THE IMPACT OF GENERATIVE AI ON JOB OPPORTUNITIES FOR JUNIOR SOFTWARE DEVELOPERS

Samuel Westby
Alicia Modestino
Peiran Cheng

Working Paper: March 2026

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.

© 2026 by Samuel Westby, Alicia Modestino, Peiran Cheng. All rights reserved.

The Impact of Generative AI on Job Opportunities for Junior Software Developers
Samuel Westby, Alicia Modestino, Peiran Cheng
March 2026

ABSTRACT

Generative artificial intelligence may not only reshape labor demand across occupations, but also within them. We study this possibility in software developer, where large language models overlap closely with many core tasks yet may affect early-career and later-career workers differently. Using the near-universe of U.S. online job vacancies from Lightcast, we study how the public release of ChatGPT in November 2022 changed employer demand for junior versus senior software developers. We estimate event-study and difference-in-differences models comparing software developer vacancies requiring three or fewer years of experience to those requiring more experience, and extend the design to a triple-difference comparison with other computer and mathematical occupations. We find that rapid dissemination of generative AI coincided with a 14 to 15 percent relative drop in junior-level versus senior-level software developer job vacancies, with no comparable pre-trend and a larger relative decline than in related technical occupations. To examine the mechanism, we use a shift-share decomposition. The rising experience demands are driven almost entirely by employers asking for more years of experience for the same position, not a reduction in hiring for those positions. The skill content of the remaining junior vacancies shifts toward problem solving, interpersonal communication, and detail oriented rather than toward AI-specific requirements. Larger firms and larger labor markets show the largest impacts. Overall, the results suggest that firms responded to generative AI not only by reducing early-career opportunities, but also by redefining what counts as an adequate junior hire. This paper shows how the spread of a general-purpose technology changes organizational demand for human capital through within-occupation reallocation and the rewriting of job vacancies.

1 Introduction

The broad economic consequences of generative artificial intelligence (AI) remain theoretically and empirically unsettled. A central question is not simply whether generative AI increases productivity, but how those productivity gains are distributed across workers and how firms translate them into hiring decisions. Emerging evidence suggests that generative AI can produce heterogeneous productivity gains across workers within the same occupation (Cui et al. 2025, Peng et al. 2023, Brynjolfsson et al. 2025b, Hui et al. 2024, Otis et al. 2023). Recent research also suggests that generative AI's performance effects depend on expertise and the stage of work being performed (Hou et al. 2025). Yet it remains unclear whether these productivity gains expand opportunities or reduce labor market opportunities, and for which groups of workers. 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 (Demirci et al. 2025).

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 (Eloundou et al. 2024, Demirci et al. 2025, Khojah et al. 2024, Handa et al. 2025). 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. At the same time, 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 whether a general-purpose digital technology reshapes labor demand *within* an occupation rather than only across occupations.

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 to provide an especially useful perspective on the organizational responses of firms. This is because job vacancy postings capture how employers are shifting hiring standards in real-time in anticipation of changing workflows. Prior research shows that posted job requirements are responsive to labor market conditions rather than fixed features of jobs (Hershbein and Kahn 2018, Modestino et al. 2016, 2020). If generative AI disproportionately affects the tasks of junior software developers, one might expect to see a decline in both the relative number of postings for these workers as well as changes in the experience requirements and/or skill contents of the junior vacancies that remain.

In this paper, we study the public release of ChatGPT in November 2022 as a natural experiment to examine how generative AI reshaped employer demand within software development. Using a comprehensive dataset of U.S. online job vacancies from Lightcast, we compare junior software developer postings, defined as those requiring three or fewer years of experience, to more senior postings 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, 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 related occupations. In addition, there was no impact on junior workers in other STEM occupations such as mechanical engineering.

Second, we show that this adjustment operated through a redefinition of junior roles, not merely a reduction. A shift-share decomposition indicates that a post-ChatGPT increase in average required experience was driven almost entirely by rising experience requirements *within* job titles rather than a shift toward more senior job titles. Put differently, employers did not post fewer junior-style jobs after the popularization of large language models. Instead, they sought more experienced candidates to fill the same role. We also find that the junior vacancies that remained shifted toward complementary higher-order skills, with relative increases in demand for problem solving, interpersonal communication, and attention to detail. Notably, the evidence does not suggest that AI-specific skills made junior workers broadly more employable. Instead, generative AI appears to have increased the value of skills that complement AI-assisted work rather than the value of junior labor. A similar pattern appears in online freelancing markets, where the jobs that remain after generative AI diffusion tend to be more complex rather than simply fewer in number (Demirci et al. 2025).

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, and larger labor markets that are more competitive, are more able to adopt and reorganize around new technologies. In addition, 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 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 in three distinct areas. First it contributes to the emerging literature on generative AI in the labor market by connecting worker-level productivity effects to employer-side labor demand. Existing studies show that generative AI can improve performance, often with large gains for less-experienced workers. We show that these gains coincide with a reduction in employer demand for junior workers. Second, the paper contributes to research on technological change and skill-biased labor demand by showing that change can occur *within* an occupation through shifting vacancy requirements, in addition to aggregate employment shifts. This extends foundational accounts of task-based technological change into a new setting where firms redefine work in response to a new technology. Third, the paper shows how the diffusion of generative AI reshapes organizational demand for human capital.

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. This can occur even without reducing the number of entry-level job openings if firms increase the skill requirements for entry-level positions. 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 (Modestino et al. 2026a).

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 generative AI raises a broad debate whether it will reduce employment, we focus on a more specific question: How does a general-purpose coding technology shift labor demand across workers with different skill sets within software development? Within-occupation upskilling is particularly salient in this case because junior and senior software developers vary significantly in the daily tasks that they perform and their exposure to generative AI as a substitute. As a result, firms may respond to adopting this new technology by changing which types of workers they seek (e.g., junior versus senior) and what they require from them (e.g., specific skills), rather than shrinking the total number of workers employed within the occupation. We argue that the public release of ChatGPT created a task shock that overlapped more with junior than senior software work, creating a cheaper substitute for tasks typically done by junior software engineers. In contrast, the introduction of generative AI likely increased the relative value of senior software engineers whose greater years of experience and higher-order skill sets are more complementary to directing, overseeing, and verifying AI-generated output. In the short run, this should reduce employer demand for junior developers relative to senior developers. That adjustment should appear as relatively fewer entry- and junior-level vacancies, and also through an increasing shift in experience requirements and job tasks for those vacancies.

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 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 (Eloundou et al. 2024, Demirci et al. 2025, Khojah et al. 2024). However, in the early years following the dissemination of large language models, software developers did not initially use these systems as fully autonomous substitutes. Rather, developers often used them for

examples, direction, and preliminary solution paths, while continuing to evaluate outputs for correctness and trustworthiness (Xiao et al. 2024). This combination of high task overlap and semi-automation set 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 sooner rather than later (Bick et al. 2026). For employers, these adjustments may initially take the form of changes in hiring practices rather than layoffs or other formal personnel reorganizations. Changing posted job requirements for new hires is likely faster, 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 theoretically meaningful early signal of how firms respond to a new general-purpose technology (Hershbein and Kahn 2018, Modestino et al. 2020, Dillender and Forsythe 2025).

2.2 Uneven Productivity Gains and Within-Occupation Substitution

The underlying identifying assumption in our approach is that the introduction of generative AI does not equally raise the productivity of all software developers–at least not initially. 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. (Ju et al. 2021). Recent evidence suggests that AI assistance provided a sizable productivity gain for software work and these gains were often large for less-experienced workers (Peng et al. 2023, Brynjolfsson et al. 2025b, Cui et al. 2025). In task-based terms, generative AI should therefore substitute more strongly for junior task bundles than for senior ones (Acemoglu and Restrepo 2019). This interpretation is consistent with online freelancing evidence showing larger post-ChatGPT declines in automation-prone writing and coding jobs than in more manual or intensive work (Demirci et al. 2025).

In contrast, senior developers should remain 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. Recent field evidence even suggests that AI tools may even slow experienced developers in realistic open-source settings, showing that AI's productivity effects are not solely positive for software development (Becker et al. 2025). 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 (Meyer et al. 2019). However, this does not imply that senior developers are untouched by generative AI nor are they immune to future iterations of this technology. Rather, we argue that during the initial adoption period generative AI should raise the relative value of workers who can supervise, evaluate, and integrate AI-assisted work compared to those who are assigned more routine entry-level tasks.

This paper studies the immediate transition period of the rapid diffusion of generative AI in the years following the sudden public release of ChatGPT. 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 (Autor et al. 2024, Autor 2015, Bessen 2018). But during the 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 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 rather than through immediate 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 (Modestino et al. 2026b). During weak labor markets, firms raise requirements within occupations, and those increases may appear directly in vacancy postings even before they affect employment totals (Hershbein and Kahn 2018, Modestino et al. 2020). When labor markets tighten, posted requirements can fall again, which suggests that hiring standards are adjustable rather than fixed properties of the job itself (Modestino et al. 2016).

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 judgment, independent problem solving, and the ability to work collaboratively. 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.

Our hypotheses are consistent with task-based theories of technological change, which argue that new technologies typically substitute for routine tasks while complementing non-routine skills (Autor et al. 2003). More recently, research has also shown that the labor market also rewards roles that combine both quantitative (hard) and social (soft) skills, as higher-order capabilities may become more valuable when routine work is offloaded (Deming 2017). This prior literature sets the stage for two testable hypotheses with regards to the impact of generative AI on software developers: First, the relative decline in junior opportunities post-ChatGPT is likely to manifest in higher years of experience requirements *within existing job titles*, not just through a reduction in the number of postings. Second, the remaining junior opportunities should also place relatively less emphasis on routine coding skills and more emphasis on higher-order cognitive and interpersonal capabilities.

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 vary in their ability to experiment with new tools, reorganize work, and adjust hiring in response to new technologies. 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 move first 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 (Acemoglu et al. 2022, Hall and Khan 2003, McElheran et al. 2024). Recent evidence also suggests that firms may treat AI investment as a strategic response to labor-related pressures in addition to a productivity upgrade (Li et al. 2025). 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 matter for similar reasons. Larger labor markets can facilitate the spread of new technologies and work practices (Bettencourt et al. 2007). 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 rather than as a primary mechanism.

Finally, in some industries, particularly those that engage in software development, generative AI may act as a direct substitute for routine junior developer tasks, making within-occupation substitution more likely. But in industries with relatively fewer software developers, generative AI may instead expand the ability of smaller technical teams and increase demand for junior specialists. The broader implication is that the effect of generative AI on a particular industry is more nuanced and likely depends on its exposure to software work, and also how easily organizations can reorganize around it.

2.5 Summary of Expectations

Taken together, the literature implies that the public release of ChatGPT created a task shock in software development that should affect junior and senior workers differently. Because junior developers are more likely to perform narrower and more explicit tasks, generative AI should raise their productivity while also reducing firms' needs for their labor in the short term. Firms should first respond with changes in hiring by reducing demand for entry- and junior-level developers relative to senior-level developers. They may also increase years of experience and skill requirements of the remaining junior vacancies. These adjustments should be more prevalent where technology adoption and work reorganization is easiest, particularly in larger firms, larger labor markets, and software-heavy industries where substituting for routine software tasks using generative AI is least costly.

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) (Modestino et al. 2016, Hershbein and Kahn 2018, Modestino et al. 2020, Alekseeva et al. 2021).

These data are well suited to our research question for two reasons. First, we seek to understand how firms adjust labor demand and skill requirements in response to this new general-purpose technology, both of which can be observed in real-time using vacancy postings long before they can be detected using publily available data on occupational employment. Second, software development is a white-collar, highly digital occupation that is well represented in online vacancy data. This reduces any bias arising from under-representation in the online job posting data relative to occupations like welding which are less likely to advertise through online job boards.

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 a longer window for descriptive context. Our main empirical design uses a narrow window covering the 12 months before and 12 months after the public release of ChatGPT in November 2022.

3.2 Sample Characteristics and Key Measures

We restrict the main analysis to software developer postings with 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).

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 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 software hiring such as population quintile rank, firm-size, and

industry mix that are measured using pre-ChatGPT data so that they are independent of the generative AI technological shock.

The choice of November 2022 as the inflection point captures the public release of ChatGPT rather the first appearance of AI-assisted coding tools. 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 (Hu and Tong 2023) with wide adoption in computer and mathematical occupations (Bick et al. 2026), making it ideal for studying changes in employer hiring behavior as a discrete and measurable event.

3.3 Research Design

For our empirical strategy, we first estimate an event-study specification to assess pre-trends and trace the monthly evolution of junior relative to senior software vacancies around he November 2022 ChatGPT release date. We initially estimate the average post-period effect using a difference-in-differences specification between junior versus senior software engineers as well as a triple-difference relative to all computer science occupations. In the appendix, we also present robustness checks using a synthetic control for our triple-difference, using a longitudinal dataset of job-titles by firm, and a placebo occupation (mechanical engineers). We then use a shift-share decomposition to distinguish between two mechanisms behind rising experience requirements: changes in the composition of job titles versus rising experience demands within titles. We conclude with some descriptive evidence regaridng which skills are changing for junior versus senior software engineer vacancies.

3.3.1 Event Study

Our event study tests our key assumption that the relative gap between junior and senior software vacancies would not have shifted over time without the public release of ChatGPT. This implies that the pre-event coefficients should be close to zero and insignificant showing a steady relative ratio between junior and senior postings before November 2022, and a significant downward trend thereafter.

Specifically, we estimate:

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

where $Vacancies_{cjm}$ denotes the number of 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 β_k are interpreted relative to the month immediately preceding November 2022. Further, μ_m and γ_c denote month and CSA fixed effects, respectively, X_{cjm} is a vector of controls, and ε_{cjm} is a stochastic error term and $\alpha$is the regression constant.

The coefficients β_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. Specificaly, we examine the trend in 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 β_k trace the dynamic path of junior software vacancies relative to November 2022.[1]

3.3.2 Difference-in-Differences Specification

Specifically, we estimate the post-ChatGPT impact on junior software engineers relative to other computer science occupations using the following difference-in-differences specification:

[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 Pierce and Schott (2016), Frazer and Van Biesebroeck (2010), Sytsma (2022).

$$\ln Vacancies_{cjm} = \alpha + \delta\, Junior_{cjm} + \theta\Big(PostChatGPT_m \times Junior_{cjm}\Big) + \mu_m + \gamma_c + X_{cjm} + \varepsilon_{cjm}. \quad (3)$$

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, and θ. This is the coefficient of interest that captures the average change in junior relative to senior software vacancies after November 2022. A negative θ value indicates that junior postings declined relative to senior postings 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 A..1, and summary statistics for the control variables are reported in Appendix Table B.1. Standard errors are clustered at the job title level to allow for within-category correlation in outcomes over time.

3.3.3 Synthetic Control

Although comparing the impact of junior relative to senior roles for software engineers (SOC 15-1252) to the broader SOC 15 category is informative, the broader category contains six digit occupations of different sizes that have different junior-to-senior ratios and different pre-treatment trends. Rather than weighting all six-digit occupations within SOC-15 using their relative size representation in the dataset, we also construct a synthetic control from the raw data to better match the pre-trend that is observed for the ratio of junior relative to senior software developer postings. Specifically, the synthetic control re-weights the other SOC 15 six-digit occupations so that the aggregated SOC 15 category matches the pre-trend data of SOC 15-1252. This approach produces more precise estimates of the differential impact of generative AI on junior software engineers relative to other SOC 15 six-digit occupations.

We construct the synthetic SOC 15 comparison group using the a 12 month pre-treatment window from November 2021 through October 2022. Weights must be nonnegative and contribute no more than 20 percent of the variation. We find the combination of weights that best match 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 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 A..2

3.3.4 Shift-Share Analysis of Experience Requirements

The regressions above determine whether employer demand shifted away from junior software developers. They do not reveal the mechanism behind this shift, which could be a reduction in roles, or an increase in employer hiring standards. To distinguish between these options, we conduct a shift-share decomposition of average required experience between job titles within software development.

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:

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

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}.

3.4 Interpretation

Two caveats are worth noting. First, online vacancies measure employer demand rather than actual hires. This is appropriate for our analyses because the paper studies how firms update hiring standards and vacancy composition. However, Brynjolfsson et al. (2025a) find evidence in data from a large U.S. payroll provider for an employment decline in early-career AI-exposed workers.

Second, the public release of ChatGPT occurred during a broader cooling in the software labor market following a strong post-pandemic hiring surge. For that reason, the design is intentionally relative between early-career and later-career opportunities.

In addition to the designs above, later analyses compare software developers to other computer, mathematical, and engineering occupations and examine heterogeneity across industries, firms, locations, and skill requirements. These analyses help assess whether the results reflect a software-specific labor-demand adjustment rather than a general decline in entry-level technical hiring.

4 Results

4.1 The Relative Decline in Junior Software Developer Vacancies

We begin with the event study evidence. 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 through the following year. The pattern is more consistent with a discrete shift in relative employer demand than with a slow-moving preexisting divergence between junior and senior software roles.

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, and in the fully controlled specification it is -0.154. Taken together, the

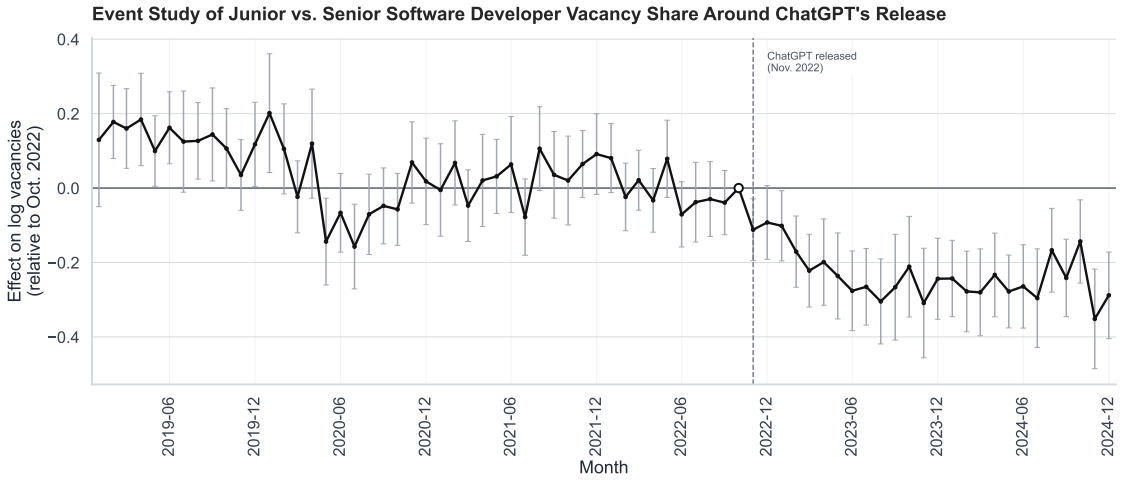


Figure 1. Event-study estimates of the change in junior relative to senior software developer vacancies around the public release of ChatGPT. The outcome is the log number of postings in an experience-group × title × CSA × month cell. Coefficients correspond to the interaction between event time and 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.

estimates imply a decline of roughly 14 to 16 percent in junior software developer vacancies relative to senior vacancies after November 2022. The stability of the estimate across specifications indicates that the main result is not driven by shifts in local labor-market size, firm-size composition, or industry mix.

The main result is not that software hiring collapsed uniquely after ChatGPT. Senior software hiring also softened during the broader post-pandemic cooling of the tech labor market. Rather, the result is that hiring shifted *away from junior roles relative to senior roles*. This distinction is central to the paper's contribution. The early labor-market effect of generative AI appears as a within-occupation reallocation in demand rather than only as an aggregate contraction in software employment.

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 technical hiring during post-COVID normalization of the labor market. First, Figure 2 extends the dynamic design to a triple-difference comparison between software developers and other computer and mathematical occupations. Second, we look at mechanical engineers. This occupation has many similarities to software engineers (required

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

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 A..1 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.

education, cognitive demands, salaries, corporate budgets), but is much less exposed to AI task substitution. Third, Table 2 summarizes the average post-period effect.

Figure 2 again shows 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.

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 throughout. These estimates imply that junior software developer vacancies declined by an additional roughly 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 technical labor market. It is disproportionately concentrated in software development, which in theory is the occupation in which large language models most directly overlap with core tasks and skills.

Figure 3 provides a useful benchmark 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. In contrast to software developers, the figure shows no clear break in junior relative to senior mechanical engineer vacancies around November 2022. The pre-period coefficients fluctuate around zero, and the post-period pattern remains noisy. This strengthens the interpretation that the main results are showing labor-demand shift concentrated in occupations where generative AI more directly overlaps with core production tasks, not just a broader decline in junior-level demand.

4.3 Shift Share Analysis of Junior Software Engineering Roles

The preceding analyses show that relative employer demand shifted away from junior software developers. We next examine whether that shift reflects a change in the mix of posted job titles or a change in requirements within titles. (Job titles within software development can include roles like software engineer, senior software engineer, staff software engineer, etc.) This matters because the theoretical argument is not that employers stopped posting junior roles. It is that generative AI may have raised the bar for what counts as a junior-level hire,

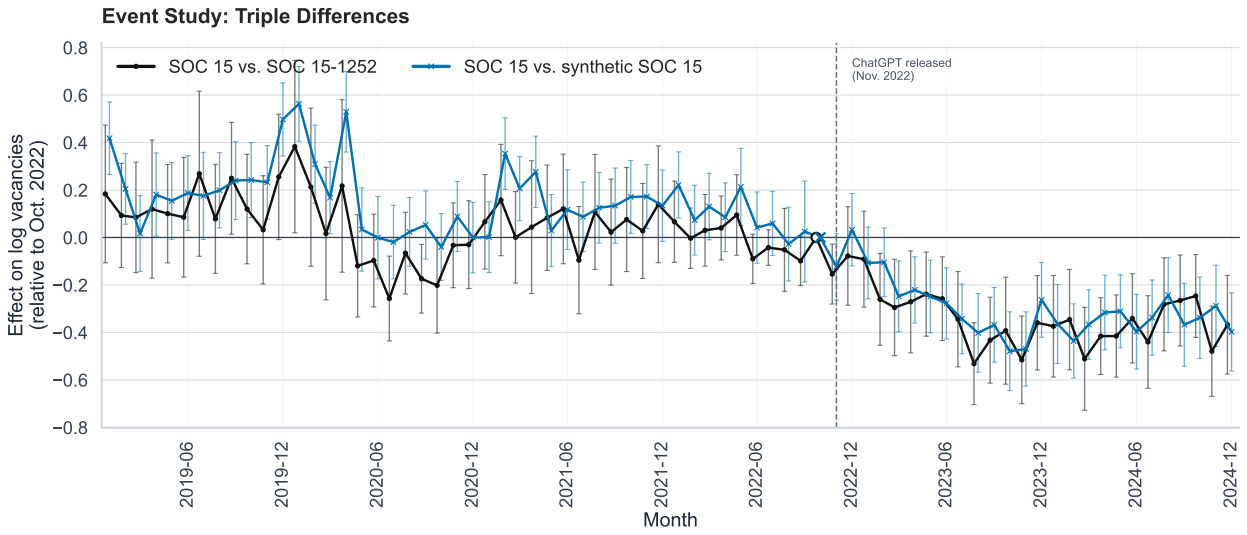


Figure 2. Triple-difference event-study estimates comparing junior software developer vacancies to junior vacancies in other computer and mathematical occupations. Coefficients correspond to the event-time interactions for $Junior \times Software$, with October 2022 ($k = -1$) omitted as the reference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

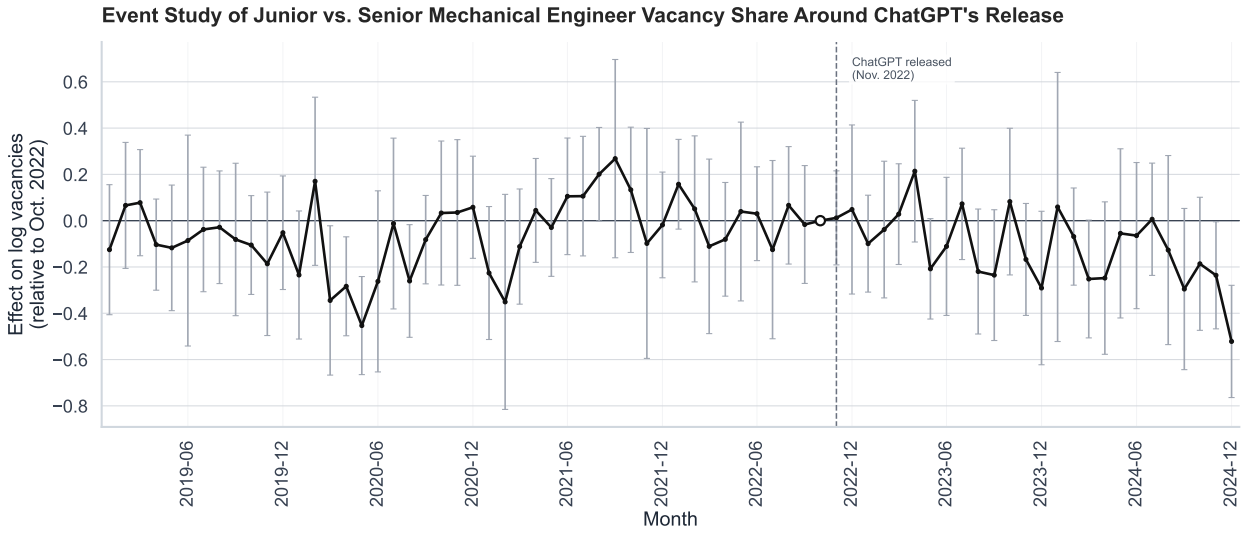


Figure 3. For *mechanical engineers*, there is no observable difference before or after generative AI, unlike what we observe for software development. This plot shows DiD event-study estimates of the change in junior relative to senior mechanical engineer vacancies around the public release of ChatGPT. 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.

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 coefficient listed is from a separate regression as specified by equation **??** *in the text. Each regression uses the first sample listed in Appendix Table A..1 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.*

The shift-share results strongly support that interpretation. We find an overall increase in experience requirements across job titles. In Q3 2022, the average required experience for a job title was 5.30 years. By Q3 2023, it had increased to 5.64 years, a shift of 0.34 years. Our shift-share calculations indicate that the largest contributor (0.35 years) is the within-job-title effect, which reflects more employers increasing experience requirements within each job title. Only 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.

Figure 4 provides a closer look within the Lightcast Titles Taxonomy groupings to capture changes in experience both within and between specific job titles, revealing a more nuanced transformation within the software developer labor market. 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, more senior roles such as Principal Software Engineers actually show decreasing experience requirements, largely *across* specific job titles, perhaps because these workers are complementary to the adoption of ChatGPT. Finally, titles that have a mix of junior and senior roles, such as Software Development Engineers, manifest both phenomena simultaneously occurring with increasing experience requirements within some job titles and decreasing experience requirements across others.

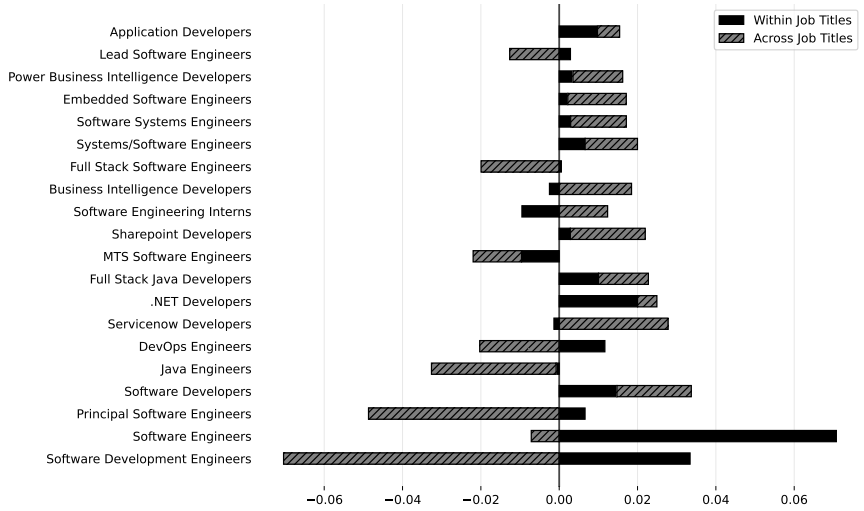


Figure 4. Decomposing the change in average years of experience pre- versus post-ChatGPT occurring within versus between job titles. Our shift-share analysis shows that junior roles such as Software Engineer exhibit increasing experience requirements within job titles compared to senior roles such as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

Isolating the top 100 skills for junior software developers, very few became more associated with juniors than pre-trends would predict. Those that increased from what their pre-generative AI trend would predict are mostly general skills like Interpersonal Communication, Problem Solving, and Detail Oriented. The increase in employer demand for these particular skillsets aligns with the theory that routine coding tasks are being replaced by ChatGPT. Firms may require the new era junior workers to do more higher-order thinking (e.g., problem-solving) as well as excel in team setting with senior software developers (e.g., interpersonal communication). One should interpret this finding with caution due to one notable exogenous effect: recruiters may be using generative AI to add required skills. We include more details about this method as well as the top and bottom 5 skills in the Appendix Section E.2

Surprisingly, it does not appear that learning any particular artificial intelligence/machine learning (AI/ML) skills will make inexperienced software developers more employable. Between October 2022 and October 2024, the monthly share of software developer vacancies asking for AI/ML skills[2] rose from 10.5% to 13.1%, suggesting some increase in relative employer demand for AI/ML expertise amid a decline in the monthly count of vacancies. During this same period, the number of junior-level vacancies requiring at least one AI/ML skill dropped by 86.1% (from 1,915 to 1,029) while senior-level AI/ML vacancies only declined by 10.3% (from 4,338 to 3,934). This finding aligns with prior research showing an "escalator" effect (Autor et al. 2006), where technological progress raises the demand for high-skill workers, often leaving middle-skill workers behind.

4.4 Heterogeneous Effects by Firms Size, Geography, and Industry

Given that new technology tends to diffuse unevenly, we explore whether the public release of ChatGPT in November 2022 affected the relative demand for junior versus senior software engineers across industries, firms of different sizes, and geographies. 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 fully saturated model shown in Equation (3).

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

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

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 τ_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 D.2). 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 D.3). 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.

Our industry-level analysis reveals that 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 D.1 in the Supplementary Materials). Perhaps not 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.

[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).

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 92), and Public Administration (NAICS 72);
- **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), Fiance 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 Shifted Demand within Software Development

This paper began with a narrow question within the broad conversation about the impact of generative AI: how does a general-purpose coding technology reshape labor demand within an occupation? The results point to a clear answer. Following the public release of ChatGPT, employer demand shifted away from junior software developers relative to senior developers. That shift appears suddenly, remains stable across several model specifications, and persists when software developers are compared to other technical occupations. This pattern matches emerging economy-wide evidence that early-career workers in highly AI-exposed occupations have experienced disproportionate employment declines (Brynjolfsson et al. 2025a, Demirci et al. 2025).

This phenomenon happened in a specific way. Employers raised experience requirements within existing titles, not merely by eliminating junior-level titles altogether. The remaining early-career vacancies also shifted toward problem solving, interpersonal communication, and attention to detail rather than toward a broad set of AI-specific technical requirements. Software developers show this effect much stronger than the rest of SOC 15 and mechanical

Table 3. 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 B.1. Each regression uses the first sample listed in Appendix Table A..1 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 *: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$.*

engineers (a comparable technical occupation less exposed to generative AI). These results support a theory of within-occupation reallocation which suggests that generative AI altered which parts of software work remained valuable at the entry level and thus changed who counted as an adequate junior hire.

This contribution matters because it shifts attention from worker-level performance gains to the employer's organizational response. This complements recent evidence that generative AI has stage-specific and expertise-dependent effects in creative work by showing how firms translate such heterogeneity into hiring decisions (Hou et al. 2025). Our findings show how productivity gains from generative AI affect broader opportunities. We find that firms may translate those gains into changing labor demands.

5.2 From Task Substitution to Rewriting Vacancies

The results also contribute to task-based accounts of technological change. Classic theories looks at differential effects when technologies automate versus augment tasks. Generative AI balances on this boundary within software development. Our evidence suggests that the early organizational response was to *rewrite job vacancies* to find candidates that were augmented, not automated, by generative AI. Rather than responding only through layoffs or headcount reductions, firms may first respond that the margin of recruitment by increasing experience thresholds and altering skill requirements.

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 can not currently automate: diagnosing problems, having human interactions, and designing larger systems. The empirical pattern in this paper closely fits that mechanism. 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 interpretation also sheds light on a macro-economic trend. Generative AI can improve the productivity of less-experience developers, while simultaneously firms reduce demand for junior work. 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 are clear. 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 (Li et al. 2025).

This creates an issue. 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 (Brynjolfsson et al. 2025a, Demirci et al. 2025).

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 (Kahn 2010, Oreopoulos et al. 2012). 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". The skill results suggest that the durable response is to develop capabilities that complement AI-assisted work: problem solving, communication, and 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 rather than simply fewer in number (Demirci et al. 2025). At the same time, firms that rely on generative AI may need to invest in the early talent pipeline if the want to maintain an internal pathway from junior to 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 employment. 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 but rather reshape the labor market by altering what is most valuable. 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 postings rather than 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 opportunity may be from a short-run adjustment while firms learn where generative AI is and is not useful. For example, recent field evidence has shown the productivity effects can even be negative for experienced developers in a context-heavy scenario (Becker et al. 2025). Over a longer horizon, organizations may redesign roles and new entry-level tasks may emerge.

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

Acemoglu D, Anderson G, Beede DN, Buffington C, Childress EE, Dinlersoz EM, Foster LS, Goldschlag N, Haltiwanger JC, Kroff Z, et al. (2022) *Automation and the workforce: A firm-level view from the 2019 annual business survey*, volume 30659 (National Bureau of Economic Research).

Acemoglu D, Restrepo P (2019) Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives* 33(2):3–30.

Alabdulkareem A, Frank MR, Sun L, AlShebli B, Hidalgo C, Rahwan I (2018) Unpacking the polarization of workplace skills. *Science advances* 4(7):eaao6030.

Alekseeva L, Azar J, Giné M, Samila S, Taska B (2021) The demand for ai skills in the labor market. *Labour economics* 71:102002.

Anderson KA (2017) Skill networks and measures of complex human capital. *Proceedings of the National Academy of Sciences* 114(48):12720–12724.

Arntz M, Gregory T, Zierahn U (2017) Revisiting the risk of automation. *Economics Letters* 159:157–160.

Autor D, Chin C, Salomons A, Seegmiller B (2024) New frontiers: The origins and content of new work, 1940–2018*. *The Quarterly Journal of Economics* 139(3):1399–1465, URL `http://dx.doi.org/10.1093/qje/qjae008`.

Autor DH (2015) Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives* 29(3):3–30, URL `http://dx.doi.org/10.1257/jep.29.3.3`.

Autor DH, Katz LF, Kearney MS (2006) The polarization of the us labor market. *American economic review* 96(2):189–194.

Autor DH, Levy F, Murnane RJ (2003) The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics* 118(4):1279–1333.

Becker J, Rush N, Barnes E, Rein D (2025) Measuring the impact of early-2025 ai on experienced open-source developer productivity. *arXiv preprint arXiv:2507.09089* .

Bessen J (2018) Artificial intelligence and jobs: The role of demand. *The economics of artificial intelligence: an agenda*, 291–307 (University of Chicago Press).

Bettencourt LM, Lobo J, Helbing D, Kühnert C, West GB (2007) Growth, innovation, scaling, and the pace of life in cities. *Proceedings of the national academy of sciences* 104(17):7301–7306.

Bick A, Blandin A, Deming DJ (2026) The rapid adoption of generative ai. *Management Science* .

Brynjolfsson E, Chandar B, Chen R (2025a) Canaries in the coal mine?: Six facts about the recent employment effects of artificial intelligence .

Brynjolfsson E, Li D, Raymond L (2025b) Generative ai at work. *The Quarterly Journal of Economics* 140(2):889–942, URL `http://dx.doi.org/10.1093/qje/qjae044`.

Cui ZK, Demirer M, Jaffe S, Musolff L, Peng S, Salz T (2025) The effects of generative ai on high-skilled work: Evidence from three field experiments with software developers. *Available at SSRN 4945566* .

Deming DJ (2017) The growing importance of social skills in the labor market. *The quarterly journal of economics* 132(4):1593–1640.

Demirci O, Hannane J, Zhu X (2025) Who is ai replacing? the impact of generative ai on online freelancing platforms. *Management Science* 71(10):8097–8108.

Dillender M, Forsythe E (2025) Computerization of white collar jobs. Accepted at the Journal of Labor Economics.

Eloundou T, Manning S, Mishkin P, Rock D (2024) Gpts are gpts: Labor market impact potential of llms. *Science* 384(6702):1306–1308.

Frank MR, Moro E, South T, Rutherford A, Pentland A, Taska B, Rahwan I (2024) Network constraints on worker mobility. *Nature Cities* 1(1):94–104.

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

Hall BH, Khan B (2003) Adoption of new technology.

Handa K, Tamkin A, McCain M, Huang S, Durmus E, Heck S, Mueller J, Hong J, Ritchie S, Belonax T, et al. (2025) Which economic tasks are performed with ai? evidence from millions of claude conversations. *arXiv preprint arXiv:2503.04761* .

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

Hidalgo CA, Klinger B, Barabási AL, Hausmann R (2007) The product space conditions the development of nations. *Science* 317(5837):482–487.

Hou J, Wang L, Wang G, Wang HJ, Yang S (2025) The double-edged roles of generative ai in the creative process: experiments on design work. *Information Systems Research* .

Hu K, Tong A (2023) Chatgpt sets record for fastest-growing user base - analyst note. *Reuters* ChatGPT reached 100 million monthly active users within two months of launch.

Hui X, Reshef O, Zhou L (2024) The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science* 35(6):1977–1989.

Ju A, Sajnani H, Kelly S, Herzig K (2021) A case study of onboarding in software teams: Tasks and strategies. *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, 613–623 (IEEE).

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

Khojah R, Mohamad M, Leitner P, de Oliveira Neto FG (2024) Beyond code generation: An observational study of chatgpt usage in software engineering practice. *Proceedings of the ACM on Software Engineering* 1(FSE):1819–1840.

Li B, Huang N, Shi W (2025) Forced to change? media exposure of labor issues and firm artificial intelligence investment. *Information systems research* .

McElheran K, Li JF, Brynjolfsson E, Kroff Z, Dinlersoz E, Foster L, Zolas N (2024) Ai adoption in america: Who, what, and where. *Journal of Economics & Management Strategy* 33(2):375–415.

Meyer AN, Barr ET, Bird C, Zimmermann T (2019) Today was a good day: The daily life of software developers. *IEEE Transactions on Software Engineering* 47(5):863–880.

Modestino AS, Burke MA, Sadighi S, Sederberg R, Stern T, Taska B (2026a) No longer qualified? changes in the supply and demand for skills within occupations. *ILR Review* 79(3):486–522.

Modestino AS, Burke MA, Sadighi S, Sederberg R, Stern T, Taska B (2026b) No longer qualified? changes in the supply and demand for skills within occupations. *ILR Review* 79(3):486–522.

Modestino AS, Shoag D, Ballance J (2016) Downskilling: changes in employer skill requirements over the business cycle. *Labour Economics* 41:333–347.

Modestino AS, Shoag D, Ballance J (2020) Upskilling: Do employers demand greater skill when workers are plentiful? *Review of Economics and Statistics* 102(4):793–805.

Oreopoulos P, Von Wachter T, Heisz A (2012) The short-and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics* 4(1):1–29.

Otis NG, Clarke R, Koning R (2023) The uneven impact of generative ai on entrepreneurial performance. Working Paper 24-042, Harvard Business School, Cambridge, MA.

Peng S, Kalliamvakou E, Cihon P, Demirer M (2023) The impact of ai on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590* URL `https://arxiv.org/abs/2302.06590`.

Pierce JR, Schott PK (2016) The surprisingly swift decline of us manufacturing employment. *American economic review* 106(7):1632–1662.

Stephany F, Teutloff O (2024) What is the price of a skill? the value of complementarity. *Research Policy* 53(1):104898.

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

Xiao T, Treude C, Hata H, Matsumoto K (2024) Devgpt: Studying developer-chatgpt conversations. *Proceedings of the 21st International Conference on Mining Software Repositories*, 227–230 (New York, NY, USA: Association for Computing Machinery), ISBN 9798400705878, URL `http://dx.doi.org/10.1145/3643991.3648400`.

Supplementary Materials

A. Sample Characteristics

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 $\times$ software job title $\times$ CSA cell within a given month. The third sample includes only postings for which a given employer $\times$ software job title $\times$ 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 A..1. 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 job 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	

Table A..2. 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

Below 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.

Table B.1. Summary Statistics for Independent Variables.

	Quintile 1	*Quintile 2*	*Quintile 3*	*Quintile 4*	*Quintile 5*
Share CSA population as of Oct. 2022	19.4	20.3	24.3	16.4	19.6
Share firm size as of 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 requred 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 Chat-GPT 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 (Hershbein and Kahn 2018) or left these job postings as unassigned (e.g., missing) year of experience (Modestino et al. 2020) 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 C.1 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 C.1. 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.

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 C..2a, 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 Chat-GPT in November 2022 that is similar to what we documented using only the explicitly experience-labeled

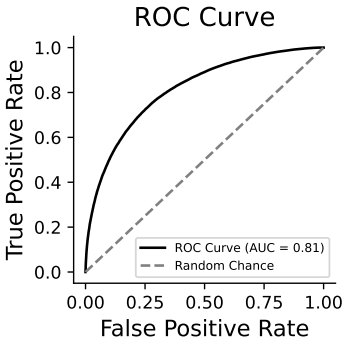
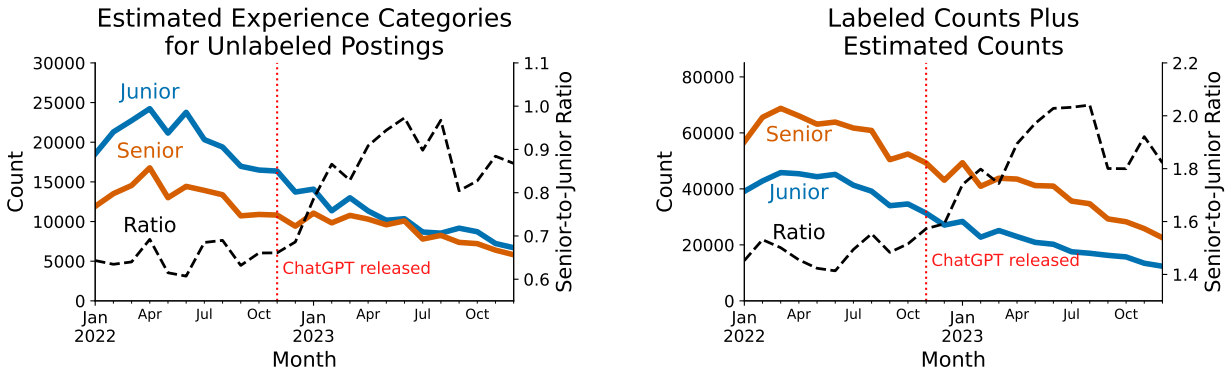


Figure C..1. The receiver operating characteristic curve for the logistic regression classifier with an AUC=0.81. This indicates the experience level classification has a satisfactory false positive and true positive rate.

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 C..2b).

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 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 is does not reflect changes in how employers phrase or omit experience requirements when writing job descriptions.



(a) **(b)**

Figure C..2. (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.

Table C.2. 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 A..1 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 Chat-GPT 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 C.3. 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 A..1 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 Chat-GPT 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 C.4. Difference in Difference Estimates by Skill Level.

	No Experience	>0 but ≤ 3 Years	≥ 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 A..1 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 Chat-GPT 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 D.1. 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 A..1 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 Chat-GPT 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 D.2. 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 A..1 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 Chat-GPT 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 D.3. 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 A..1 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 Chat-GPT 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

E.1 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

E.2 Skill Deviation from Pre-ChatGPT Trends

To study whether particular skills became more or less junior-associated after the release of ChatGPT, we estimate skill-specific deviations from pre-period trends in the share of postings classified as junior. We begin with monthly skill-level counts for SOC 15-1252, where each skill-month ob-

servation contains the total number of postings mentioning the skill, the number of junior postings mentioning the skill, and the number of senior postings mentioning the skill. We restrict this analysis to the top 100 skills by posting volume in each year from 2019 through 2024.

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

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

and analogously the senior share

$$S_{st} = \frac{\text{senior postings}_{st}}{\text{total postings}_{st}}.$$

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

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

where τ_t is a month index normalized so that $\tau_t = 0$ in October 2021. 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. Table D.1

E.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

Table D.1. Skills with the largest positive and negative deviations in junior-posting share relative to their pre-ChatGPT trend. Positive values indicate that the skill appeared in a larger share of junior postings in November 2023 than predicted by its October 2021 to October 2022 trend. Negative values indicate the opposite.

Rank	Skill	Junior share deviation (pp)	Actual junior share (%)	Predicted junior share (%)
1	Interpersonal Communications	24.49	20.85	-3.64
2	Problem Solving	3.57	16.66	13.09
3	Programming Languages	2.22	18.97	16.75
4	Detail Oriented	1.97	25.17	23.19
5	Computer Engineering	1.75	23.66	21.91
96	Artificial Intelligence	-13.51	14.33	27.84
97	Business Process	-13.66	14.93	28.59
98	.NET Framework	-14.14	11.14	25.28
99	Machine Learning	-16.71	16.80	33.50
100	Operations	-16.89	14.09	30.98

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 (Anderson 2017, Stephany and Teutloff 2024), assess the growing cognitive-physical skill polarization (Alabdulkareem et al. 2018), and to predict worker mobility (Frank et al. 2024). 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 Alabdulkareem et al. (2018), 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 Hidalgo et al. (2007) 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 $e(\text{occupation},$

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 θ 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 D.2 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 (Arntz et al. 2017). This means displaced junior software developers could transition into these adjacent roles without extensive retraining, rather than exiting the tech sector altogether.

Table D.2. 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