


Working Paper 25-039

# Displacement or Complementarity? The Labor Market Impact of Generative AI

Wilbur Xinyuan Chen  
Suraj Srinivasan  
Saleh Zakerinia



**Harvard  
Business  
School**

# Displacement or Complementarity? The Labor Market Impact of Generative AI

Wilbur Xinyuan Chen

Hong Kong University of Science and Technology

Suraj Srinivasan

Harvard Business School

Saleh Zakerinia

Harvard Business School

**Working Paper 25-039**

Copyright © 2025 by Wilbur Xinyuan Chen, Suraj Srinivasan, and Saleh Zakerinia.

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Funding for this research was provided in part by Harvard Business School.

# **Displacement or Complementarity? The Labor Market Impact of Generative AI \***

**Wilbur Xinyuan Chen**

Hong Kong University of Science and Technology

**Suraj Srinivasan**

Harvard Business School  
Digital, Data, Design Institute at Harvard

**Saleh Zakerinia**

Harvard Business School  
Digital, Data, Design Institute at Harvard

December 2024

## **Abstract**

Generative AI is poised to reshape the labor market, affecting cognitive and white-collar occupations in ways distinct from past technological revolutions. This study examines whether generative AI displaces workers or augments their jobs by analyzing labor demand and skill requirements across occupations. Our findings reveal a heterogeneous effect: generative AI-driven automation reduces labor demand and skill requirements in structured cognitive-task jobs, while increasing both demand and skill complexity in positions that involve human-AI collaboration. These results highlight the importance of understanding generative AI's nuanced impact on the labor market and designing targeted policies to mitigate job displacement while supporting skills development for human-AI collaboration.

**Keywords:** Generative Artificial Intelligence, Labor Market, Automation and Augmentation

**JEL Codes:** J23, J24, L86, O33

# 1 Introduction

Artificial intelligence (AI) has advanced rapidly in recent decades and is widely regarded as the next industrial revolution, with the potential to transform the economy and redefine the nature of work. This transformation accelerated with the public release of generative AI technologies such as ChatGPT in late 2022. However, the impact of this breakthrough on the labor market remains dynamic and complex. Unlike previous technological revolutions, generative AI can perform sophisticated cognitive tasks such as problem-solving and decision-making that traditionally required human expertise. This capability raises a fundamental question: does generative AI displace an unprecedented number of workers, or does it unlock new productivity gains for them?

While many argue that generative AI, as a capital-augmenting technology, will displace workers by automating tasks previously performed by humans, thereby reducing labor demand (Eisfeldt et al., 2023; Frank et al., 2023), others suggest it could have a positive effect by increasing worker productivity and labor demand (Agrawal et al., 2023b; Noy & Zhang, 2023). This could occur through productivity gains from automating parts of tasks within an occupation (Acemoglu & Restrepo, 2018a) as well as the creation of entirely new tasks enabled by generative AI (Acemoglu & Restrepo, 2019). The net effect of these opposing forces on labor demand determines how generative AI will impact the labor market, making the actual outcomes in labor markets an empirical question.

Theoretical research on automation's labor market effects is well-established (Acemoglu & Restrepo, 2018a, 2019, 2020), but recent studies are only beginning to explore how generative AI may impact jobs, productivity, and firms (Eisfeldt et al., 2023; Eloundou et al., 2024; Felten et al., 2023; Noy & Zhang, 2023). However, most of these studies rely on task-exposure models to estimate job automation rather than analyzing actual labor market outcomes. Moreover, these studies have primarily focused on generative AI's potential to

automate work, relatively little attention has been given to its potential to augment labor. To address this gap, we develop indices that measure both the augmentation and automation potential of generative AI across occupations. Using a near-universe dataset of U.S. job postings, we empirically validate these indices and analyze how generative AI is reshaping labor demand and skill requirements. Specifically, we examine how firms adjust their demand for automation-prone versus augmentation-prone occupations and how skill requirements evolve in response to generative AI adoption.

Our analysis begins by examining the tasks associated with each occupation in the Occupational Information Network (O\*Net) v25.1 dataset. Following the methodology of Eloundou et al. (2024), we utilize OpenAI's ChatGPT to assess the potential for task automation by generative AI. Based on these assessments, we assign an automation score to each occupation. Furthermore, we develop an augmentation score for each task, recognizing that occupations comprising a mix of automatable tasks and those unaffected by generative AI are likely to experience the highest productivity gains. This approach aligns with the concept of labor-augmenting technological change, as discussed by Acemoglu (2002b, 2003) and Acemoglu & Autor (2011).

By combining these automation and augmentation scores with the Lightcast dataset, we study the changes in job postings before and after the introduction of generative AI across various occupational groups. Using a synthetic difference-in-differences approach, we find a 17% decrease in job postings per quarter per firm for occupations in the top quartile of automation potential following the introduction of generative AI, compared to the control group. Conversely, job postings for augmentation-prone occupations increase by 22% per quarter per firm. This heterogeneous impact of generative AI on labor demand underscores the technology's potential as a labor-augmenting force, contrasting with prevailing concerns about its role in reducing overall labor demand. These results suggest that while generative AI may

decrease demand for highly automatable jobs, it simultaneously increases demand for jobs that can be augmented by this technology. This nuanced perspective contributes to the ongoing discourse on the impact of AI on the labor market, highlighting both the challenges and opportunities presented by this emerging technology.

Our next set of analyses focuses on the impact of generative AI on the skills associated with job postings. Utilizing the comprehensive required skill data for each job posting from LightCast, we employ ChatGPT to classify skills and identify those exposed to generative AI. We hypothesize that the number of AI-exposed skills, total required skills, and newly required skills will decrease for jobs highly susceptible to automation, while increasing for jobs prone to augmentation. This hypothesis aligns with the arguments presented by Acemoglu & Restrepo (2018b, 2019, 2020a) and Autor et al. (2024).

Our findings support this hypothesis. Following the introduction of generative AI, we observe a significant 24% decrease in generative AI-exposed skills per firm per quarter among jobs in the top quartile of automation exposure. In contrast, there is a 15% increase in generative AI-exposed skills per firm per quarter for jobs most susceptible to augmentation. Similar trends are mirrored in both the total number of required skills and newly required skills. These results underscore the differential impact of generative AI on skill requirements across various job categories. They suggest that while generative AI may reduce the need for certain skills in highly automatable jobs, it simultaneously creates demand for new skills in jobs that can be effectively augmented by the technology.

Our study contributes to multiple strands of the literature. First, it builds on research examining how technological change affects the labor market by empirically examining whether generative AI leads to job displacement or augmentation. Technological advancements can be factor-specific, augmenting either capital or labor (Acemoglu et al., 2024; Acemoglu & Restrepo, 2019). They can also create new tasks that reshape workforce dynamics, with

implications for job creation and skill evolution (Acemoglu & Restrepo, 2019; Autor Caroline Chin Anna Salomons Bryan Seegmiller et al., 2022). Notably, automation has historically favored capital, and diminished labor's share of production, which has caused downward pressure on wages, especially in the manufacturing sector over the past several decades (Acemoglu & Restrepo, 2019). However, given the distinct characteristics of AI, some researchers argue that it could diverge from previous technological trends by expanding labor demand in areas involving non-automatable tasks (Fossen et al., 2022). On the other hand, other studies suggest that AI may replicate earlier technologies' displacement effects, reducing labor demand overall as automation intensifies (Acemoglu et al., 2022; Agrawal et al., 2023a). This study contributes to this debate by empirically examining generative AI's effects on labor demand, providing evidence on whether generative AI causes displacement or augmentation. To our knowledge, this is the first paper to empirically examine the displacement and augmentation effects of generative AI using the near-universe dataset of U.S. job postings.

Second, this study contributes to the nascent and rapidly growing literature on the diverse impacts of AI. The body of research on AI's labor market impact is substantial, e.g. (Acemoglu et al., 2022; Agrawal et al., 2023a; Fossen et al., 2022; Frank et al., 2023; Green, 2024). Much of this literature relies on AI exposure indices developed by Felten et al. (2018, 2023), Brynjolfsson & Mitchell (2017), Brynjolfsson et al. (2018) and Webb (2019) to measure AI's impact at the firm or industry level, focusing on outcomes such as employment, wage inequality, and productivity. However, these approaches provide limited insight into AI's potential to complement, rather than to replace, human labor. By developing an augmentation index, our study offers a more comprehensive view of generative AI's impact on labor markets.

Recently, scholars have also started to examine the specific impact of generative AI on the labor market. Many of these analyses focus on controlled or specialized settings, such as customer support roles (Brynjolfsson et al., 2023) or online platforms (Hui et al., 2024; Liu et

al., 2023), and often employ lab-based experiments, e.g. (Doshi & Hauser, 2024; Noy & Zhang, 2023), to more precisely evaluate the productivity impact. While these studies demonstrate productivity gains in narrowly defined settings, they stop short of offering a comprehensive, economy-wide view of generative AI’s impact on labor demand and skills. Our study addresses these limitations by empirically analyzing generative AI’s potential for both automation and augmentation across a wide range of occupations. Building on the task-based AI exposure index developed by Eloundou et al. (2024), which has been used to study AI’s economic impacts (Acemoglu, 2024; Eisefeldt et al., 2023), we introduce an augmentation index to capture generative AI’s productivity-enhancing effects alongside its automation potential. By applying these indices to the near-universe dataset of U.S. job postings, we provide a comprehensive analysis of how generative AI reshapes labor demand and skill requirements. Our results offer the first empirical insights into both the displacement and augmentation effects of generative AI across all occupations in the U.S. economy.

## **2 Conceptual Framework**

### **2.1 Generative AI as a Task-Augmenting and Task-Automating Technology**

The literature widely acknowledges that technological change can be either capital-intensive, replacing labor through task automation, or labor-intensive, increasing labor demand by creating new tasks and boosting productivity (Acemoglu, 2002a; Acemoglu & Restrepo, 2018b). Generative AI distinguishes itself from previous waves of technological change by primarily targeting cognitive tasks in white-collar jobs, rather than manual, blue-collar work. This dual capacity—automating cognitive tasks while complementing human capabilities—creates heterogeneous effects across occupations, contingent on their task composition.



Occupations predominantly composed of automatable tasks are likely to experience job displacement as generative AI substitutes human labor. This aligns with the theoretical framework of automation, where technological advancements reduce labor demand by replacing routine and codifiable tasks (Acemoglu & Restrepo, 2019; D. H. Autor et al., 2003a). Generative AI's proficiency in tasks such as text generation, summarization, and language translation diminish the need for human input in automation-prone roles, potentially leading to a decline in job postings.

Conversely, occupations that mix automatable and non-automatable tasks are more likely to benefit from generative AI through productivity enhancements. In these roles, generative AI could complement human labor by improving task efficiency, allowing workers to focus on higher-value activities that require human judgment, creativity, and problem-solving (Acemoglu & Autor, 2011). As productivity increases in augmentation-prone roles, firms may expand hiring to capitalize on these gains, potentially leading to an increase in job postings. Thus, we examine the following hypothesis:

*H1: Generative AI reduces the number of job postings in automation-prone occupations and increases the number of job postings in augmentation-prone occupations.*

## **2.2 Effect of Generative AI on Skills: A Skill-Biased Technological Change Perspective**

Historically, technological change has demonstrated skill-biased effects, which favors skilled workers while diminishing demand for routine tasks (Autor et al., 2003b, 2006). Previous waves of automation largely impacted manual and blue-collar jobs, displacing workers who performed routine physical tasks. In contrast, generative AI marks a departure from this trend by targeting cognitive tasks that are traditionally performed by highly educated, white-collar workers. These tasks include activities such as language translation, text generation, coding, and problem-solving—areas previously insulated from automation (Ellingrud et al., 2023).

At the same time, generative AI has the potential to complement labor in tasks that require creativity, judgment, and advanced reasoning. As a result, its effects on skill demand will differ based on the task composition of occupations. In occupations dominated by automatable tasks, generative AI is likely to reduce the demand for specific skills and simplify job roles. In contrast, occupations that combine automatable tasks with non-automatable tasks are more likely to see an expansion in skill demand as generative AI enhances productivity and creates new opportunities for skill development. Given these dual capabilities, the impact of generative AI on skill requirements depends on the task composition of occupations.

### **2.2.1 Effect on AI-Exposed Skills**

AI-exposed skills are defined as tasks that generative AI can either automate or enhance them. In automation-prone occupations, generative AI automates tasks that involve routine, codifiable cognitive work, which reduces the demand for AI-exposed skills, as workers are replaced by AI systems that can efficiently perform these tasks. In augmentation-prone occupations, however, generative AI enhances AI-exposed skills by amplifying workers' productivity rather than replacing them entirely (Acemoglu & Restrepo, 2020).

To illustrate the augmentation effect, consider a physician interpreting diagnostic tests, such as radiology scans or lab results. Generative AI can automate the routine aspect of analyzing test images or summarizing results, but the physician's expertise remains essential for contextualizing findings, confirming diagnoses, and determining treatment plans. In this scenario, the physician's ability to interact with and effectively oversee AI outputs becomes a valuable skill. Rather than reducing the need for AI-exposed skills, generative AI raises their importance, as workers must integrate AI tools into workflows and use their judgment to add value beyond AI's capabilities (Noy & Zhang, 2023). Thus, generative AI increases the demand for AI-exposed skills in augmentation-prone occupations because workers are

required to interpret, manage, and validate AI-assisted outputs while focusing on higher-order tasks that combine human expertise with AI-enhanced efficiency.

*H2: Generative AI reduces the demand for AI-exposed skills in automation-prone occupations while increasing the demand for these skills in augmentation-prone occupations.*

### **2.2.1 Effect on Total Skills**

The total number of required skills reflects the overall complexity of job roles (D. H. Autor et al., 2003a). In automation-prone occupations, generative AI automates a substantial portion of tasks, simplifying workflows and reducing the need for a broad set of skills. This aligns with theories of deskilling, where automation consolidates job roles and diminishes the range of required competencies (Acemoglu & Restrepo, 2019). In augmentation-prone occupations, generative AI introduces new tools and workflows, expanding the breadth of required skills. Workers must develop complementary capabilities, such as AI literacy, advanced analytical skills, and creative problem-solving, to effectively integrate AI into their tasks. This reflects a process of upskilling, where technology increases the complexity and diversity of work (Bessen, 2018).

*H3: Generative AI decreases the total number of required skills in automation-prone occupations while increasing the total number of required skills in augmentation-prone occupations.*

### **2.2.1 Effect on New Skills**

As argued by D. Autor et al. (2024), new work emerges in response to demand shocks that raise occupational demand. Generative AI, as a technological innovation, acts as a demand shock in the labor market, driving the need for new skills that align with its emerging capabilities. In automation-prone occupations, where tasks are increasingly automated, the demand for new skills declines. Job functions become standardized and consolidated, leaving

workers with fewer opportunities for skill acquisition. This stagnation arises because generative AI replaces routine, codifiable tasks, reducing the scope for skill development. In contrast, in augmentation-prone occupations, generative AI enhances productivity by complementing workers' capabilities, leading to a rise in the demand for new skills. These skills are often linked to the effective use of generative AI technologies, such as AI tool proficiency, creative reasoning, and higher-order problem-solving. As firms integrate AI systems into workflows, workers acquire these new complementary skills to fully leverage AI's capabilities, which drives an increased demand for new specialized and adaptive skill sets.

*H4: Generative AI reduces the emergence of new required skills in automation-prone occupations while increasing the emergence of new required skills in augmentation-prone occupations.*

### **3 Data and Measurement**

We use two primary data sources: task descriptions from the Occupational Information Network (O\*Net) v 25.1 database<sup>1</sup> and U.S. job postings from LightCast, covering the period from 2019 to June 2024. The O\*NET dataset maintained by the U.S. Department of Labor, Employment and Training Administration (USDOL/ETA). This dataset uses the O\*NET-SOC taxonomy, which is based on the Standard Occupational Classification (SOC) system to classify different occupations. The dataset contains 923 occupations of U.S. workers. Additionally, it also provides a rich set of variables for each occupation that describes the nature of work and worker characteristics. In total, the dataset has 19,265 tasks, with detailed descriptions of each of these tasks.

O\*Net dataset captures the near-universe set of U.S. online job vacancies from approximately 40,000 company websites and recruiting websites and includes detailed descriptions of tasks associated with each occupation, allowing us to assess the exposure of

---

<sup>1</sup> <https://www.onetonline.org>

occupations to generative AI. The job posting data includes the posting date, job titles, and skill requirements, enabling us to track changes in labor demand and skill requirements over time.

In order to create our exposure indices, we use task descriptions from the O\*Net dataset. Using these task descriptions, we follow the methodology established by Eloundou et al. (2024) to calculate each occupation's generative AI exposure score. Specifically, we employ OpenAI's GPT-4o model to assess whether each task can be effectively performed by generative AI.

In our prompt, we ask ChatGPT to assign each task to one of the following categories based on its potential exposure to large language models (LLMs):

**E0:** No exposure to LLMs.

**E1:** Direct exposure to LLMs, where access to an LLM can reduce task completion time by at least half without compromising quality.

**E2:** Exposure via LLM-powered applications, where the task cannot be accelerated by direct access to an LLM alone but can be completed more efficiently if supported by additional software developed on top of the LLM.

**E3:** Exposure given image capabilities, where completion time can be reduced by at least half if the LLM is combined with E2 software and systems capable of reading, creating, and interpreting images.

This classification rubric is designed to capture the varying degrees of AI exposure across different tasks, and the full exposure rubric can be shown in Appendix A.

After classifying all tasks, we calculate each occupation's AI exposure score, as shown in Table 1. We assign weights of 1 for E1 tasks (direct exposure) and 0.5 for E2 tasks (exposure through LLM-powered applications), while E0 (no exposure) and E3 (exposure given image capabilities) tasks receive a weight of zero. Each task is further weighted by its importance score in O\*NET, emphasizing tasks that are more critical to the occupation. This weighting scheme captures intermediate exposure, balancing the technology's impact when applied directly or through complementary applications.

Eloundou et al. (2024) validated this methodology by comparing AI exposure scores with human expert evaluations, finding a high correlation. They also developed an alternative rubric for assessing automation potential and reported high correlation scores between the two indices, though only the exposure score was validated with human raters. Therefore, we use the exposure score as a proxy for the automation potential of each occupation, reflecting the extent to which tasks could be automated or effectively performed by generative AI.

For the augmentation score, we calculate the Herfindahl-Hirschman Index (HHI) for each occupation. The intuition behind this method, is that generative AI augments an occupation when it consists of some tasks that are automatable by generative AI but also others that are not automatable by these technologies. To compute this score, we group E1, E2, and E3 tasks as “AI-exposed” tasks and consider E0 tasks as “non-exposed.” Weighted by each task’s importance score, the augmentation score is defined as:

$$\text{Augmentation Score} = 1 - (\text{Share of gen AI exposed tasks}^2 + \text{Share of non gen AI exposed tasks}^2)$$

Table 2 summarizes the distribution of augmentation and automation indices across occupations. Higher scores on the augmentation index indicate greater potential for generative AI to complement, rather than to replace, human labor, while higher scores on the automation index suggest a higher likelihood of automation through generative AI. The larger standard deviation in the augmentation index compared to the automation index also suggests greater variability in augmentation potential across occupations.

Table 3 lists the top 10 occupations with the highest scores in both automation and augmentation indices, showcasing the occupations most prone to automation and those most amenable to augmentation by generative AI.

### 3.1 Skills Classification

To classify different skill demands, we utilize the LightCast dataset and employ ChatGPT to identify which skills in each job posting are exposed to generative AI. While traditional methods often rely on keyword matching to assess technology exposure, e.g., Acemoglu et al. (2022) and Alekseeva et al. (2021), this approach is less suitable for generative AI. Given its potential to impact a broad spectrum of skills, including those not directly related to AI, such as writing, keyword-based methods lack the necessary precision for our analysis. LightCast dataset contains over 33,000 unique skills, making a comprehensive and nuanced classification crucial. To address this challenge, we leverage ChatGPT's advanced language understanding capabilities to categorize skills into four distinct classes, as shown below. The full rubric for this classification is provided in Appendix A.

**S0:** Skills irrelevant to generative AI, with limited impact.

**S1:** Generative AI-relevant skills, essential for developing AI technologies.

**S2:** Skills complemented by generative AI, fundamentally human tasks that can be enhanced by AI.

**S3:** Skills substituted by generative AI, tasks that can be fully performed by AI.

To simplify the analysis, we focus on generative AI-exposed and non-generative AI-exposed skills. For this purpose, we aggregate all S1, S2, and S3 skills as “AI-exposed skills,” while categorizing S0 skills as “non-generative AI-exposed skills.” This aggregation allows us to assess the overall impact of generative AI on skill demands without needing to differentiate between skills that are directly relevant, complemented by, or substitutable through generative AI.

Table 4 shows that approximately 49% of skills in the dataset are not exposed to generative AI, while 51% have some level of AI exposure. This distribution highlights the broad potential reach of generative AI across a variety of skill types, suggesting a significant influence on workforce skill requirements.

Additionally, to calculate the number of new skills required within each firm-occupation, we establish a baseline of required skills posted for each firm-occupation between 2015 and 2019. Starting from 2020, we track required skills on a quarterly basis; any skill appearing within a firm-occupation that was not present in the baseline set is classified as a new skill. This new skill is then added to the baseline set for tracking in subsequent quarters. We repeat this process through the first quarter of 2024 to capture the emergence of new skills for each firm-occupation.

Using this data, we aggregate the number of job postings and associated required skills and new skills by firm and quarter, beginning in 2019. In the main analyses, occupations are divided into quartiles based on their augmentation and automation scores, and we calculate the total number of job postings and required skills for each quartile within each firm. This aggregation enables us to compare trends in job demand, new skills, and skill requirements across occupations with varying degrees of automation and augmentation potential.

Table 5 presents summary statistics for the sample used in the automation analysis, which includes firm-quarter observations for occupations in the top quartile of automation exposure, along with those in the third and fourth quartiles, serving as the comparison group. The dataset contains 76,890 firm-quarter observations, capturing trends in job postings, required skills, generative AI-exposed skills, and new skills. On average, firms in this sample post 450.9 job openings per quarter (median = 94), while those in the top quartile of automation exposure post 565.0 job openings per quarter. Firms in this sample require 756.0 distinct skills per quarter, with those in the top quartile of automation exposure listing 1,450.8 skills per quarter. Generative AI-exposed skills account for 494.8 AI-exposed skills per quarter on average, increasing to 1,070.1 generative AI-exposed skills per quarter in the top quartile of automation exposure. New skill introduction averages 36.7 new skills per firm per



quarter, with firms in the top quartile of automation exposure introducing 70.6 new skills per quarter.

Table S6 reports summary statistics for the sample used in the augmentation analysis, which includes firm-quarter observations for occupations in the top quartile of augmentation exposure, along with those in the third and fourth quartiles as the comparison group. This dataset contains 95,634 firm-quarter observations, providing an overview of job postings and skill measures within this sample. On average, firms in this sample post 458.4 job openings per quarter (median = 97), while those in the top quartile of augmentation exposure post 479.0 job openings per quarter. Firms require 666.2 distinct skills per quarter, with those in the top quartile of augmentation exposure listing 492.6 skills per quarter. Generative AI-exposed skills are less concentrated in this sample compared to automation-exposed samples, averaging 411.0 AI-exposed skills per quarter, with firms in the top augmentation quartile requiring 239.7 AI-exposed skills per quarter. New skill introduction averages 31.2 new skills per firm per quarter, slightly lower than in the automation sample, with firms in the top quartile of augmentation exposure introducing 21.5 new skills per quarter.

## **4 Research Design**

We employ a difference-in-differences approach as our primary research design. The introduction of ChatGPT by OpenAI in November 2022 serves as a source of exogenous variation, providing a plausible shock to the availability of generative AI technology for firms. This event is considered as an exogenous shock in our context because ChatGPT's launch represented a substantial and publicly accessible advancement in generative AI capabilities, which offered generative AI functionality broadly to firms across sectors, regardless of individual firms' prior AI readiness or investment levels. Consequently, its introduction led to a sudden, widespread shift in generative AI technology adoption that was not directly influenced by specific firm-level factors or occupation-specific characteristics.

Our treatment groups consist of occupations in the top quartile of either automation scores or augmentation scores, while occupations in the third and fourth quartiles serve as our control groups. The second quartile is excluded from the control group to mitigate potential spillover effects, as occupations in this middle quartile may still be partially impacted by automation or augmentation forces, which would violate the Stable Unit Treatment Value Assumption (SUTVA).

To yield unbiased estimates, the difference-in-differences (DiD) approach requires treated and control groups to follow parallel pre-treatment trends. However, if pre-treatment trends differ between the treated and control groups, it could bias the results. To address this, we use synthetic difference-in-differences (synthetic DiD), a method that relaxes the parallel trends assumption by combining synthetic control methods (SCM) with DiD (Arkhangelsky et al., 2021).

SCM predicts the outcomes of treated occupations as if they had not been exposed to generative AI by creating a weighted average of control units (Abadie et al., 2010). The weights are chosen to match pre-treatment trends, capturing underlying patterns without requiring strict parallel trends. Unlike SCM, which is typically used for a few treated units and limited control units, synthetic DiD is more scalable and for larger datasets. It constructs a synthetic control for the treated group's average outcome, making it suitable for larger panels and allowing for tighter inferences. Synthetic DiD, like traditional DiD, is invariant to unit-level shifts and produces more reliable estimates in this context of heterogeneous pre-treatment trends.

In our main analysis, we work with a balanced panel of  $N$  units and  $T$  time periods, where the outcome is defined as  $Y_{it}$  for unit  $i$  in period  $t$ . The synthetic DiD procedure estimates parameters by solving the following optimization:

$$(\hat{\tau}^{sdid}, \hat{\mu}, \hat{\alpha}, \hat{\beta}) = \arg \min_{\tau, \mu, \alpha, \beta} \left\{ \sum_{i=1}^N \sum_{t=1}^T (Y_{it} - \mu - \alpha_i - \beta_t - \text{Post Treatment}_{it}\tau)^2 \hat{\omega}_i \hat{\lambda}_t \right\}$$

Here,  $\tau$  represents the average causal effect of the treatment, while  $\mu$  is an intercept term. The parameters  $\alpha_i$  and  $\beta_t$  denote fixed effects for firm-occupation quartiles and quarters, respectively. The indicator variable  $\text{Post Treatment}_{it}$  identifies whether firm-quartile  $i$  is impacted by ChatGPT event at time  $t$ .

The synthetic DiD procedure optimally selects unit weights  $\hat{\omega}_i$ , to ensure that pre-treatment outcomes in the control group mimic the average pre-treatment trajectory of the treated units. Time weights  $\hat{\lambda}_t$  are similarly chosen so that the average post-treatment outcome for each control unit maintains a consistent difference from the weighted average of its pre-treatment outcomes. This method allows us to relax the strict parallel trends assumption required in standard DiD approaches by constructing a synthetic control group tailored to the treated units.

Under general conditions, including cases where treatment assignment correlates with unit-level time trends and treatment effects vary across units, Arkhangelsky et al. (2021) show that synthetic DiD yields a consistent and asymptotically normal estimate of the average treatment effect, provided there are sufficient control units and pre-treatment periods relative to treated units and post-treatment periods. In our study, these assumptions hold, as each firm's treated quartile (quartile 1) has two control quartiles (quartiles 3 and 4), and the pre-treatment period spans from 2019 to the third quarter of 2022, while the post-treatment period extends from the fourth quarter of 2022 to the second quarter of 2024.

As an alternative specification, we also apply a Poisson DiD model, which has a functional form that is better suited for count data such as job postings (Cohn et al., 2022). To account for differences in pre-treatment trends, we weight control units using synthetic weights obtained from SDID, ensuring that pre-treatment trends between treated and control groups are balanced. The model is specified as follows:

$$E(Y_{it}|\cdot) = \exp(\tau \text{Post Treatment}_{it} \times \text{Treated}_i + \alpha_i + \beta_t + \epsilon_{it})$$

Here,  $\text{Treated}_i$  is an indicator variable identifying whether group  $i$  belongs to the treated group. In this specification,  $\tau$  approximates the percentage change in the outcome variable for the treated group post-treatment, compared to the control group post-treatment. By weighting control units using SDID-derived synthetic weights, this approach mitigates biases arising from pre-existing trends and ensures a more robust estimation of treatment effects.

#### 4.1 Constructing Event study plot of synth DiD

To examine how treatment effects evolve over time and verify pre-treatment parallel trends, we construct an event study-style plot using synthetic difference-in-differences (SDID), following the approach in Clarke et al. (2023). This visualization captures evolving differences between treated and synthetic control units, accounting for baseline differences. For each period  $t$ , the event study plot displays the differential outcome:

$$(\bar{Y}_t^{Tr} - \bar{Y}_t^{Co}) - (\bar{Y}_{baseline}^{Tr} - \bar{Y}_{baseline}^{Co})$$

where  $\bar{Y}_t^{Tr}$  and  $\bar{Y}_t^{Co}$  represent the average outcomes of the treated group and synthetic control units at time  $t$ , respectively.  $\bar{Y}_{baseline}^{Tr}$  and  $\bar{Y}_{baseline}^{Co}$  are the optimally weighted pre-treatment means (for details, see Arkhangelsky et al. (2021)). With these weights, this method allows a dynamic comparison of treated and control units relative to baseline differences.

We use block-bootstrap resampling to generate confidence intervals, enabling statistical inference on the differential estimates over time. For further methodological details and applications, see Arkhangelsky et al. (2021).

## 5 Results

### 5.1 Impact on Labor Demand

The introduction of generative AI is reshaping labor demand, but its effects vary across occupations. As shown in Table 7, firms reduce hiring for automation-prone occupations while increasing demand for augmentation-prone jobs. Occupations in the top quartile of exposure to automation experience an average decrease of 95 job postings per firm per quarter, while those in the top quartile of exposure to augmentation see an average increase of 80 job postings. Poisson regression estimates indicate that these changes correspond to a 17% decline in automation-prone jobs and a 22% rise in augmentation-prone jobs at the firm level.

Figure 1 presents the dynamic effects of generative AI's introduction on job postings, using the SDID methodology. We use the launch of ChatGPT by OpenAI in November 2022 as the key event marking the introduction of widely accessible generative AI technology. The results reveal a clear divergence in hiring trends between automation-prone and augmentation-prone occupations following this event. Before November 2022, job postings in these two occupation groups followed similar trajectories, providing evidence that our control groups satisfy the parallel trends assumption. However, after the introduction of generative AI, we observe a consistent and stable divergence: job postings decline in automation-prone occupations while rising in augmentation-prone occupations.

The decline in job postings for automation-prone occupations is likely driven by generative AI's ability to automate repetitive and standardized cognitive tasks more efficiently than humans. As firms integrate generative AI into their workflows, they reduce their reliance on human labor for these tasks, resulting in lower demand for these occupations. This mechanism aligns with prior research on capital-augmenting technologies, which reduce labor demand by automating tasks previously performed by workers

In contrast, the increase in job postings for occupations prone to augmentation supports the hypothesis that generative AI enhances productivity in these roles. These occupations often consist of a mix of tasks—some of which can be automated, while others remain unaffected by AI. The partial automation of tasks allows workers to focus on a smaller set of tasks, which likely shortens the time required to complete these tasks and thereby increases their productivity. Thus, these results support H1

## **5.2 Impact on Skills Demand**

To better understand the mechanisms driving these changes in labor demand, we next analyze how generative AI affects the specific skills required for these occupations. If automation is displacing jobs by automating repetitive tasks, we expect a corresponding reduction in demand for certain skills in automation-prone occupations. Likewise, if AI augments productivity in other roles, we expect an increase in demand for skills that enable workers to effectively leverage AI.

As shown in Table 8, we observe a significant reduction in total unique skills, generative AI-exposed skills, and new skills associated with top automation-exposed occupations following the introduction of generative AI.

These results suggest that automation by generative AI reduces not only AI-exposed skills but also all required skills for these jobs, supporting the mechanism discussed earlier. As generative AI automates routine tasks, the overall complexity of these jobs decreases, leading to fewer required skills. Tasks that previously demanded specialized skills are now automated, simplifying roles and reducing skill demands in automation-prone occupations.

In addition to reducing existing skills, the decline in the emergence of new skills among occupations in the top automation quartile suggests that automation limits the development of new capabilities. As jobs become more streamlined, the need for innovation and the introduction of new skills diminishes, as routine tasks are fully automated. This

further illustrates how generative AI reshapes labor demand, not only by decreasing required skills but also by constraining skill evolution in automation-prone roles. Furthermore, Panel (A) in Figures 2, 3 and 4 confirm we have a parallel trend prior to the introduction of generative AI among treated and control groups, which makes the results meaningful.

Looking at the effect of generative AI on skills in augmentation-prone occupations in Table 9, we observe a significant increase in AI-exposed skills, all required skills, and new skills. The rise in AI-exposed skills aligns with the augmentation mechanism, where generative AI complements rather than replaces human labor. While AI automates certain routine tasks in these roles, workers are still required to manage, collaborate with, and oversee AI systems. As a result, demand for AI-exposed skills expands, as workers must integrate AI into their workflows, to fully leverage the productivity gains in this hybrid human-AI interaction.

The increase in total required skills suggests that as generative AI augments these roles, the overall complexity of the jobs rises. This trend aligns with the augmentation mechanism, where AI automates routine tasks, creating opportunities for workers to integrate AI into other aspects of their jobs. Consequently, these roles require a broader skill set, including both technical AI-exposed skills and additional skills that enable workers to effectively collaborate with AI systems. As firms increasingly adopt AI technologies, the demand for these skills expands, reflecting the growing extent of human-AI collaboration.

The rise in new skills suggests that the introduction of ChatGPT is fostering the development of new skills in augmentation-prone occupations, in contrast to the decline seen in automation-prone roles. This supports the view that generative AI acts as a complementary force for innovation in certain occupations, as described in Autor et al. (2024), by promoting skill development and complementing human labor rather than replacing it. Additionally,

Panel (B) in Figures 2, 3, and 4 show stable and statistically insignificant pre-trends and also the dynamic evolution of the treatment effects.

Our findings strongly support Hypotheses 2, 3, and 4, by providing compelling evidence that the introduction of generative AI impacts skills among top augmented and top automated occupations in ways that align with established theories of technological change. The observed trends in skill demands correspond closely with the predictions put forth by these theoretical frameworks, underscoring the transformative role of generative AI in reshaping the labor market landscape.

## **6 Limitations**

Generative AI is a rapidly evolving technology, and many firms are still in the early stages of adoption of these technologies. Our results so far have only unveiled the short-term effects on the labor market, and so the long-term impacts remain uncertain as adoption scales.

Additionally, using job postings as a proxy for labor demand has its limitations, as some "ghost job postings" (Hao & Simon, 2023) may distort the true picture of hiring trends.

Furthermore, this study focuses on the U.S. labor market, and the effects of generative AI could vary across regions with different technological adoption rates and labor market structures.

Despite these limitations, our findings demonstrate that generative AI has distinct effects on different types of jobs. In automation-prone occupations, AI leads to a reduction in skill demand as tasks are automated. Conversely, in augmentation-prone roles, generative AI increases skill requirements, as workers adapt to new tasks that involve managing and interacting with AI systems. These results highlight the immediate impact of AI on reshaping both job structures and skill demand in a rapidly evolving labor market.



## 7 Conclusion

Our study demonstrates the dual impact of generative AI on the labor market. In automation-prone occupations, generative AI simplifies tasks and reduces the demand for specialized skills, while in augmentation-prone roles, it enhances productivity, increasing the need for more advanced skill sets. These findings highlight the transformative role of generative AI in reshaping occupations, both by changing the demand for labor and the skill sets required. As generative AI continues to evolve, understanding its heterogeneous effects on labor demand is critical. Policymakers and business practitioners must recognize the dual forces of automation and augmentation to ensure that workers are equipped to adapt and thrive in an increasingly AI-integrated labor market.


## References

- Abadie, A., Diamond, A., & Hainmueller, A. J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's Tobacco control program. *Journal of the American Statistical Association*, *105*(490), 493–505. <https://doi.org/10.1198/JASA.2009.AP08746>
- Acemoglu, D. (2002a). Directed Technical Change. *The Review of Economic Studies*, *69*(4), 781–809. <https://doi.org/10.1111/1467-937X.00226>
- Acemoglu, D. (2002b). Technical Change, Inequality, and the Labor Market. *Journal of Economic Literature*, *40*(1), 7–72. <https://doi.org/10.1257/0022051026976>
- Acemoglu, D. (2003). Labor- and capital-augmenting technical change. *Journal of the European Economic Association*, *1*(1), 1–37. <https://doi.org/10.1162/154247603322256756>
- Acemoglu, D. (2024). *The Simple Macroeconomics of AI*. <https://www.nber.org/papers/w32487>
- Acemoglu, D., & Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. *Handbook of Labor Economics*, *4*(PART B), 1043–1171. [https://doi.org/10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5)
- Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, *40*(S1), S293–S340. <https://doi.org/10.1086/718327>
- Acemoglu, D., Kong, F., & Restrepo, P. (2024). *Tasks At Work: Comparative Advantage, Technology and Labor Demand*. <https://doi.org/10.3386/W32872>
- Acemoglu, D., & Restrepo, P. (2018a). Artificial intelligence, automation, and work. In *nber.orgD Acemoglu, P RestrepoThe economics of artificial intelligence: An agenda, 2018•nber.org*. <https://www.nber.org/system/files/chapters/c14027/c14027.pdf>

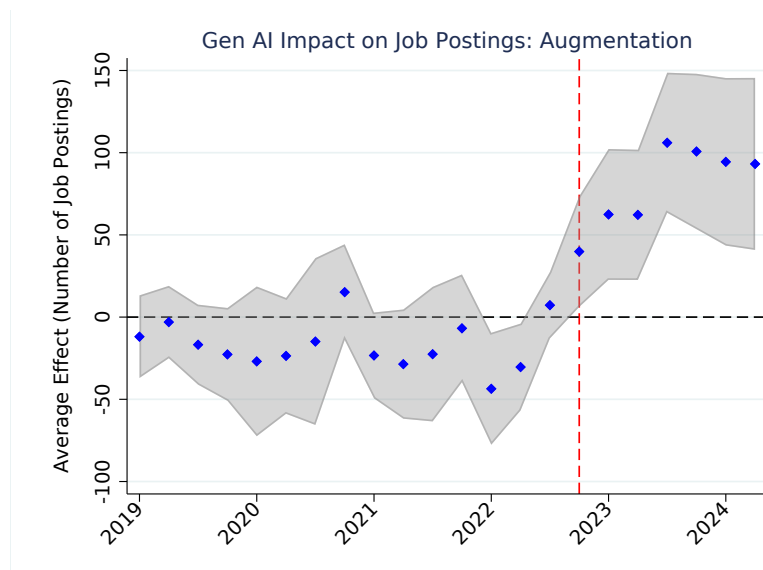
- Acemoglu, D., & Restrepo, P. (2018b). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review*, *108*(6), 1488–1542. <https://doi.org/10.1257/AER.20160696>
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, *33*(2), 3–30. <https://doi.org/10.1257/JEP.33.2.3>
- Acemoglu, D., & Restrepo, P. (2020). The wrong kind of AI? Artificial intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society*, *13*(1), 25–35. <https://doi.org/10.1093/CJRES/RSZ022>
- Agrawal, A., Gans, J., & Goldfarb, A. (2023a). *Similarities and Differences in the Adoption of General Purpose Technologies*. <https://doi.org/10.3386/W30976>
- Agrawal, A., Gans, J. S., & Goldfarb, A. (2023b). Do we want less automation? AI may provide a path to decrease inequality. *Science*, *381*(6654), 155–158. <https://doi.org/10.1126/SCIENCE.ADH9429>
- Alekseeva, L., Azar, J., Giné, M., Samila, S., & Taska, B. (2021). The demand for AI skills in the labor market. *Labour Economics*, *71*, 102002. <https://doi.org/10.1016/J.LABECO.2021.102002>
- Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2021). Synthetic Difference-in-Differences. *American Economic Review*, *111*(12), 4088–4118. <https://doi.org/10.1257/AER.20190159>
- Autor, D. H., Chin, C., Salomons, A., & Seegmiller, B. (2022). New frontiers: The origins and content of new work, 1940–2018. *Academic.Oup.Com*. <https://academic.oup.com/qje/advance-article-abstract/doi/10.1093/qje/qjae008/7630187>
- Autor, D. H., 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. <https://doi.org/10.1093/QJE/QJAE008>
- Autor, D. H., Katz, L. F., Kearney, M. S., Berman, E., & Chandra, A. (2006). The polarization of the U.S. labor market. *American Economic Review*, *96*(2), 189–194. <https://doi.org/10.1257/000282806777212620>
- Autor, D. H., Levy, F., & Murnane, R. J. (2003a). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, *118*(4), 1279–1333. <https://doi.org/10.1162/003355303322552801>
- Autor, D. H., Levy, F., & Murnane, R. J. (2003b). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, *118*(4), 1279–1333. <https://doi.org/10.1162/003355303322552801>
- Bessen, J. (2018). *AI and Jobs: the role of demand*. <https://doi.org/10.3386/W24235>
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at Work*. <https://doi.org/10.3386/W31161>
- Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications: Profound change is coming, but roles for humans remain. *Science*, *358*(6370), 1530–1534. <https://doi.org/10.1126/SCIENCE.AAP8062>

- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What Can Machines Learn, and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings*, 108, 43–47. <https://doi.org/10.1257/PANDP.20181019>
- Clarke, D., Pailaňir, D., Athey, S., arXiv:2301.11859, G. I. preprint, & 2023, undefined. (2023). Synthetic difference in differences estimation. *Arxiv.OrgD Clarke, D Pailaňir, S Athey, G ImbensarXiv Preprint ArXiv:2301.11859, 2023•arxiv.Org*. <https://arxiv.org/abs/2301.11859>
- Cohn, J. B., Liu, Z., & Wardlaw, M. I. (2022). Count (and count-like) data in finance. *Journal of Financial Economics*, 146(2), 529–551. <https://doi.org/10.1016/J.JFINECO.2022.08.004>
- Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28), 5290. [https://doi.org/10.1126/SCIADV.ADN5290/SUPPL\\_FILE/SCIADV.ADN5290\\_SM.PDF](https://doi.org/10.1126/SCIADV.ADN5290/SUPPL_FILE/SCIADV.ADN5290_SM.PDF)
- Eisfeldt, A. L., Schubert, G., Zhang, M. Ben, & Taska, B. (2023). The Labor Impact of Generative AI on Firm Values. *SSRN Electronic Journal*. <https://doi.org/10.2139/SSRN.4436627>
- Ellingrud, K., Sanghvi, S., Dandona, G. S., Madgavkar, A., Chui, M., White, O., & Hasebe, P. (2023). *Generative AI and the future of work in America*. <https://www.mckinsey.com/mgi/our-research/generative-ai-and-the-future-of-work-in-america>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science (New York, N.Y.)*, 384(6702), 1306–1308. <https://doi.org/10.1126/SCIENCE.ADJ0998>
- Felten, E. W., Raj, M., & Seamans, R. (2018). A Method to Link Advances in Artificial Intelligence to Occupational Abilities. *AEA Papers and Proceedings*, 108, 54–57. <https://doi.org/10.1257/PANDP.20181021>
- Felten, E. W., Raj, M., & Seamans, R. (2023). Occupational Heterogeneity in Exposure to Generative AI. *SSRN Electronic Journal*. <https://doi.org/10.2139/SSRN.4414065>
- Fossen, F. M., Samaan, D., & Sorgner, A. (2022). How Are Patented AI, Software and Robot Technologies Related to Wage Changes in the United States? *Frontiers in Artificial Intelligence*, 5, 869282. <https://doi.org/10.3389/FRAI.2022.869282/BIBTEX>
- Frank, M. R., Ahn, Y.-Y., & Moro, E. (2023). *AI exposure predicts unemployment risk*. <http://arxiv.org/abs/2308.02624>
- Green, A. (2024). *Artificial intelligence and the changing demand for skills in the labour market*. [https://www.oecd-ilibrary.org/science-and-technology/artificial-intelligence-and-the-changing-demand-for-skills-in-the-labour-market\\_88684e36-en](https://www.oecd-ilibrary.org/science-and-technology/artificial-intelligence-and-the-changing-demand-for-skills-in-the-labour-market_88684e36-en)
- Hao, S., & Simon, L. (2023, October). *Ghost Job Postings*. <https://www.reveliolabs.com/news/macro/ghost-job-postings/>
- 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. <https://doi.org/10.1287/ORSC.2023.18441>

- Liu, J., Foster, M. G., Nan, X., Li, Y., & Tan, Y. (2023). “Generate” the Future of Work through AI: Empirical Evidence from Online Labor Markets. *SSRN Electronic Journal*.  
<https://doi.org/10.2139/ssrn.4529739>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192.  
[https://doi.org/10.1126/SCIENCE.ADH2586/SUPPL\\_FILE/SCIENCE.ADH2586\\_MDA\\_R\\_REPRODUCIBILITY\\_CHECKLIST.PDF](https://doi.org/10.1126/SCIENCE.ADH2586/SUPPL_FILE/SCIENCE.ADH2586_MDA_R_REPRODUCIBILITY_CHECKLIST.PDF)
- Webb, M. (2019). *The impact of artificial intelligence on the labor market*.  
<http://i8.hexun.com/2019-12-25/199788177.pdf>

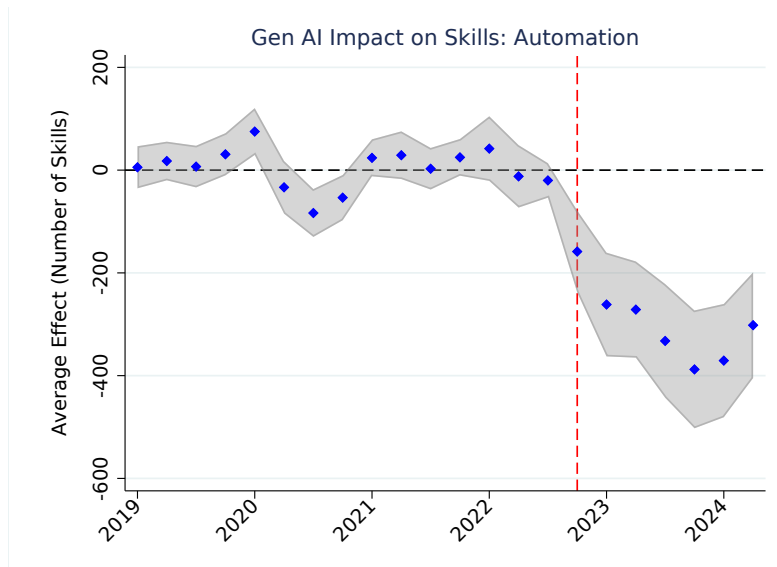


(A) Automation




(B) Augmentation

**Figure 1. Impact of generative AI on job postings across occupations with high automation and augmentation exposure.** Panel (A) shows the decrease in job postings for automation-prone occupations following the introduction of generative AI, while Panel (B) illustrates the increase in job postings for augmentation-prone occupations.




(A) Automation




(B) Augmentation

**Figure 2. Synthetic Difference-in-Differences Estimates of the Impact of Generative AI on Total Required Skills.** (A) Automation-prone occupations. (B) Augmentation-prone occupations.




(A) Automation



(B) Augmentation

**Figure 3. Synthetic Difference-in-Differences Estimates of the Impact of Generative AI on Required AI-Exposed Skills.** (A) Automation-prone occupations. (B) Augmentation-prone occupations.



**Figure 4. Synthetic difference-in-differences estimates of the impact of generative AI on new required skills per job posting in automation-prone and augmentation-prone occupations. (A) Automation-prone occupations. (B) Augmentation-prone occupations.**



**Table 1. Distribution of GPT-4 task classifications across all tasks.**

Exposure (Automation)	Task Count	Task Share
No exposure (E0)	10,611	55.08%
Direct exposure (E1)	5458	28.33%
Exposure by LLM-powered applications (E2)	3111	16.15%
Exposure given image capabilities (E3)	85	0.44%
Total	19,265	100%

**Table 2. Summary statistics of augmentation scores across all occupations.**

Index	Number of Observations	Mean	Standard deviation
Automation Index	923	0.316	0.157
Augmentation Index	923	0.298	0.197

**Table 3. Top occupations exposed to automation and augmentation by generative AI**

Occupation	Automation score	Occupation	Augmentation score
Correspondence Clerks	0.859608	Clinical Neuropsychologists	0.500000
Interpreters and Translators	0.797865	Medical Dosimetrists	0.500000
Court, Municipal, and License Clerks	0.761533	Fish and Game Wardens	0.499999
Medical Transcriptionists	0.742068	Agricultural Engineers	0.499993
Telemarketers	0.734706	Cartographers and Photogrammetrists	0.499991
Word Processors and Typists	0.722573	Traffic Technicians	0.499986
Climate Change Policy Analysts	0.717479	Microbiologists	0.499982
Order Clerks	0.717092	First-Line Supervisors of Police and Detectives	0.499965
Sustainability Specialists	0.71032	Arbitrators, Mediators, and Conciliators	0.499964
Payroll and Timekeeping Clerks	0.703069	First-Line Supervisors of Helpers, Laborers, and Material Movers, Hand	0.499938

**Table 4. Distribution of GPT-4 skill classifications across all skills.**

Exposure	Skills Count	Skill Share
Skills irrelevant to generative AI (s0)	16,422	48.91%
Generative AI relevant skills (s1)	4,097	12.19%
Skills complemented by generative AI (s2)	12,256	36.45%
Skills substituted by generative AI (s3)	825	2.45%
Total	33,620	100%

**Table 5. Summary statistics for the sample used in the automation analyses.** This table reports summary statistics for firm-quarter observations in occupations within the top quartile of automation exposure, along with those in the third and fourth quartiles, which serve as the comparison group.

Variable	Observation	Mean	Median	Std. Dev.	Min	Max
<b>Job Postings (per firm per quarter)</b>						
Total job postings	76,890	450.89	94	1,886.90	0	142,861
Job postings in top quartile of automation	29,634	564.98	90	2,437.85	0	142,861
<b>Skill Measures (per firm per quarter)</b>						
Total required skills	76,890	755.99	156	3,143.18	0	182,020
Total required skills in top quartile of automation	29,634	1,450.79	264	4,871.26	0	182,020
<b>Generative AI-Exposed Skills (per firm per quarter)</b>						
Generative AI-exposed skills	76,890	494.84	71	2,405.15	0	141,418
Generative AI-exposed skills in top quartile of automation	29,634	1,070.08	169	3,774.22	0	141,418
<b>New Skills (per firm per quarter)</b>						
New skills	76,868	36.74	3	205.33	0	20,277
New skills in top quartile of automation	29,612	70.63	7	313.67	0	20,277

**Table 6. Summary statistics for the sample used in the augmentation analyses.** This table reports summary statistics for firm-quarter observations in occupations within the top quartile of augmentation exposure, along with those in the third and fourth quartiles, which serve as the comparison group.

Variable	Observation	Mean	Median	Std. Dev.	Min	Max
<b>Job Postings (per firm per quarter)</b>						
Total job postings	95,634	458.44	97	1,588.06	0	115,855
Job postings in top quartile of automation	39,380	478.99	105	1,395.50	0	34,562
<b>Skill Measures (per firm per quarter)</b>						
Total required skills	95,634	666.18	188	1,981.26	0	101,057
Total required skills in top quartile of automation	39,380	492.65	202	1,097.38	0	32,380
<b>Generative AI-Exposed Skills (per firm per quarter)</b>						
Generative AI-exposed skills	95,634	411.04	87	1,432.82	0	76,019
Generative AI-exposed skills in top quartile of automation	39,380	239.73	90	596.24	0	21,824
<b>New Skills (per firm per quarter)</b>						
New skills	95,634	31.19	4	138.82	0	12,924
New skills in top quartile of automation	39,380	21.51	4	92.11	0	9,295

**Table 7. Impact of ChatGPT introduction on job postings.** This table reports the effects of ChatGPT’s introduction on job postings for occupations in the top quartile of generative AI exposure, compared to those in the third and fourth quartiles. Synthetic Difference-in-Differences (SDID) estimates (Columns 1 and 3) use a synthetic control group based on pre-treatment trends. Poisson regression estimates (Columns 2 and 4) weight control units using SDID-derived synthetic weights to balance pre-treatment trends. SDID standard errors are bootstrapped (1,000 replications); Poisson regression standard errors are clustered at the firm level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

	Automation		Augmentation	
	(1)	(2)	(3)	(4)
	Synth. DiD estimate on Job Postings	DiD estimate on Job Postings (Poisson Regression)	Synth. DiD estimate on Job Postings	DiD estimate on Job Postings (Poisson Regression)
Top Exposed Quartile x Post ChatGPT	-95.43*** (32.55)	-0.169*** (0.0882)	79.84*** (19.33)	0.225*** (0.0406)
Observations	76,890	76,802	95,634	95,634

**Table 8. Impact of ChatGPT introduction on skills demand in automation-prone occupations.** This table reports changes in skill demand for occupations in the top quartile of generative AI exposure to automation, compared to those in the third and fourth quartiles. Synthetic Difference-in-Differences (SDID) estimates (Columns 1, 3, and 5) use a synthetic control group based on pre-treatment trends. Poisson regression estimates (Columns 2, 4, and 6) weight control units using SDID-derived synthetic weights to balance pre-treatment trends. SDID standard errors are bootstrapped (1,000 replications); Poisson regression standard errors are clustered at the firm level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

	All Skills		Generative AI-Exposed Skills		New Skills	
	(1)	(2)	(3)	(4)	(5)	(6)
	Synth. DiD estimate	DiD estimate (Poisson Regression)	Synth. DiD estimate	DiD estimate (Poisson Regression)	Synth. DiD estimate	DiD estimate (Poisson Regression)
Top Exposed Quartile to Automation x Post ChatGPT	-297.7*** (46.08)	-0.238*** (0.0331)	-221.3*** (40.33)	-0.240*** (0.0568)	-30.57*** (8.967)	-0.381*** (0.0123)
Observations	76,890	76,186	76,890	74,580	76,868	62,766



**Table 9. Impact of ChatGPT introduction on skills demand in augmentation-prone occupations.** This table reports changes in skills demand for occupations in the top quartile of generative AI exposure to augmentation, compared to those in the third and fourth quartiles. Synthetic Difference-in-Differences (SDID) estimates (Columns 1, 3, and 5) use a synthetic control group based on pre-treatment trends. Poisson regression estimates (Columns 2, 4, and 6) weight control units using SDID-derived synthetic weights to balance pre-treatment trends. SDID standard errors are bootstrapped (1,000 replications); Poisson regression standard errors are clustered at the firm level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

	All Skills		Generative AI-Exposed Skills		New Skills	
	(1) Synth. DiD estimate	(2) DiD estimate (Poisson Regression)	(3) Synth. DiD estimate	(4) DiD estimate (Poisson Regression)	(5) Synth. DiD estimate	(6) DiD estimate (Poisson Regression)
Top Exposed Quartile to Augmentation x Post ChatGPT	67.29*** (10.23)	0.148*** (0.0114)	38.02*** (6.702)	0.151*** (0.0128)	3.628*** (0.930)	0.167*** (0.0201)
Observations	95,634	95,238	95,634	94,974	95,634	78,246

# **Internet Appendix for Displacement or Complementarity? The Labor Market Impact of Generative AI**

**Wilbur Xinyuan Chen**

Hong Kong University of Science and Technology

**Suraj Srinivasan**

Harvard Business School  
Digital, Data, Design Institute at Harvard

**Saleh Zakerinia**

Harvard Business School  
Digital, Data, Design Institute at Harvard

## Appendix A: GPT Prompts

### GPT prompt for task exposure scores

## E0 - No exposure:

Label tasks E0 if none of the above clearly decrease the time it takes for an experienced worker to complete the task with high quality by at least half. Some examples:

- If a task requires a high degree of human interaction (for example, in-person demonstrations) then it should be classified as E0.
- If a task requires precise measurements then it should be classified as E0.
- If a task requires reviewing visuals in detail then it should be classified as E0.
- If a task requires any use of a hand or walking then it should be classified as E0.
- Tools built on top of the LLM cannot make any decisions that might impact human livelihood (e.g. hiring, grading, etc.). If any part of the task involves collecting inputs to make a final decision (as opposed to analyzing data to inform a decision or make a recommendation) then it should be classified as E0. The LLM can make recommendations.
- Even if tools built on top of the LLM can do a task, if using those tools would not save an experienced worker significant time completing the task, then it should be classified as E0.
- The LLM and systems built on top of it cannot do anything that legally requires a human to perform the task.
- If there is existing technology not powered by an LLM that is commonly used and can complete the task then you should mark the task E0 if using an LLM or LLM-powered tool will not further reduce the time to complete the task. When in doubt, you should default to E0.

## E1 - Direct exposure:

Label tasks E1 if direct access to the LLM through an interface like ChatGPT or the OpenAI playground alone can reduce the time it takes to complete the task with equivalent quality by at least half. This includes tasks that can be reduced to:

- Writing and transforming text and code according to complex instructions,
- Providing edits to existing text or code following specifications,
- Writing code that can help perform a task that used to be done by hand,
- Translating text between languages,
- Summarizing medium-length documents,
- Providing feedback on documents,
- Answering questions about a document,
- Generating questions a user might want to ask about a document,

- Writing questions for an interview or assessment,
- Writing and responding to emails, including ones that involve refuting information or engaging in a negotiation (but only if the negotiation is via written correspondence),
- Maintain records of written data, - Prepare training materials based on general knowledge, or
- Inform anyone of any information via any written or spoken medium.

### ## E2 - Exposure by LLM-powered applications:

Label tasks E2 if having access to the LLM alone may not reduce the time it takes to complete the task by at least half, but it is easy to imagine additional software that could be developed on top of the LLM that would reduce the time it takes to complete the task by half. This software may include capabilities such as:

- Summarizing documents longer than 2000 words and answering questions about those documents,
- Retrieving up-to-date facts from the Internet and using those facts in combination with the LLM capabilities,
- Searching over an organization's existing knowledge, data, or documents and retrieving information,
- Retrieving highly specialized domain knowledge,
- Make recommendations given data or written input,
- Analyze written information to inform decisions,
- Prepare training materials based on highly specialized knowledge,
- Provide counsel on issues, and
- Maintain complex databases.

### ## E3 - Exposure given image capabilities:

Suppose you had access to both the LLM and a system that could view, caption, and create images as well as any systems powered by the LLM (those in E2 above). This system cannot take video as an input and it cannot produce video as an output. This system cannot accurately retrieve very detailed information from image inputs, such as measurements of dimensions within an image. Label tasks as E3 if there is a significant reduction in the time it takes to complete the task given access to a LLM and these image capabilities:

- Reading text from PDFs,
- Scanning images, or
- Creating or editing digital images according to instructions. The images can be realistic but they should not be detailed. The model can identify objects in the image but not relationships between those options

## **GPT prompt for skill exposure score**

You are an expert in AI and technology skill analysis with a deep understanding of generative AI and its applications across various

industries. Your task is to analyze a list of skills and classify each into one of four categories:

1. "Gen AI Relevant Skill" (s1) - Software and technical skills directly related to the development, deployment, or application of generative AI technologies.

Examples include: AI ChatBot, AI KIBIT, ANTLR, Apertium, Artificial Intelligence, Automatic Speech Recognition (ASR), Caffe Deep Learning Framework, Chatbot, Computational Linguistics, Computer Vision, Decision Trees, Deep Learning, Google Cloud Machine Learning Platform, Keras, Latent Dirichlet Allocation, Lexical Semantics, Machine Learning, Microsoft Cognitive Toolkit, MLPACK.

2. "Complemented by Gen AI" (s2) - Augmented skills that are "fundamentally human," but can be enhanced by generative AI tools. Examples include analytical thinking, problem solving, creativity, research, data visualization, strategic planning, predictive analysis, and rapid prototyping.

3. "Substituted by Gen AI" (s3) - Automated skills that will be entirely undertaken by generative AI, with minimal or no human interaction necessary. Examples include image and content generation, data sorting and categorization, forecasting, language translation, simple graphic design, and basic trend spotting.

4. "Irrelevant to Gen AI" (s0) - Limited impact skills that require a human touch, such as complex judgment or nuanced decision-making, that generative AI cannot accomplish. Examples include persuasion and negotiation, motivational leadership, ethical judgment and integrity, compassion, building human relationships, and physical dexterity.

In your classification, consider both the current state of generative AI and emerging trends that might influence the relevance of certain skills in the near future. Provide a classification code (s0, s1, s2, s3) and a brief justification for each classification.