparing these counterfactuals to the actual change in the overall average required years of experience, we can assess the relative contributions of changes in job title composition versus changes in experience requirements for specific job titles. Our analysis focuses on which counterfactual better approximates the true post-ChatGPT average, $E_{post}$.

#### 3.3.5 Skill Requirement Analysis

Finally, we examine whether the skill content of junior software vacancies changed after ChatGPT's release. For each skill, we calculate the monthly share of postings requiring that skill that are junior postings. We then estimate skill-specific linear trends using the pre-ChatGPT period and compare the observed post-ChatGPT junior share to the value predicted by the pre-period trend. This approach identifies skills that became more or less associated with junior software vacancies after ChatGPT's release, relative to their own pre-existing trajectories. We focus on whether remaining junior vacancies shifted toward AI-complementary skills or toward AI-specific technical skills.

### 3.4 Caveats for Interpretation

Two points guide the interpretation of our research design. First, online vacancies measure employer demand and posted hiring standards, not realized hires. This is appropriate for our research question because we study how firms revise the requirements and composition of entry-level software work. The postings are the empirical artifacts through which firms state experience thresholds and skill requirements. At the same time, vacancy data cannot show whether posted changes translate one-for-one into realized hiring or employment outcomes.

Second, the public release of ChatGPT occurred during a broader cooling in the software labor market following the post-pandemic hiring surge. Our design addresses this concern by comparing junior and senior vacancies within software development, by comparing software developers to other computer and mathematical occupations, and by examining a lower-exposure technical occupation as a placebo. These analyses help distinguish a software-specific redefinition of entry-level work from a general decline in junior technical hiring. The design supports an interpretation centered on changing vacancy requirements and relative demand, while leaving open the possibility that broader macroeconomic conditions also shaped the level of software hiring during this period.

# 4 Results

We present the results in four steps. First, we examine whether entry-level software opportunities became less available relative to senior opportunities after ChatGPT's public release. Second, we assess whether this pattern is specific to software development or a broader decline in junior technical hiring. Third, we examine whether firms redefined entry-level software work by raising experience requirements within job titles and changing the skill content of remaining junior vacancies. Finally, we examine whether these changes were strongest in organizational contexts where firms had greater capacity to adopt and reorganize around generative AI.

## 4.1 The Relative Decline in Junior Software Developer Vacancies

We begin by examining if junior software developer vacancies became less available relative to senior vacancies after ChatGPT's public release. This is the first indicator of entry-level work redefinition. Figure 1 plots the monthly coefficients from Equation 1, where each coefficient captures the relative difference in junior versus senior software developer vacancies at event time $k$, normalized to October 2022 ($k = -1$). The figure shows little evidence of a systematic break prior to November 2022. The pre-period coefficients hover around zero without a clear trend, while the post-period coefficients turn negative after ChatGPT's public release and remain negative throughout the following 12-month post-period. This pattern that we observe appears to be more consistent with a discrete shift in relative employer demand than with a slow-moving preexisting divergence between junior and senior software roles.

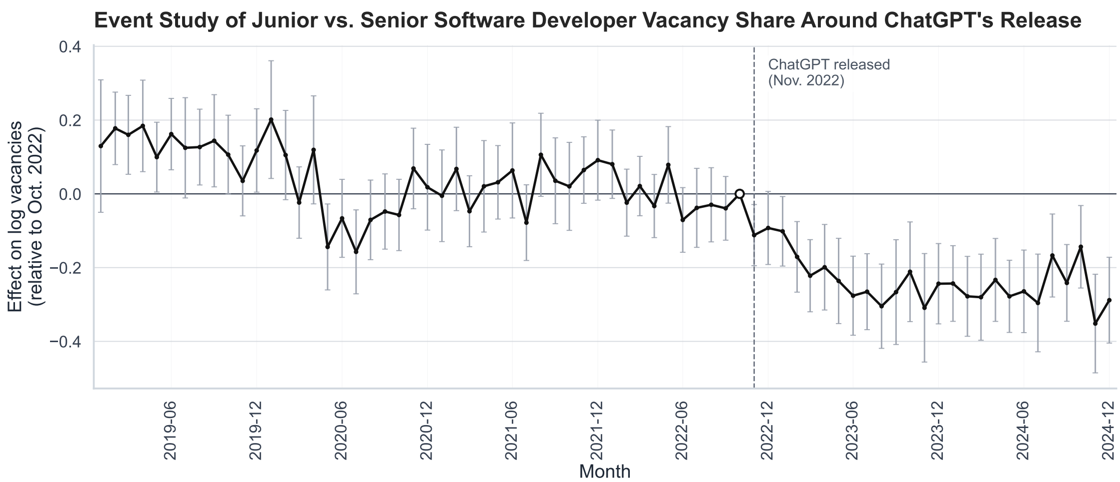


**Figure 1.** Event-study estimates of the change in junior relative to senior software developer vacancies around the public release of ChatGPT. The outcome is the log number of postings in an experience-group × title × CSA × month cell. Coefficients correspond to the interaction between event time and the junior indicator, with October 2022 ($k = -1$) omitted as the reference period. Vertical bars denote 95% confidence intervals based on standard errors clustered at the job-title level.

Table 1 summarizes the average post-ChatGPT effect using the baseline difference-in-differences specification. Across all five specifications, the coefficient on $PostChatGPT \times Junior$ is negative and statistically significant. In the simplest specification, the coefficient estimate is $-0.151$, representing a decline of roughly 14 percent (since the dependent variable is in logs) from the base period. The stability of the estimate across specifications indicates that the main result is not driven by shifts in which specific job titles were posted nor in which CSAs they were located (columns 2 and 3). Moreover, controlling for local labor-market size, firm-size composition, and industry mix in column (5) yields similar estimates ($-0.154$), suggesting that junior software developer vacancies declined by approximately 14 percent relative to senior vacancies after the public release of ChatGPT in November 2022.

The main takeaway is that while both junior and senior hiring cooled after the post-pandemic tech labor market boom, hiring shifted *away from junior roles relative to senior* *roles*. This is central to the paper's contribution. The early labor-market effect of generative AI appears as a within-occupation reallocation and overall decline in demand for software developers.

**Table 1.** Difference-in-differences regression estimates of the impact of ChatGPT's public release on junior relative to senior software developer vacancies. Observations are based on a total of 575,397 underlying individual job vacancies aggregated to 54,456 job title x month x location cells. Across all specifications, the PostChatGPT x Junior interaction is negative and statistically significant, indicating a meaningful decline in junior relative to senior software developer job postings after the widespread introduction of generative AI.

| | *Dependent variable: Log Number of Postings* | | | | |
|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) |
| Junior | -0.760*** | -0.800*** | -0.760*** | -0.760*** | -0.763*** |
| | (0.074) | (0.070) | (0.077) | (0.069) | (0.067) |
| PostChatGPT × Junior | -0.151*** | -0.167*** | -0.151*** | -0.152*** | -0.154*** |
| | (0.027) | (0.029) | (0.028) | (0.028) | (0.028) |
| Month Fixed Effects | Yes | Yes | Yes | Yes | Yes |
| Job Title Fixed Effects | | Yes | | | |
| CSA Fixed Effects | | | Yes | | |
| Population Quintile Controls | | | | Yes | Yes |
| Firm Size Quintile Controls | | | | Yes | Yes |
| Industry Mix Controls | | | | | Yes |
| Observations | 54456 | 54456 | 54456 | 54456 | 54456 |
| Adjusted $R^2$ | 0.168 | 0.305 | 0.192 | 0.170 | 0.178 |
| *Note:* | | | | | *p<0.1; **p<0.05; ***p<0.01 |

*Note: Each coefficient listed is from a separate regression as specified by equation 3 in the text. Each regression uses the first sample listed in Appendix Table A1 that aggregated data from all job postings into Job Title × Month x CSA cells containing at least 5 total postings. Cells that are missing a CSA code or are located in Guam or Puerto Rico are excluded from the analysis. Baseline controls include pre ChatGPT characteristics for each cell including the average firm size across job titles by quintile, the average population size across by quintile, and the industry mix by two-digit NAICS code.*

## 4.2 Software Developers versus Other Mathematical and Computer Science Occupations

The main event study design compares juniors to seniors within the same software development occupation. A remaining concern is that the results might simply detect a broader decline in entry-level computer-science and/or technical hiring during post-COVID normalization of the labor market. To explore this possibility, we perform two robustness checks.

First, we extend our design to a triple-difference comparison between software developers and all other computer and mathematical occupations. Second, we compare the ratio of junior to senior roles within a different "placebo" technical occupation that has characteristics that are similar to software engineering but is less exposed to AI task substitution: mechanical engineering.

Mechanical engineers show many similarities to software engineers (required education, cognitive demands, salaries, corporate budgets). However, mechanical engineers perform more hands-on tasks to design, analyze, test, and manufacture physical devices, engines, and machines. This makes junior roles less substitutable by AI than those in software engineering.

Figure 2 shows graphically the comparison between software developers and all other computer and mathematical occupations, confirming little evidence of a systematic pre-period divergence. Prior to November 2022, the triple-difference coefficients are noisy but do not display a clear break favoring or disadvantaging junior software developers relative to junior workers in the rest of SOC 15. After November 2022, however, the coefficients become negative and remain so. The timing of the shift matches the baseline event study while ruling out the interpretation that the main effect is merely a sector-wide decline in junior technical hiring among computer science and mathematical occupations.

The regression estimates in Table 2 confirm this interpretation. Across specifications, the coefficient on $PostChatGPT \times Junior \times Software$ ranges from $-0.118$ to $-0.140$ and is statistically significant despite the addition of our control variables. These estimates imply that junior software developer vacancies declined by roughly an 11 to 13 percent relative to comparable junior roles in the rest of the computer and mathematical occupations. In other words, the post-ChatGPT shift away from junior labor is not simply a generic feature of the broader computer science labor market. It is disproportionately concentrated in software development, which in theory is the occupation in which large language models most directly overlap with the core tasks and skills of junior workers.

Figure 3 provides a useful placebo check by repeating the event-study analysis within *mechanical engineering* instead of software development. This occupation exists in the same broader macroeconomic environment but is less directly exposed to generative AI, thus we would not expect to see any significant initial decline in junior versus senior postings in the period immediately after the public introduction of ChatGPT. In contrast to software developers, the event study for mechanical engineers shows no clear break in the ratio of junior relative to senior vacancies around November 2022. The pre-period coefficients fluctuate around zero, and the post-period pattern remains noisy.

Together, this strengthens the interpretation that our prior results for software engineers are capturing a labor-demand shift in jobs where generative AI more directly overlaps with core production tasks, not just a broader decline in the demand for junior versus senior technical roles. This supports the view that the post-ChatGPT shift reflects a software-specific adjustment in how firms allocate and define early-career technical work.

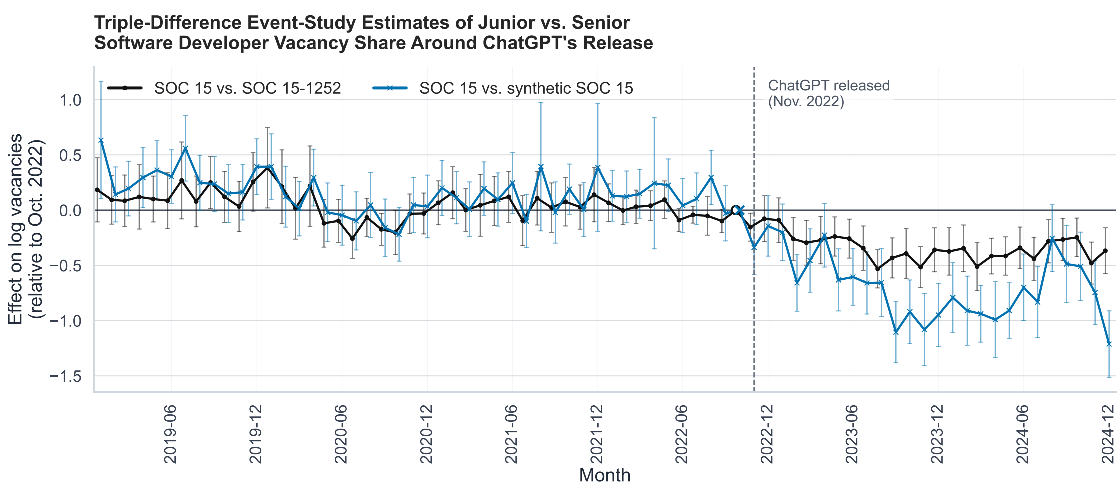


**Figure 2.** Triple-difference event-study estimates comparing junior software developer vacancies to junior vacancies in other computer and mathematical occupations. Coefficients correspond to the event-time interactions for $Junior \times Software$, with October 2022 ($k = -1$) omitted as the reference period. Vertical bars denote 95% confidence intervals based on standard errors clustered at the job-title level.

**Table 2.** Difference-in-difference-in-differences regression for junior software developer postings (SOC 15-1252) versus those in all other computer and mathematical occupations (SOC 15). The negative and significant coefficient on the triple-interaction term indicates a relatively larger decline for junior versus senior software developer vacancies compared to junior versus senior roles in other related occupations *PostChatGPT*.

| | *Dependent variable: Log Number of Postings* | | | | |
|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) |
| Junior | -0.379*** | -0.389*** | -0.376*** | -0.358*** | -0.353*** |
| | (0.081) | (0.084) | (0.083) | (0.080) | (0.080) |
| Software | 0.534*** | 0.418*** | 0.667*** | 0.524*** | 0.516*** |
| | (0.112) | (0.105) | (0.145) | (0.111) | (0.116) |
| PostChatGPT × Junior | -0.155*** | -0.153*** | -0.155*** | -0.159*** | -0.161*** |
| | (0.031) | (0.031) | (0.031) | (0.031) | (0.031) |
| PostChatGPT × Software | -0.090*** | -0.100** | -0.112*** | -0.096*** | -0.098*** |
| | (0.033) | (0.039) | (0.039) | (0.033) | (0.033) |
| PostChatGPT × Junior × Software | -0.136** | -0.140** | -0.138** | -0.122** | -0.118** |
| | (0.058) | (0.059) | (0.059) | (0.056) | (0.056) |
| Month Fixed Effects | Yes | Yes | Yes | Yes | Yes |
| Job Title Fixed Effects | | Yes | | | |
| CSA Fixed Effects | | | Yes | | |
| Population Quintile Controls | | | | Yes | Yes |
| Firm Size Quintile Controls | | | | Yes | Yes |
| Industry Mix Controls | | | | | Yes |
| Observations | 66572 | 66572 | 66572 | 66572 | 66572 |
| Adjusted $R^2$ | 0.106 | 0.186 | 0.168 | 0.110 | 0.114 |
| *Note:* | | | | | *p<0.1; **p<0.05; ***p<0.01 |

*Each regression uses the first sample listed in Appendix Table A1 that aggregated data from all job postings into Job Title × Month x CSA cells containing at least 5 total postings. Cells that are missing a CSA code or are located in Guam or Puerto Rico are excluded from the analysis. Baseline controls include pre ChatGPT characteristics for each cell including the average firm size across job titles by quintile, the average population size across by quintile, and the industry mix by two-digit NAICS code.*

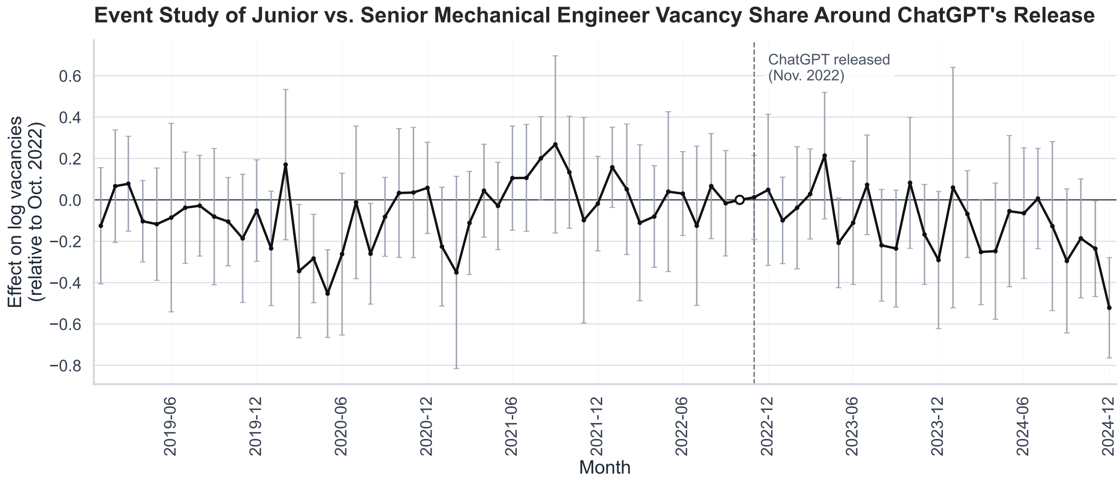


**Figure 3.** For *mechanical engineers*, there is no observable difference before or after generative AI, unlike what we observe for software development. This plot shows DiD event-study estimates of the change in junior relative to senior mechanical engineer vacancies around the public release of ChatGPT. The outcome is the log number of postings in an experience-group × job title × CSA × month cell. Coefficients correspond to the interaction between event time and the junior indicator, with October 2022 ($k = -1$) omitted as the reference period. Vertical bars denote 95% confidence intervals based on standard errors clustered at the job-title level.

## 4.3 Redefining Entry-Level Work Through Experience Requirements

The preceding analyses show that entry-level software vacancies became less available relative to senior vacancies. In this subsection, we examine whether this shift reflects a change in the mix of posted job titles or a change in requirements within those job titles. A decline in junior vacancies could occur because firms post fewer explicitly junior titles, because they shift toward more senior titles, or because they keep posting similar titles while asking for more experienced candidates. The latter possibility is central to our argument because it indicates that firms are raising the bar for similar software roles.

Within the broader software development occupation, we document an overall increase in experience requirements after the public introduction of ChatGPT. In 2022 Q3, the average required experience for a software job title was 5.30 years. Just 12 months later, the average required experience had increased to 5.64 years, a statistically significant shift of 0.34 years or roughly 4 months. Simply looking across the titles listed within the software engineering occupation in Figure 4, our shift-share calculations indicate that the largest contributor to the increase in experience requirements was the *within-job-title effect* (roughly 0.35 years). Only about 0.01 years is attributable to a change in the mix of job titles. This pattern underscores that while hiring managers may still post the same titles, they increasingly prefer candidates with greater experience.

This pattern is especially informative because software job titles vary in how tightly they map onto experience levels. Some titles, such as Software Engineering Intern, are strongly associated with junior roles, while others, such as Lead Software Engineer or Principal Software Engineer, are strongly associated with senior roles. Other titles, such as Software Engineer or Software Developer, are more flexible and can be used for workers at different levels of experience. This matters because in theory employers may have responded to the widespread adoption of generative AI by raising the bar for what requirements are needed for junior-level roles, not simply hiring fewer of them.

Figure 4 provides a closer look *within* the Lightcast Titles Taxonomy groupings to capture changes in experience both within and between the specific job titles listed for that group. Largely junior roles such as Software Engineer exhibit large increases in experience requirements post-ChatGPT due to increasing years of experience requested *within* the underlying specific job titles. In contrast, late-career roles such as Principal Software Engineers show a shift in composition towards including other job titles with less experience, possibly substituting for the more junior Software Engineer roles. This is perhaps because these senior workers are complementary to the adoption of ChatGPT and employers were seeking to fill more of these roles as quickly as possible.

## 4.4 Redefining Entry-Level Work Through Skill Requirements

The remaining junior software vacancies do not appear to have been preserved by increased demand for artificial intelligence or machine learning (AI/ML) skills. If generative AI had primarily created a new technical specialty for entry-level developers, we would expect junior

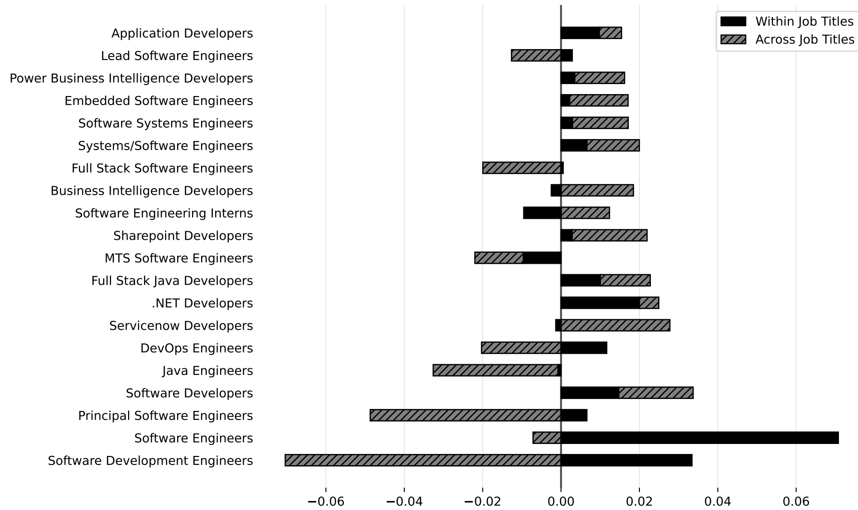


**Figure 4.** Decomposing the change in average years of experience pre- versus post-ChatGPT occurring within versus between job titles. Our shift-share analysis shows that junior roles such as Software Engineer exhibit increasing experience requirements within job titles compared to senior roles such as Principal Software Engineer that show decreasing experience requirements across job titles.

opportunities to shift toward AI-specific skills such as artificial intelligence, machine learning, PyTorch, or related tools. Instead, between October 2022 and October 2024, the monthly share of software developer vacancies asking for AI/ML skills[2] rose only modestly, from 10.5% to 13.1%. Over the same period (amid market contraction), the total number of junior-level vacancies requiring at least one AI/ML skill fell by 86.1%, from 1,915 to 1,029, while senior-level AI/ML vacancies declined by only 10.3%, from 4,338 to 3,934. This pattern suggests that AI-specific skills did not broadly protect entry-level software opportunities, and is consistent with an "escalator" effect in which technological change raises demand for higher-skill workers while leaving less-experienced workers behind [10].

Table 3 provides a more granular view of how skill requirements shifted among junior software vacancies. The skills that became more associated with junior postings after ChatGPT's release were not primarily AI-specific technical skills. Instead, the largest positive deviations from pre-ChatGPT trends were interpersonal communication, problem solving, programming languages, detail orientation, and computer engineering. In contrast, artificial

[2]See Section E of the Appendix for a full list of AI/ML skills as defined by Lightcast.

**Table 3.** This table looks at what skills are associated with junior software developer postings. For example, 20.85% of junior postings asked for Interpersonal Communication in November 2023. We compare this against predicted shares using linear trends from October 2021 to October 2022. A positive deviation from the predicted trend indicates that a skill increased in relative importance for employers. A negative deviation indicates that employers are not as interested in juniors having that skill.

| Rank | Skill | Junior Share Nov 2023 (%) | Predicted Share Nov 2023 (%) | Deviation (pp) |
|---|---|---|---|---|
| 1 | Interpersonal Communication | 20.85 | *0 | 20.85 |
| 2 | Problem Solving | 16.66 | 13.09 | 3.57 |
| 3 | Programming Languages | 18.97 | 16.75 | 2.22 |
| 4 | Detail Oriented | 25.17 | 23.19 | 1.97 |
| 5 | Computer Engineering | 23.66 | 21.91 | 1.75 |
| 96 | Artificial Intelligence | 14.33 | 27.84 | -13.51 |
| 97 | Business Process | 14.93 | 28.59 | -13.66 |
| 98 | .NET Framework | 11.14 | 25.28 | -14.14 |
| 99 | Machine Learning | 16.80 | 33.50 | -16.71 |
| 100 | Operations | 14.09 | 30.98 | -16.89 |

**Note: Interpersonal Communication was declining in usage prior November 2022, such that the linear trend predicted it's usage to be 0% by November 2023.*

intelligence and machine learning appear among the largest negative deviations. This is consistent with our theoretical argument that the remaining junior roles were reorganized around capabilities that complement AI-assisted production.

## 4.5 Organizational Context for Redefining Entry-Level Work

Finally, we examine whether the shift away from junior software vacancies was stronger in organizational contexts where firms had greater capacity to adopt and reorganize around generative AI. The theory does not imply that all firms should adjust hiring standards at the same pace. Instead, firms with more resources, larger technical workforces, and access to deeper labor markets may be better positioned to experiment with generative AI and redesign entry-level work. To estimate heterogeneous impacts separately across industry, firm size, and geography, we restrict our analyses to the period of November 2021 through October 2023, centering on ChatGPT's release, and use the following fully saturated model.

$$Ratio_t = \alpha + \beta PostChatGPT_t + \sum_{n}^{5} (\delta S_n + \theta_n (PostChatGPT_t \times S_n)) + \tau_t, \quad (4)$$

where $Ratio_t$ is the junior-to-senior ratio of job postings in month $t$, $PostChatGPT_t$ is 1 for any month on or after November 2022, $S_n$ represents the specific subgroups within each area (e.g., NAICS industry groups[3], firm size quintiles, population quintiles) and $\tau_t$ are month-specific fixed effects. We find that while some industries are affected more than others, the increase in demand for more experienced software developers is fairly widespread across firms and CSAs of different sizes. We summarize these findings below and refer the reader to the Supplementary Materials for the individual estimates.

The clearest heterogeneity appears by firm size. Among firms in the largest pre-period size quintile, the coefficient on $PostChatGPT \times Junior$ is $-0.229$, compared to $-0.136$ for firms in quintiles 1 through 4, a difference of $-0.093$ (see Table D2). This is consistent with the idea that larger firms were better positioned to experiment with generative AI and to reorganize teams around AI-assisted production. Larger organizations may be especially able to preserve output while reducing the need for junior labor, either by augmenting senior workers or by shifting lower-level coding work onto a smaller number of more experienced developers.

The geographic pattern is directionally similar, though more modest. In the largest CSA population quintile, the estimate is $-0.264$, compared to $-0.199$ elsewhere, a difference of $-0.065$ (see Table D3). This pattern is consistent with faster diffusion of generative AI in larger labor markets, where frontier firms, specialized technical labor, and complementary organizational capabilities are more concentrated.

The industry-level results are more exploratory. Although all industries experienced a general contraction in junior versus senior software developer vacancies starting in mid-2022, some industries were affected more than others (see Table D1 in the Appendix). Perhaps not

[3] We drop industries that do not have at least 100 software developer vacancies per month. This eliminates Agriculture, Forestry, Fishing and Hunting (NAICS 11), Mining, Quarrying, and Oil and Gas Extraction (NAICS 21), and Arts, Entertainment, and Recreation (NAICS 71).

surprisingly, both Information (NAICS 51) and Professional, Scientific and Technical Services (NAICS 54), industries that employ a disproportionately large share of software developers, displayed post-ChatGPT impacts that were twice as large as other sectors. However, Retail Trade (NAICS 44-45) and Accommodation and Food Services (NAICS 72), industries that had rapidly expanded their use of software applications during the pandemic, also showed out-sized impacts. In contrast, the one industry that did not demonstrate any significant impacts during the post-period was Management of Companies and Enterprises (NAICS 55). These are companies that primarily engage in influencing management decisions which may have distinctive business models or risk tolerances that delay the adoption of generative AI on their software hiring. The redefinition of entry-level software work appears to depend on how software tasks are embedded in production and how easily organizations can reallocate those tasks.

Finally, to test the differential impacts by industry we categorize sectors into three groups based on the number of software developer postings per industry:

- **Light** - Management of Companies and Enterprises (NAICS 55), Accommodation and Food Services (NAICS 72), and Public Administration (NAICS 92);
- **Medium** - Construction (NAICS 23), Wholesale Trade (NAICS 42), Transportation and Warehousing (NAICS 48-49), Real Estate and Rental and Leasing (NAICS 53), Educational Services (NAICS 61), Health Care and Social Assistance (NAICS 62), and Other Services (except Public Administration) (NAICS 81).
- **Heavy** - Manufacturing (NAICS 31), Retail Trade (NAICS 44-45), Information (NAICS 51), Finance and Insurance (NAICS 52), Professional, Scientific, and Technical Services (NAICS 54), Administrative, and Support and Waste Management and Remediation Services (NAICS 56).

Surprisingly, we find that the impacts of ChatGPT by industry appear to be nonlinear with the greatest impacts being felt by industries with the fewest and the most software developer postings.

# 5 Discussion

## 5.1 Generative AI and the Redefinition of Entry-Level Software Work

This paper began with a question about how organizations respond when a new information system changes an occupation. Our results point to a clear answer. It changed how firms defined entry-level software work. After ChatGPT's release, junior software developer vacancies became less available relative to senior vacancies, the shift was stronger than in related technical occupations, and employers increasingly asked for more experience within similar job titles.

The most important implication is that generative AI affected the boundary between junior and senior work. Firms did not only post fewer junior vacancies. They also raised experience requirements within job titles and shifted the remaining junior vacancies toward skills that complement AI-assisted work, including problem solving, interpersonal communication, and attention to detail. At the same time, AI-specific skills did not broadly preserve entry-level opportunities. This pattern supports a theory of within-occupation redefinition: generative AI changed which parts of software work remained valuable at the entry level and, in turn, who counted as an adequate junior hire.

Our findings contribute to both information systems research and the broader literature examining technology adoption and changing skill requirements in economics. Specifically, we shift attention from worker-level performance gains to the employer's organizational response and the subsequent impacts on the labor market and talent pipelines. This complements recent information systems evidence that generative AI has stage-specific and

**Table 4.** Differences-in-Differences Estimates of ChatGPT Impacts across Industry, Firm Size, and Geography. Formally testing for heterogeneity across subgroups within each area confirms that junior software developer roles in larger firms and bigger cities showed greater impacts while industries with moderate demand for software developers demonstrated the smallest impacts.

| Dependent variable: Log Number of Postings | *Coefficient/SE on PostChatGPT x Junior* | | |
|---|---|---|---|
| | **Firm Size** | **CSA Population** | **Industry Type** |
| Quintiles 1-4 | -0.136** | -0.199** | |
| *Standard error* | (0.03) | (0.039) | |
| *Number in sub-group* | 6198 | 29081 | |
| Quintile 5 | -0.229* | -0.264* | |
| *Standard error* | (0.05) | (0.069) | |
| *Number in sub-group* | 28247 | 5364 | |
| Difference | -0.093** | -0.065** | |
| *Standard error* | (0.043) | (0.038) | |
| SDE Light | | | -0.258* |
| *Standard error* | | | (0.076) |
| *Number in sub-group* | | | 198 |
| SDE Medium | | | -0.039** |
| *Standard error* | | | (0.048) |
| *Number in sub-group* | | | 1495 |
| SDE Heavy | | | -0.218** |
| *Standard error* | | | (0.042) |
| *Number in sub-group* | | | 32752 |
| Difference (Medium-Light) | | | 0.219* |
| *Standard error* | | | (0.089) |
| Difference (Heavy-Medium) | | | -0.179* |
| *Standard error* | | | (0.08) |
| Difference (Heavy-Light) | | | 0.04* |
| *Standard error* | | | (0.082) |
| Month Fixed Effects | Yes | Yes | Yes |
| Population Quintile Controls | Yes | Yes | Yes |
| Firm Size Quintile Controls | Yes | Yes | Yes |
| Industry Mix Controls | Yes | Yes | Yes |
| Observations | 34445 | 34445 | 34445 |
| Adjusted $R^2$ | 0.032 | 0.029 | 0.032 |

*Note: We run separate regressions for each group listed and report the coefficients on the interactions of the main PostChatGPT effect with a set of dummy variables that fully specify the sample according to the subgroup of interest. Each regression also includes the main subgroup effect as well as the full set of controls for the pre-ChatGPT characteristics listed in Table B1. Each regression uses the first sample listed in Appendix Table A1 that aggregated data from all job postings into Job Title × Month x CSA cells containing at least 5 total postings. Cells that are missing a CSA code or are located in Guam or Puerto Rico are excluded from the analysis. Baseline controls include pre ChatGPT characteristics for each cell including the average firm size across job titles by quintile, the average population size across by quintile, and the industry mix by two-digit NAICS code. Statistical significance is indicated by *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.*

expertise-dependent effects by showing how firms translate such heterogeneity into hiring decisions [29]. Our findings show how firms translate the productivity gains from generative AI into changing labor demands that affect broader entry-level and advancement opportunities. From that perspective, ChatGPT is not only a productivity tool. It is also an information system whose adoption changed how organizations defined expertise and allocated tasks. This framing connects to cognitive reapportionment by showing how generative AI can shift the allocation of work, judgment, and responsibility across workers of different experience levels [35].

## 5.2 From Task Substitution to Rewriting Vacancies

The findings contribute to information systems research by shifting attention from the direct productivity effects of generative AI to the organizational responses that follow from those effects. Much of the early evidence on generative AI examines whether AI tools make individual workers more productive or impacts aggregate employment. Our evidence suggests that productivity effects are only one part of the organizational consequence of generative AI. When a technology changes the relative value of tasks, firms may respond by revising role requirements, reallocating work across experience levels, and changing the hiring standards attached to entry-level positions.

If generative AI lowers the cost of routine coding, then the marginal value of a junior software developer depends on the worker's ability to perform tasks that AI cannot currently automate: diagnosing problems, having human interactions, and designing larger systems. The empirical pattern in this paper closely fits that mechanism. Firms translated this task-level change into vacancy-level changes: fewer junior opportunities, higher experience thresholds within titles, and remaining junior vacancies that emphasized complementary skills. The labor market did not blindly reject junior talent. It narrowed the kinds of early career roles that remain attractive to employers and resembles a lower-friction adjustment where firms limit junior inflows before displacing current employees.

This mechanism also sheds light on larger macro-economic trends. Generative AI can improve the productivity of less-experienced developers, while simultaneously reducing the demand for work done by junior workers. Both facts can be true if the technology increases the output for each junior-equivalent unit of work so that firms need fewer units to produce the same output. In the framework, productivity gains for the marginal junior worker do not imply an expansion of junior hiring in the short run. Although, in the long run, we may see firms increase hiring for juniors as lower unit costs decrease the costs of production allowing for increased product demand.

## 5.3 Implications for Organizations and Talent

The organizational implications center on talent pipelines. Firms appear to value experienced developers in this context because they are better positioned to supervise and integrate AI-assisted work. In the short run, generative AI may increase the relative value of workers who can coordinate across systems and manage ambiguity. Organizations that adopt these AI-assisted systems may save on junior labor while increasing dependence on experienced workers who can build and manage AI workflows. AI adoption can be both a worker-side tool choice and also a firm-level strategic investment that can reflect labor-saving motives [36].

This creates tension. Junior roles are not only low-cost production, they are also how organizations build and train experienced senior talent. If generative AI constricts the bottom of the talent pipeline, firms may benefit from short-run efficiency gains while disrupting the flow of talent. Over time, a labor market with fewer junior entry points could reduce the future supply of experience software developers, especially if fewer students choose to study this topic in fear of a difficult early-career job market. That concern is consistent with recent evidence that entry-level employment declines are strongest in applications where AI automates work rather than augment it [15, 19, 46].

The difficult early-career job market can also hurt new talent. Research confirms that this scarcity of entry-level roles can also lead to "scarring" effects for young workers [33, 42]. Graduating into a weak labor market depresses earnings and stunts professional growth. These effects can persist for decades with talented graduates ending up underemployed or working outside of their field of study, making it harder for them to compete when labor market conditions improve.

For firms, educational institutions, and workforce development programs, the implication is not simply to "teach AI". Our findings examining more granular skill shifts suggest that a better approach would be to help workers develop capabilities that complement AI-assisted work: problem solving, communication, and being detail oriented. Those skills are harder to automate and more likely to remain valuable as routine coding becomes cheaper and generative AI advances. This pattern also aligns with evidence that the jobs remaining after generative AI diffusion tend to be more complex [19]. At the same time, firms that rely on generative AI may need to invest in early talent pipelines if they want to maintain an internal pathway to grow their own senior talent.

## 5.4 Conclusions, Limitations, and Future Research

This paper presents evidence on an early transition period, not a long-run equilibrium effect, of generative AI on software work. The public release of ChatGPT coincided with a distinct shift from junior to senior software developer roles. This finding is consistent with both recent industry observations as well as the vast literature about task-based accounts of technological change. Our findings reinforce prior research suggesting that technological disruptions do not uniformly eliminate employment. They reshape the labor market by altering what is most valuable to employers. The event-study and triple-difference evidence make that interpretation more credible, but they do not imply that generative AI is the only force affecting software labor demand during this period.

A second limitation is that we measure vacancies, not actual hires, such that the employment rate of junior candidates might deviate from the trends implied by our analysis of job vacancies. That is appropriate for the paper's theoretical focus on hiring standards and vacancy composition, but future work should link posting changes to downstream employment outcomes.

Third, some of the observed decline in junior opportunities may be from a short-run adjustment while firms learn where generative AI is most useful within specific workflows. For example, recent field evidence has shown the productivity effects associated with generative AI can be negative for experienced developers in context-heavy scenarios where training the LLM can take more time than simply drawing on human experience [11]. Over a longer horizon, organizations may also redesign roles and new entry-level tasks may emerge.

Fourth, the mechanism documented here may evolve as organizations learn how to use generative AI. The early response may reflect a period in which firms adjusted hiring standards before fully redesigning teams or workflows. Over a longer horizon, lower production costs may expand software demand, new entry-level tasks may emerge, and organizations may develop new apprenticeship models for AI-assisted work. Future research should therefore examine whether the redefinition of entry-level work changes form as generative AI becomes embedded in organizational routines.

These limitations point to promising directions for future research. One is to track the careers of displaced workers. Do they move into adjacent technical occupations or graduate studies? More broadly, future research should also test whether the mechanism documented in this paper generalizes to other AI-exposed occupations like paralegals and sales representatives. If so, the broader consequence of generative AI may be the redefinition of how workers enter the labor market.

Overall, by studying the ways in which AI can alter entry-level opportunities and skill requirements, this paper offers a window into how technology reshapes the workforce. Firms, educational institutions, and public agencies should consider how these findings can guide long-term strategic decisions to prepare the next generation of workers.

## References


[1] Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives*, 33(2):3–30, 2019.

[2] Daron Acemoglu, Gary Anderson, David N Beede, Catherine Buffington, Eric E Childress, Emin Murat Dinlersoz, Lucia S Foster, Nathan Goldschlag, John C Haltiwanger, Zachary Kroff, et al. *Automation and the workforce: A firm-level view from the 2019 annual business survey*, volume 30659. National Bureau of Economic Research, 2022.

[3] Ahmad Alabdulkareem, Morgan R Frank, Lijun Sun, Bedoor AlShebli, César Hidalgo, and Iyad Rahwan. Unpacking the polarization of workplace skills. *Science advances*, 4 (7):eaao6030, 2018.

[4] Liudmila Alekseeva, José Azar, Mireia Giné, Sampsa Samila, and Bledi Taska. The demand for ai skills in the labor market. *Labour economics*, 71:102002, 2021.

[5] Katharine A Anderson. Skill networks and measures of complex human capital. *Proceedings of the National Academy of Sciences*, 114(48):12720–12724, 2017.

[6] Melanie Arntz, Terry Gregory, and Ulrich Zierahn. Revisiting the risk of automation. *Economics Letters*, 159:157–160, 2017.

[7] David Autor, Caroline Chin, Anna Salomons, and Bryan Seegmiller. New frontiers: The origins and content of new work, 1940–2018*. *The Quarterly Journal of Economics*, 139 (3):1399–1465, 03 2024. doi: 10.1093/qje/qjae008. URL `https://doi.org/10.1093/qje/qjae008`.

[8] David H. Autor. Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30, September 2015. doi: 10.1257/jep.29.3.3. URL `https://www.aeaweb.org/articles?id=10.1257/jep.29.3.3`.

[9] David H Autor, Frank Levy, and Richard J Murnane. The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4):1279–1333, 2003.

[10] David H Autor, Lawrence F Katz, and Melissa S Kearney. The polarization of the us labor market. *American economic review*, 96(2):189–194, 2006.

[11] Joel Becker, Nate Rush, Elizabeth Barnes, and David Rein. Measuring the impact of early-2025 ai on experienced open-source developer productivity. *arXiv preprint arXiv:2507.09089*, 2025.

[12] James Bessen. Artificial intelligence and jobs: The role of demand. In *The economics of artificial intelligence: an agenda*, pages 291–307. University of Chicago Press, 2018.

[13] Luís MA Bettencourt, José Lobo, Dirk Helbing, Christian Kühnert, and Geoffrey B West. Growth, innovation, scaling, and the pace of life in cities. *Proceedings of the national academy of sciences*, 104(17):7301–7306, 2007.

[14] Alexander Bick, Adam Blandin, and David J Deming. The rapid adoption of generative AI. *Management Science*, 2026. doi: 10.1287/mnsc.2025.02523.

[15] Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen. Canaries in the coal mine?: Six facts about the recent employment effects of artificial intelligence. 2025. Dated November 13, 2025.

[16] Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. Generative ai at work. *The Quarterly Journal of Economics*, 140(2):889–942, 02 2025. doi: 10.1093/qje/qjae044. URL `https://doi.org/10.1093/qje/qjae044`.

[17] Kevin Zheyuan Cui, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. The effects of generative AI on high-skilled work: Evidence from three field

experiments with software developers. *Management Science*, 2026. doi: 10.1287/mnsc. 2025.00535.

[18] David J Deming. The growing importance of social skills in the labor market. *The quarterly journal of economics*, 132(4):1593–1640, 2017.

[19] Ozge Demirci, Jonas Hannane, and Xinrong Zhu. Who is ai replacing? the impact of generative ai on online freelancing platforms. *Management Science*, 71(10):8097–8108, 2025.

[20] Marcus Dillender and Eliza Forsythe. Computerization of white collar jobs. Accepted at the Journal of Labor Economics, 2025.

[21] Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308, 2024.

[22] Morgan R Frank, Esteban Moro, Tobin South, Alex Rutherford, Alex Pentland, Bledi Taska, and Iyad Rahwan. Network constraints on worker mobility. *Nature Cities*, 1(1): 94–104, 2024.

[23] Garth Frazer and Johannes Van Biesebroeck. Trade growth under the african growth and opportunity act. *The Review of Economics and Statistics*, 92(1):128–144, 2010.

[24] Bronwyn H. Hall and Beethika Khan. Adoption of new technology. NBER Working Paper 9730, National Bureau of Economic Research, 2003.

[25] Kunal Handa, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, et al. Which economic tasks are performed with ai? evidence from millions of claude conversations. *arXiv preprint arXiv:2503.04761*, 2025.

[26] Brad Hershbein and Lisa B Kahn. Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7):1737–1772, 2018.

[27] César A Hidalgo, Bailey Klinger, A-L Barabási, and Ricardo Hausmann. The product space conditions the development of nations. *Science*, 317(5837):482–487, 2007.

[28] Seyed Mahdi Hosseini Maasoum and Guy Lichtinger. Generative ai as seniority-biased technological change: Evidence from us résumé and job posting data. *Available at SSRN 5425555*, 2025.

[29] Jinghui (Jove) Hou, Lei Wang, Gang Wang, Harry Jiannan Wang, and Shuai Yang. The double-edged roles of generative AI in the creative process: Experiments on design work. *Information Systems Research*, 2025. doi: 10.1287/isre.2024.0937.

[30] Krystal Hu and Anna Tong. Chatgpt sets record for fastest-growing user base - analyst note. *Reuters*, February 2023. ChatGPT reached 100 million monthly active users within two months of launch.

[31] Xiang Hui, Oren Reshef, and Luofeng Zhou. The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science*, 35(6):1977–1989, 2024.

[32] An Ju, Hitesh Sajnani, Scot Kelly, and Kim Herzig. A case study of onboarding in software teams: Tasks and strategies. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, pages 613–623. IEEE, 2021.

[33] Lisa B Kahn. The long-term labor market consequences of graduating from college in a bad economy. *Labour economics*, 17(2):303–316, 2010.

[34] Ranim Khojah, Mazen Mohamad, Philipp Leitner, and Francisco Gomes de Oliveira Neto. Beyond code generation: An observational study of chatgpt usage

in software engineering practice. *Proceedings of the ACM on Software Engineering*, 1 (FSE):1819–1840, 2024.

[35] Benn R Konsynski, Abhishek Kathuria, and Prasanna P Karhade. Cognitive reapportionment and the art of letting go: a theoretical framework for the allocation of decision rights. *Journal of Management Information Systems*, 41(2):328–340, 2024.

[36] Boshuo Li, Ni Huang, and Wei Shi. Forced to change? media exposure of labor issues and firm artificial intelligence investment. *Information Systems Research*, 2025. doi: 10.1287/isre.2022.0402.

[37] Kristina McElheran, J Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas. Ai adoption in america: Who, what, and where. *Journal of Economics & Management Strategy*, 33(2):375–415, 2024.

[38] André N Meyer, Earl T Barr, Christian Bird, and Thomas Zimmermann. Today was a good day: The daily life of software developers. *IEEE Transactions on Software Engineering*, 47(5):863–880, 2019.

[39] Alicia Sasser Modestino, Daniel Shoag, and Joshua Ballance. Downskilling: changes in employer skill requirements over the business cycle. *Labour Economics*, 41:333–347, 2016.

[40] Alicia Sasser Modestino, Daniel Shoag, and Joshua Ballance. Upskilling: Do employers demand greater skill when workers are plentiful? *Review of Economics and Statistics*, 102(4):793–805, 2020.

[41] Alicia Sasser Modestino, Mary A. Burke, Sina Sadighi, Rachael Sederberg, Tristan Stern, and Bledi Taska. No longer qualified? changes in the supply and demand for skills within occupations. *ILR Review*, 79(3):486–522, 2026.

[42] Philip Oreopoulos, Till Von Wachter, and Andrew Heisz. The short-and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics*, 4(1):1–29, 2012.

[43] Nicholas Otis, Rowan Clarke, Solène Delecourt, David Holtz, and Rembrand Koning. The uneven impact of generative AI on entrepreneurial performance, 2024. SSRN working paper.

[44] Sida Peng, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. The impact of ai on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590*, 2023. URL `https://arxiv.org/abs/2302.06590`.

[45] Justin R Pierce and Peter K Schott. The surprisingly swift decline of us manufacturing employment. *American economic review*, 106(7):1632–1662, 2016.

[46] Dandan Qiao, Huaxia Rui, and Qian Xiong. Artificial intelligence and jobs: Has the inflection point arrived? evidence from an online labor platform. *Journal of Management Information Systems*, 43(1):5–37, 2026.

[47] Guohou Shan and Liangfei Qiu. Examining the impact of generative AI on users' voluntary knowledge contribution: Evidence from a natural experiment on Stack Overflow. *Information Systems Research*, 2025. doi: 10.1287/isre.2023.0332.

[48] Fabian Stephany and Ole Teutloff. What is the price of a skill? the value of complementarity. *Research Policy*, 53(1):104898, 2024.

[49] Tobias Sytsma. Improving preferential market access through rules of origin: Firm-level evidence from bangladesh. *American Economic Journal: Economic Policy*, 14(1): 440–472, 2022.

[50] Tao Xiao, Christoph Treude, Hideaki Hata, and Kenichi Matsumoto. Devgpt: Studying developer-chatgpt conversations. In *Proceedings of the 21st International Conference on*

*Mining Software Repositories*, page 227–230, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400705878. doi: 10.1145/3643991.3648400. URL https://doi.org/10.1145/3643991.3648400.

[51] Mei Xue, Xing Cao, Xu Feng, Bin Gu, and Yongjie Zhang. Is college education less necessary with ai? evidence from firm-level labor structure changes. *Journal of Management Information Systems*, 39(3):865–905, 2022.

# Appendix

## A Sample Characteristics

**Table A1.** Summary Statistics for Employer Skill Requirements.

| | Oct. 21 | Oct. 22 | Oct. 23 | Difference (*Oct. 22 - Oct. 21*) | Difference (*Oct. 23 - Oct. 22*) | Total |
|---|---|---|---|---|---|---|
| *Sample 1: Cross-sectional sample of all postings aggregated by CSA × software job title cells* | | | | | | |
| Conceptual number of observations | 46000 | 46000 | 46000 | | | 1104000 |
| Actual number of observations | 3177 | 2795 | 1485 | -382 | -1310 | 95953 |
| Total postings per CSA×title | 54969 | 45473 | 18923 | -9496 | -26550 | 1555766 |
| *Mean* | 17.3 | 16.27 | 12.74 | -1.03 | -3.53 | |
| *Std. dev* | 31.57 | 29.94 | 16.61 | | | |
| Percent of job postings requesting: | | | | | | |
| Bachelor's degree or higher | 53.1 | 55.1 | 54.3 | 1.9 | -0.8 | |
| No education listed | 41.3 | 41.8 | 41.9 | 0.5 | 0.1 | |
| Four or more years of experience | 47.8 | 49.6 | 50.9 | 1.8 | 1.3 | |
| Three or fewer years of experience | 20.9 | 20.2 | 14.5 | -0.8 | -5.6 | |
| No experience listed | 31.2 | 30.2 | 34.6 | -1.0 | 4.3 | |
| *Sample 2: Cross-sectional sample of postings with non-missing employer names* | | | | | | |
| Total number of job postings | 49821 | 40052 | 17561 | -9769 | -22491 | 1431527 |
| Percent of job postings requiring: | | | | | | |
| Bachelor's degree or higher | 52.5 | 56.4 | 54.7 | 3.9 | -1.6 | |
| Four or more years of experience | 46.6 | 48.8 | 50.9 | 2.2 | 2.1 | |
| Three or fewer years of experience | 21.2 | 20.8 | 14.7 | -0.4 | -6.1 | |
| *Sample 3: Panel sample of repeated employer × title × CSA observations* | | | | | | |
| Total number of job postings | 28274 | 29221 | 10483 | 947 | -18738 | 868622 |
| Percent of job postings requiring: | | | | | | |
| Bachelor's degree or higher | 53.6 | 58.0 | 57.3 | 4.5 | -0.8 | |
| Four or more years of experience | 45.8 | 48.5 | 50.9 | 2.7 | 2.4 | |
| Three or fewer years of experience | 24.0 | 21.7 | 14.7 | -2.3 | -7.0 | |

***Note****: Employer skill requirements are constructed using online job posting data provided by Lightcast. This sample uses data from all 5.7 million job postings for software developers aggregated into CSA x software job title cells containing at least 5 total postings. The second sample uses data from the subset of job postings that identify employer name and contain at least 5 total postings in each employer × software job title × CSA cell within a given month. The third sample includes only postings for which a given employer × software job title × CSA cell is observed at least once before November 2022 and once on or after November 2022. In all three of the above samples, we exclude postings that are missing a CSA code or are located in Guam or Puerto Rico.*

**Table A2.** Donor occupation composition of the synthetic SOC 15 sample. The table reports the donor SOC weights and the average monthly junior and senior vacancies in the one-year pre-treatment window from November 2021 through October 2022.

| SOC-5 occupation | Weight (%) | Mean junior (pre) | Mean senior (pre) | Junior share (%) |
|---|---|---|---|---|
| 15-1242 | 20.000 | 824.083 | 1690.000 | 32.779 |
| 15-2051 | 20.000 | 1964.750 | 2305.417 | 46.011 |
| 15-1255 | 20.000 | 735.167 | 1320.417 | 35.764 |
| 15-1299 | 20.000 | 1163.417 | 3695.167 | 23.946 |
| 15-1241 | 13.062 | 477.083 | 1211.833 | 28.248 |
| 15-1232 | 6.938 | 1834.250 | 633.417 | 74.331 |

# B Independent Variables

**Table B1.** Summary Statistics for Independent Variables. In this table we provide summary statistics for each of our independent variables used in the regression presented in the text. The population quintile is generated using data from the Census for 2022. The firm size quintile is generated as a firm's total postings for any job title from November 2021 to October 2022. The industry mix is by NAICS code for October 2022.

| | *Quintile 1* | *Quintile 2* | *Quintile 3* | *Quintile 4* | *Quintile 5* |
|---|---|---|---|---|---|
| *Share CSA population (Oct. 2022)* | 19.4 | 20.3 | 24.3 | 16.4 | 19.6 |
| *Share firm size (Oct. 2022)* | 1.3 | 2.7 | 5.0 | 10.9 | 80.2 |

| | |
|---|---|
| *Industry mix as of Oct. 2022* | |
| NAICS 11: Agriculture, Forestry, Fishing, and Hunting | 0.0 |
| NAICS 21: Mining | 0.0 |
| NAICS 22: Utilities | 0.3 |
| NAICS 23: Construction | 0.2 |
| NAICS 31: Manufacturing | 17.4 |
| NAICS 42: Wholesale Trade | 1.0 |
| NAICS 44: Retail Trade | 8.3 |
| NAICS 48: Transportation and Warehousing | 0.3 |
| NAICS 51: Information | 4.3 |
| NAICS 52: Finance and Insurance | 9.5 |
| NAICS 53: Real Estate and Rental and Leasing | 0.1 |
| NAICS 54: Professional, Scientific, and Technical Services | 18.4 |
| NAICS 55: Management of Companies and Enterprises | 0.1 |
| NAICS 56: Admin, Support, Waste Management and Remediation Services- | 38.0 |
| NAICS 61: Education Services | 0.2 |
| NAICS 62: Health Care and Social Assistance | 0.4 |
| NAICS 71: Arts, Entertainment, and Recreation | 0.0 |
| NAICS 72: Accommodation and Food Services | 0.0 |
| NAICS 81: Other Services (except Public Administration) | 0.0 |
| NAICS 92: Public Administration | 0.2 |
| NAICS 99: Undefined | 1.1 |

# C Robustness Checks

Below we provide several robustness checks to the main specification presented in the text. These include (1) understanding the impact of job postings with no listed required experience, (2) using the level of postings, (3) restricting the sample of postings to either those with employer names and/or those that we can observe repeatedly over time, and (4) separate estimates of the impact of ChatGPT on postings with no listed experience, 0-3 years of experience, and 4 or more years of experience.

### Robustness Check: Classification of Unlabeled Experience Levels

One common but little explored issue with using real-time labor market data is the high number of job postings that do not state any experience requirements, which is distinct from postings that explicitly state "no experience required". Roughly 54.7% of all Lightcast job postings in our sample, including 35.7% of software engineer postings, have no stated experience requirements. Prior researchers have either assumed these postings require no experience [26] or left these job postings as unassigned (e.g., missing) year of experience [40] in their analyses. Neither of these approaches is ideal for our context, given the high share of postings that do not have any stated experience requirements and the importance of years of experience in our analysis.

To explore whether job postings with no years of required experience (which is different than zero years) could bias observed differences between junior and senior positions, we built a logistic regression model to classify job postings into "0-3 required years of experience" (junior) or "4+ required years of experience" (senior) categories. We refer to this subset of postings as "unlabeled" and postings with stated experience requirements as "labeled". Logistic regression was chosen specifically for its interpretability. It allows a clear and straightforward understanding of how each feature contributes to the classification outcome.

As inputs, our model takes individual vacancies with listed skills, job title, education, salary, hiring firm, month, year, and geographic indicators. To incorporate each posting's list of required skills, we use one-hot encoding for the 1000 most common skills. This method represents each skill with a unique position in a 1000-entry vector. When a job posting lists a specific skill, the vector value at that skill's position will equal one. All skills not in the job posting will have a value of zero. This common technique creates a machine-interpretable representation of the total skill landscape. The model output provides the estimated years of required experience for the job posting as either "0-3 required years of experience" or "4+ required years of experience".

The logistic regression classifier was trained on a random sample of the postings for SOC 15-1252 (Software Developers) within our period of interest, 2021-2023. Table C1 presents the results, showing precision, recall, and F1-scores by experience category. The overall accuracy on the test set is 0.74. The model has balanced discrimination between junior and senior categories (precision of 0.73 for junior and 0.75 for senior). These results validate that the logistic regression is reasonably accurate.

**Table C1.** This table reports the classification performance on experience levels, confirming that the classifier achieves a moderate level of success in distinguishing between the two experience groups.

| | **Precision** | **Recall** | **F1-Score** | **Support** |
|---|---|---|---|---|
| 0-3 years experience | 0.73 | 0.76 | 0.74 | 69477 |
| 4+ years experience | 0.75 | 0.71 | 0.73 | 69921 |
| **Accuracy** | | | 0.74 | 139398 |
| **Macro Avg** | 0.74 | 0.74 | 0.74 | 139398 |
| **Weighted Avg** | 0.74 | 0.74 | 0.74 | 139398 |

We conducted a sensitivity analysis to quantify any uncertainty introduced by misclassification. The number of monthly unlabeled postings ranges from 12,514 to 41,005. The proportion of estimated senior roles each month ranges from 37.7% to 49.3%. Applying a binomial proportion calculation, we obtained a standard error of 0.0006. The results in a 95% confidence interval of [0.419, 0.421] under the assumption of no misclassification.

Given the classifier's 74% accuracy, we calculated monthly margins of error for our estimates of senior roles (4+ years experience) in the unlabeled set using a binomial proportion method. Across months, these margins of error range from 0.48% to 0.87%, indicating a high degree of confidence in our estimates. These calculations confirm that the uncertainty introduced by misclassification is relatively minor.

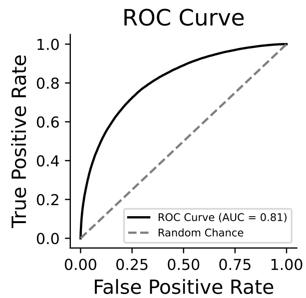


**Figure C1.** The receiver operating characteristic curve for the logistic regression classifier with an AUC=0.81. This indicates the experience level classification has a satisfactory false positive and true positive rate.

To account for potential measurement error, we use a fitted logistic regression model to estimate the experience categories for job vacancies that have no required experience. This is because job vacancies that do not explicitly state experience requirements (e.g., management positions that recruit base on skills-based hiring) may differ in important ways from job vacancies explicitly listing no required experience (e.g., zero years for an entry-level software developer). Figure C2a, shows that our logistical model labels anywhere from 30.3% to 40.6% of software developer postings with no required experience as experienced software developer jobs. Categorizing these previously non-labeled postings we find a clear spike in the estimated senior-to-junior vacancy ratio after the public release of ChatGPT in November 2022 that is similar to what we documented using only the explicitly experience-labeled vacancies. Adding these previously unlabeled vacancies to our prior counts of experience-labeled vacancies does not change the timing nor reduce the magnitude of the hiring shift from junior to senior postings beginning with the widespread introduction of generative AI (Figure C2b).

The broad consistency of these results confirms that the observed contraction in junior software developer roles that coincides with the introduction of generative AI, does indeed

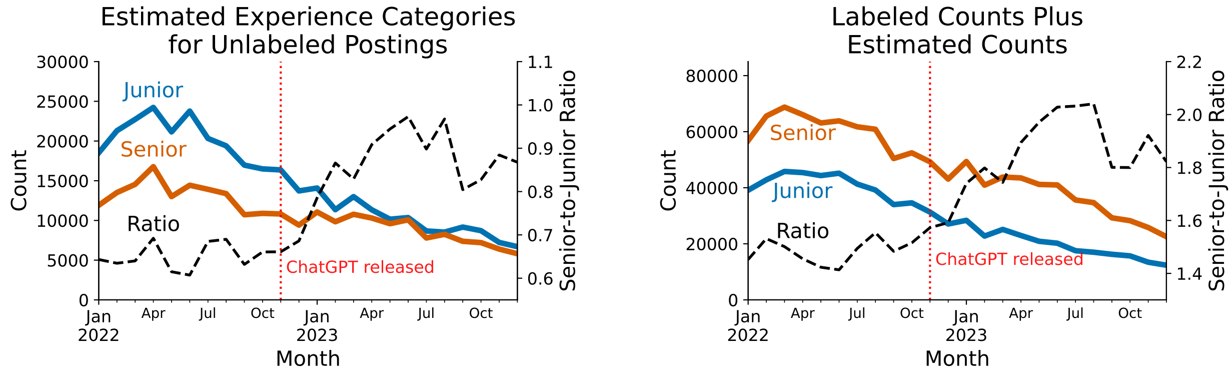


**(a)** **(b)**

**Figure C2.** (a) Counts of previously unlabeled job postings categorized by experience using a logistical model. The senior-to-junior ratio sharply increases immediately after the public release of ChatGPT version 3.5 that is similar to what is observed using only postings that explicitly list years of experience. (b) Adding the estimated counts of previously unlabeled postings to those that explicitly list experience requirements does not change the timing nor the magnitude of the observed shift between senior and junior vacancies shown in the prior figure.

reflect a change in employer demand, rather than a change in how employers write their job vacancies. Employers may have strategically avoided labeling "entry-level" thresholds while still screening for more experienced candidates. By reframing or leaving experience fields blank, firms might seek to recruit candidates with higher skill sets without admitting to a change in job requirements. The observed similarity in the post-ChatGPT shift in the relative demand for senior versus junior software developers, even after imputing experience levels for "unlabeled" vacancies, confirms that the hiring shift does not reflect changes in how employers phrase or omit experience requirements when writing job descriptions.

**Table C2.** Difference-in-Differences Analysis using Posting Levels.

| | *Dependent variable: Number of Postings* | | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| Share Junior | 6.747*** | 5.901*** | 1.414 | 1.414 |
| | (1.474) | (1.434) | (1.148) | (1.148) |
| PostChatGPT × Share Junior | -7.090*** | -6.823*** | -3.751*** | -3.751*** |
| | (1.610) | (1.606) | (1.393) | (1.393) |
| Month Fixed Effects | Yes | Yes | Yes | Yes |
| Job Title Fixed Effects | | Yes | Yes | Yes |
| Firm Fixed Effects | | | Yes | Yes |
| CSA Fixed Effects | | | | Yes |
| Observations | 35179 | 35179 | 35179 | 35179 |
| $R^2$ | 0.004 | 0.024 | 0.180 | 0.180 |
| Adjusted $R^2$ | 0.003 | 0.018 | 0.129 | 0.129 |

*Note: Each coefficient listed is from a separate regression as specified by equation 3 in the text. Each regression uses the first sample listed in Appendix Table A1 that aggregated data from all job postings into Job Title × Month x CSA cells containing at least 5 total postings. Cells that are missing a CSA code or are located in Guam or Puerto Rico are excluded from the analysis. Baseline controls include pre-ChatGPT characteristics for each cell including the average firm size across job titles by quintile, the average population size across by quintile, and the industry mix by two-digit NAICS code. Statistical significance is indicated by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

**Table C3.** Difference-in-Differences Analysis using Restricted Samples.

| | *Share of Postings Requiring 4 or More Years of Experience* | | | |
|---|---|---|---|---|
| **Panel A: Cross-sectional sample of job postings with non-missing employer names (N=589,292)** | | | | |
| | (1) | (2) | (3) | (4) |
| Share Junior | 6.747*** | 5.901*** | 1.414 | 1.414 |
| | (1.474) | (1.434) | (1.148) | (1.148) |
| PostChatGPT × Share Junior | -7.090*** | -6.823*** | -3.751*** | -3.751*** |
| | (1.610) | (1.606) | (1.393) | (1.393) |
| Observations | 35179 | 35179 | 35179 | 35179 |
| Adjusted $R^2$ | 0.003 | 0.018 | 0.129 | 0.129 |
| **Panel B: Panel sample of repeated employer × job title × state observations (N=141,164)** | | | | |
| Share Junior | 7.287*** | 3.611*** | 2.018*** | 2.311*** |
| | (1.472) | (0.800) | (0.672) | (0.679) |
| PostChatGPT × Share Junior | -8.735*** | -4.941*** | -2.878*** | -3.159*** |
| | (1.535) | (0.945) | (0.824) | (0.858) |
| Observations | 14287 | 14287 | 14287 | 14287 |
| Adjusted $R^2$ | 0.010 | 0.071 | 0.088 | 0.100 |
| Month Fixed Effects | Yes | Yes | Yes | Yes |
| Job Title Fixed Effects | | Yes | Yes | Yes |
| Firm Fixed Effects | | | Yes | Yes |
| CSA Fixed Effects | | | | Yes |

*Note: Each coefficient listed is from a separate regression as specified by equation 3 in the text. Each regression uses the first sample listed in Appendix Table A1 that aggregated data from all job postings into Job Title × Month x CSA cells containing at least 5 total postings. Cells that are missing a CSA code or are located in Guam or Puerto Rico are excluded from the analysis. Baseline controls include pre-ChatGPT characteristics for each cell including the average firm size across job titles by quintile, the average population size across by quintile, and the industry mix by two-digit NAICS code. Statistical significance is indicated by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

**Table C4.** Difference-in-Differences Estimates by Skill Level.

| | No Experience | $> 0$ but $\leq 3$ Years | $\geq 4$ Years |
|---|---|---|---|
| | (1) | (2) | (3) |
| PostChatGPT | -0.015*** | -0.026*** | 0.041*** |
| | (0.003) | (0.002) | (0.003) |
| Population Quintile Controls | Yes | Yes | Yes |
| Firm Size Quintile Controls | Yes | Yes | Yes |
| Industry Mix Controls | Yes | Yes | Yes |
| Observations | 47840 | 47840 | 47840 |
| Adjusted $R^2$ | 0.040 | 0.043 | 0.024 |

*Note: Each coefficient listed is from a separate regression as specified by equation 3 in the text. Each regression uses the first sample listed in Appendix Table A1 that aggregated data from all job postings into Job Title × Month x CSA cells containing at least 5 total postings. Cells that are missing a CSA code or are located in Guam or Puerto Rico are excluded from the analysis. Baseline controls include pre-ChatGPT characteristics for each cell including the average firm size across job titles by quintile, the average population size across by quintile, and the industry mix by two-digit NAICS code. Statistical significance is indicated by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

# D Hetergeneous Impacts by Industry, Firm Size, and Geography

In this section we report estimates exploring heterogeneity by industry, firm size, and geography as specified by equation 4 in the text.

**Table D1.** Differences-in-Differences Estimates of ChatGPT Impacts by Industry.

| | *Dependent variable: Ratio of junior to senior vacancies* | |
|---|---|---|
| | Coefficient | Std Error |
| *Coefficient on PostChatGPT x Junior x* | | |
| Construction | -0.705** | (0.319) |
| Manufacturing | -0.779** | (0.319) |
| Wholesale Trade | -0.805** | (0.319) |
| Retail Trade | -1.584*** | (0.319) |
| Transportation and Warehousing | -0.658** | (0.319) |
| Information | -1.534*** | (0.319) |
| Finance and Insurance | -0.934*** | (0.319) |
| Real Estate, Rental, and Leasing | -1.428*** | (0.319) |
| Professional, Scientific, and Technical Services | -1.766*** | (0.319) |
| Management of Companies and Enterprises | -0.187 | (0.319) |
| Admin. Support, Waste Mgmt., and Remediation Services | -0.904*** | (0.319) |
| Educational Services | -1.015*** | (0.319) |
| Health Care and Social Assistance | -0.857*** | (0.319) |
| Accommodation and Food Services | -2.337*** | (0.319) |
| Other Services (except Public Administration) | -1.115*** | (0.319) |
| Public Administration | -0.769** | (0.319) |
| *Controlling for:* | | |
| NAICS FE | Yes | |
| Month FE | Yes | |
| Observations | 1170894 | |
| Groups | 480 | |
| $R^2$ | 0.614 | |
| Adjusted $R^2$ | 0.597 | |

*Note: Each coefficient listed is from a separate regression as specified by equation 4 in the text. Each regression uses the first sample listed in Appendix Table A1 that aggregated data from all job postings into Job Title × Month x CSA cells containing at least 5 total postings. Cells that are missing a CSA code or are located in Guam or Puerto Rico are excluded from the analysis. Baseline controls include pre-ChatGPT characteristics for each cell including the average firm size across job titles by quintile, the average population size across by quintile, and the industry mix by two-digit NAICS code. Statistical significance is indicated by* * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

**Table D2.** Differences-in-Differences Estimates of ChatGPT Impacts by Firm Size.

| | *Dependent variable: Ratio of junior to senior vacancies* | |
|---|---|---|
| | Coefficient | Std Error |
| *Coefficient on PostChatGPT x Junior x* | | |
| Firm Size Quintile Q1 | -0.005 | (0.121) |
| Firm Size Quintile Q2 | -0.015 | (0.121) |
| Firm Size Quintile Q3 | -0.297** | (0.121) |
| Firm Size Quintile Q4 | -0.007 | (0.121) |
| Firm Size Quintile Q5 | -0.480*** | (0.121) |
| *Controlling for:* | | |
| Quintile FE | Yes | |
| Month FE | Yes | |
| Observations | 853538 | |
| Groups | 125 | |
| Adjusted $R^2$ | 0.853 | |

*Note: Each coefficient listed is from a separate regression as specified by equation 4 in the text. Each regression uses the first sample listed in Appendix Table A1 that aggregated data from all job postings into Job Title × Month x CSA cells containing at least 5 total postings. Cells that are missing a CSA code or are located in Guam or Puerto Rico are excluded from the analysis. Baseline controls include pre-ChatGPT characteristics for each cell including the average firm size across job titles by quintile, the average population size across by quintile, and the industry mix by two-digit NAICS code. Statistical significance is indicated by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

**Table D3.** Differences-in-Differences Estimates of ChatGPT Impacts by Geography.

| | *Dependent variable: Ratio of senior to junior vacancies* | |
|---|---|---|
| *Coefficient on PostChatGPT x Junior x* | | |
| | Coefficient | Std. Error |
| Population Quintile 1 | -0.014 | (0.122) |
| Population Quintile 2 | -0.021 | (0.122) |
| Population Quintile 3 | -0.106 | (0.122) |
| Population Quintile 4 | -0.134 | (0.122) |
| Population Quintile 5 | -0.251** | (0.122) |
| *Controlling for:* | | |
| Quintile FE | Yes | |
| Month FE | Yes | |
| Observations | 956596 | |
| Groups | 120 | |
| Adjusted $R^2$ | 0.822 | |

*Note: Each coefficient listed is from a separate regression as specified by equation 4 in the text. Each regression uses the first sample listed in Appendix Table A1 that aggregated data from all job postings into Job Title × Month x CSA cells containing at least 5 total postings. Cells that are missing a CSA code or are located in Guam or Puerto Rico are excluded from the analysis. Baseline controls include pre-ChatGPT characteristics for each cell including the average firm size across job titles by quintile, the average population size across by quintile, and the industry mix by two-digit NAICS code. Statistical significance is indicated by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.*

# E Skill Clusters and Transferability

### Lightcast's AI/ML Skill Clusters

There are a total of 110 skills in Lightcast's Artificial Intelligence and Machine Learning skill cluster that we used in our analysis of skill requirements. These include the following: 3D Reconstruction, Activity Recognition, AdaBoost (Adaptive Boosting), Adversarial Machine Learning, AIOps (Artificial Intelligence For IT Operations), Amazon Textract, Apache Mahout, Apache MXNet, Apache SINGA, Apache Spark, Applications Of Artificial Intelligence, Artificial Intelligence, Artificial Intelligence Development, Artificial Intelligence Markup Language (AIML), Artificial Intelligence Systems, Artificial Neural Networks, Association Rule Learning, Autoencoders, Automated Machine Learning, AWS SageMaker, Azure Cognitive Services, Azure Machine Learning, Baidu, Boosting, Caffe, Caffe2, Chatbot, Classification And Regression Tree (CART), Cognitive Automation, Cognitive Computing, Cognitive Robotics, Collaborative Filtering, Computational Intelligence, Confusion Matrix, Convolutional Neural Networks, Cortana, Cudnn, Dask (Software), Deep Learning, Deep Learning Methods, Deeplearning4j, Dialog Systems, Dlib (C++ Library), Ensemble Methods, Evolutionary Acquisition Of Neural Topologies, Expert Systems, Fast.ai, Feature Engineering, Feature Extraction, Feature Learning, Feature Selection, Game Ai, General-Purpose Computing On Graphics Processing Units, Genetic Algorithm, Gesture Recognition, Google AutoML, Gradient Boosting, H2O.ai, Hidden Markov Model, Inference Engine, Intelligent Agent, Intelligent Control, Intelligent Systems, Intelligent Virtual Assistant, Interactive Kiosk, IPSoft Amelia, Kaldi, Keras (Neural Network Library), Kernel Methods, Knowledge-Based Configuration, Knowledge-Based Systems, Kubeflow, Long Short-Term Memory (LSTM), Machine Learning, Machine Learning Algorithms, Machine Learning Methods, Meta Learning, Microsoft Cognitive Toolkit (CNTK), Microsoft LUIS, MLflow, MLOps (Machine Learning Operations), mlpack (C++ Library), Multi-Agent Systems, Nvidia Jetson, Objective Function, OmniPage, Open Neural Network Exchange (ONNX), OpenAI Gym, OpenCV, OpenVINO, PaddlePaddle, Pydata, PyTorch (Machine Learning Library), Random Forest Algorithm, Reasoning Systems, Recommender Systems, Recurrent Neural Network (RNN), Reinforcement Learning, Scikit-learn (Machine Learning Library), Semi-Supervised Learning, Seq2Seq, Sorting Algorithm, Speech Recognition Software, Speech Synthesis, Supervised Learning, Support Vector Machine, TensorFlow, Test Datasets, Text-To-Speech, Theano (Software), Torch (Machine Learning), Training Datasets, Transfer Learning, Unsupervised Learning, Voice Assistant Technology, Voice Interaction, Voice User Interface, Vowpal Wabbit, Watson Conversation, Watson Studio, Xgboost

### Skill Deviation from Pre-ChatGPT Trends

To observe how required skills change over time, we estimate how the share of junior postings requiring a specific skill deviates from pre-period trends. If a skill appeared to be increasing in importance for employers pre-ChatGPT, then dropped in importance post-ChatGPT, that contributes to our story about the impact of generative AI on entry-level software developers.

We begin with monthly skill-level counts for SOC 15-1252, where each skill-month observation contains the total number of postings mentioning the skill, the number of junior postings mentioning the skill, and the total number of postings mentioning the skill. We restrict this analysis to the top 100 skills by posting volume summed over 2019 through 2024.

For each skill $s$ in month $t$, we compute the junior share

$$J_{st} = \frac{\text{junior postings}_{st}}{\text{total postings}_{st}}.$$

We then estimate a skill-specific linear trend for the junior share using only the pre-ChatGPT training window, defined as October 2021 through October 2022. Specifically, for each skill $s$, we estimate

$$J_{st} = \alpha_s + \beta_s \cdot \tau_t + \varepsilon_{st},$$

Using the estimated coefficients $(\hat{\alpha}_s, \hat{\beta}_s)$, we generate a predicted junior share for a post-period target month, November 2023:

$$\widehat{J}_{s,\text{Nov 2023}} = \hat{\alpha}_s + \hat{\beta}_s \cdot \tau_{\text{Nov 2023}}.$$

Our skill-level measure of post-ChatGPT deviation is the difference between the observed and predicted junior share:

$$\Delta_s^J = J_{s,\text{Nov 2023}} - \widehat{J}_{s,\text{Nov 2023}}.$$

We report this quantity in percentage points, i.e.,

$$100 \times \Delta_s^J.$$

A positive value of $100 \times \Delta_s^J$ indicates that skill $s$ appears in a larger share of junior postings in November 2023 than would be predicted by its own pre-ChatGPT trend, while a negative value indicates that the skill became less junior-associated than predicted. We rank skills by this deviation measure to identify which skills shifted most strongly toward or away from junior postings relative to their prior trajectory. The top five and bottom five deviations are reported in Table 3.

### Skills Analysis for Job Transferability

A significant number of new computer science graduates could be excluded from consideration for, or potentially displaced from, junior software developer roles – at least in the short term. To address this issue, we compare the existing skill sets associated with pre-ChatGPT software developer job postings to those of other vacancies throughout the economy to construct a network of occupations connected by their similarities. This network could serve as a recommender system for recent computer science graduates to target related occupations in their job search and/or make appropriate career transitions if they are laid off.

**Occupation Co-Occurrence Network** Researchers have used skill co-occurrence networks to show the value of diverse skill sets [5, 48], assess the growing cognitive-physical skill polarization [3], and to predict worker mobility [22]. A skill co-occurrence network maintains the information between skills. Nodes are individual skills, and edges are the probability of co-occurrence of skills along job vacancies or job applicants. We adapt the skill co-occurrence network methodology to create an Occupational Co-Occurrence Network.

As the dataset and network density grow, we must prune away edges to maintain the network's usefulness. Following [3], we start with the frequencies of each skill for each occupation. This is the raw number of vacancies for an occupation in Lightcast asking for a specific skill divided by the total number of vacancies for that occupation. Creating edges between every occupation with co-occurring skills would result in a network too dense for useful analysis. Instead, we calculate the revealed comparative advantage of each skill for each occupation. This is similar to the method used in [27] to form a network of products in international trade:

$$rca(\text{occupation, skill}) = \frac{\text{Frequency of a skill in a specific occupation}}{\text{Frequency of a skill across all occupations}}$$

Common skills receive low scores while more distinct skills receive high scores. Next, skills are "effectively used", $e(\text{occupation}, \text{skill}) = 1$, when $rca(\text{occupation}, \text{skill}) > 1$. Otherwise $e(\text{occupation},$ $\text{skills}) = 0$. Each occupation is left with its set of defining skills. Finally, we calculate the conditional probability that two occupations $o$ and $o'$ share skills $s$,

$$\Theta(o, o') = \frac{\sum_s e(s,o)e(s,o')}{\max\left(\sum_s e(s,o), \sum_s e(s,o')\right)}$$

Let occupations be nodes and $\Theta(o, o')$ denote edge weights in our occupational co-occurrence network. We have 734 nodes representing each unique SOC 6 occupation and 263,551 edges with non-zero $\theta$ values. The majority of edge weights are small, with only 8.4% (22,012) exceeding 0.10 and only 0.58% (1,519) exceeding 0.20. This means different occupations share only a few effectively used skills and that only a minority of occupations are similar, suggesting that our model will be able to make useful recommendations for junior software developers seeking to apply their skills in other occupations.

**Results** For software developers, nine of the top ten most similar occupations are within a related subset of Computer Occupations (the SOC 15-1200 series ). While nearly all computer-related roles had some reduction in junior hiring, few experienced as severe of a vacancy contraction as junior software developers. For instance, Data Scientists (SOC 15-2051) showed only moderate declines in the total number of junior vacancies ( $-29\%$) compared to software developers ($-49\%$) alongside smaller decrease in the junior-to-senior ratio ($-11\%$ versus $-33\%$ respectively). Table E1 reports the 34 most similar occupations to software developers using this method. This suggests weaker displacement pressures for junior data scientists. Other high-similarity occupations (e.g., Computer Systems Analysts, SOC 15-1211 and Web and Digital Interface Designers (SOC 15-1255) appear to have maintained a balanced junior-to-senior hiring proportion, indicating they may be more accessible to early-career software developers seeking comparable work.

These findings highlight plausible transitions for junior software developers affected by the generative AI shift. The Occupational Co-Occurrence Network shows that many computer-related roles still require overlapping skill sets (e.g., coding, database management, or technical problem-solving). By measuring the overlap of skill sets between occupations, this network accounts for the transferability of skills across different roles [6]. This means displaced junior software developers could transition into these adjacent roles without extensive retraining, rather than exiting the tech sector altogether.

**Table E1.** Possible transition occupations identified by our Co-Occurrence Network as similar to junior software developers in terms of skill requirements. Θ represents the probability that jobs share a set of effectively used skills. Total employment share, change in the number of junior-level vacancies and the change in the ratio of senior to junior vacancies indicate the feasibility of junior software developers transitioning to a related occupation.

| **Occupation** | **SOC** | **Θ** | **Market Share** | **Vacancy Change** | **Ratio Change** |
|---|---|---|---|---|---|
| *Software Developers* | *15-1252* | *1* | *0.82%* | *-49%* | *33%* |
| Computer Occupations, All O... | 15-1299 | 0.409 | 0.78% | −39% | −26% |
| Data Scientists | 15-2051 | 0.312 | 0.52% | −29% | −11% |
| Database Administrators | 15-1242 | 0.286 | 0.17% | −37% | −11% |
| Computer User Support Speci... | 15-1232 | 0.283 | 0.73% | −38% | −13% |
| Computer Systems Analysts | 15-1211 | 0.261 | 0.28% | −22% | 3% |
| Computer Network Architects | 15-1241 | 0.26 | 0.15% | −34% | −27% |
| Network and Computer System... | 15-1244 | 0.246 | 0.14% | −37% | 2% |
| Web Developers | 15-1254 | 0.236 | 0.11% | −45% | −3% |
| Management Analysts | 13-1111 | 0.232 | 0.39% | −36% | 1% |
| Software Quality Assurance ... | 15-1253 | 0.226 | 0.08% | −53% | −10% |
| Database Architects | 15-1243 | 0.215 | 0.16% | −35% | −20% |
| Industrial Engineers | 17-2112 | 0.21 | 0.34% | −30% | 10% |
| Computer Programmers | 15-1251 | 0.185 | 0.06% | −36% | 24% |
| Marketing Managers | 11-2021 | 0.18 | 0.36% | −32% | −28% |
| Web and Digital Interface D... | 15-1255 | 0.176 | 0.1% | −46% | −2% |
| Operations Research Analyst... | 15-2031 | 0.171 | 0.12% | −40% | −32% |
| Information Security Analys... | 15-1212 | 0.171 | 0.09% | −35% | −6% |
| Mechanical Engineers | 17-2141 | 0.16 | 0.23% | −29% | −17% |
| Market Research Analysts an... | 13-1161 | 0.154 | 0.59% | −30% | −3% |
| Project Management Speciali... | 13-1082 | 0.152 | 0.75% | −10% | −19% |
| Financial and Investment An... | 13-2051 | 0.145 | 0.57% | −24% | −9% |
| Electrical Engineers | 17-2071 | 0.143 | 0.17% | −34% | −33% |
| Architectural and Engineeri... | 11-9041 | 0.141 | 0.07% | −30% | −6% |
| Financial Risk Specialists | 13-2054 | 0.128 | 0.12% | −15% | −4% |
| Business Operations Special... | 13-1199 | 0.128 | 0.07% | −26% | 34% |
| Electronics Engineers, Exce... | 17-2072 | 0.124 | 0.07% | −15% | 9% |
| Logisticians | 13-1081 | 0.124 | 0.4% | −32% | −15% |
| Buyers and Purchasing Agent... | 13-1028 | 0.121 | 0.46% | −19% | −3% |
| First-Line Supervisors of P... | 51-1011 | 0.12 | 0.47% | −25% | −9% |
| Detectives and Criminal Inv... | 33-3021 | 0.116 | 0.09% | 2% | −7% |
| Managers, All Other | 11-9199 | 0.116 | 0.47% | −17% | −17% |
| Commercial and Industrial D... | 27-1021 | 0.115 | 0.06% | −31% | −12% |
| Production, Planning, and E... | 43-5061 | 0.114 | 0.52% | −20% | −23% |
| Computer Hardware Engineers | 17-2061 | 0.113 | 0.02% | −44% | −47% |

*Note: The change in job postings and the junior-to-senior ratio is calculated during the period from September 2022 (pre-ChatGPT) through January 2025*