Generative AI and Firm Hiring Demand: Evidence from Developed and Developing Economies

Faezeh Khosravi *[1] and Elaine M. Liu[2]

[1]Georgia State University
[2]Georgia State University and NBER

March 2026

Abstract

We study how the release of ChatGPT in November 2022 affected firm hiring demand across 186 countries, using 77.8 million online job postings from 37,395 firm-country pairs over 2021Q2–2025Q3. We construct predetermined firm-country AI exposure from 2019 occupational composition and estimate a continuous difference-in-differences design. The results reveal a divergent global adjustment. In advanced economies, posting volumes decline gradually beginning three to four quarters after ChatGPT's release, with compositional shifts toward fewer high-substitution vacancies emerging later. In developing economies, posting volumes begin to decline earlier in the event window, with compositional changes following by the third quarter. In both groups, more exposed firms reduce the share of high-substitution postings and lower the posting-weighted AI substitution mix, but the adjustment is faster and more front-loaded in developing economies. These patterns are consistent with generative AI reshaping hiring demand differently across development groups, with developing economies adjusting earlier and with clearer evidence on the composition margin than on overall posting volumes.

Keywords: Generative Artificial Intelligence, Labor Demand, Firm Hiring JEL Codes: J23, O33, J24

*We thank the GSU ARCTIC High-Performance Computing Center for providing computational resources used in data processing and analysis.We also acknowledge the use of generative AI tools to assist with code debugging and editing during the research and writing process. All errors and omissions are our own.

1 Introduction

The public release of ChatGPT in November 2022 marked a salient reduction in the cost of accessing large language models (LMs), rapidly expanding the set of tasks that can be automated or augmented by AI. Within months, workers and firms gained access to tools capable of producing text, code, and structured reasoning outputs at near-zero marginal cost. This diffusion has renewed long-standing concerns about technology-driven labor displacement, while sharpening a more immediate empirical question: has generative AI begun to alter firms' hiring demand, and if so, where and how quickly do such changes emerge? Identifying these effects is challenging because firm-level adoption is rarely observed at scale, employment stocks adjust slowly, and commonly used task-based "exposure" measures do not distinguish between substitution and complementarity (Acemoglu and Restrepo, 2019; D. H. Autor, 2015).

A growing empirical literature examines the labor market impacts of generative artificial intelligence, but existing evidence remains heavily concentrated in high-income economies, particularly the United States and the United Kingdom, reflecting data availability and early adoption patterns (Acemoglu, D. Autor, et al., 2022; Babina et al., 2025; Eloundou et al., 2024; Klein Teeselink, 2025; Y. Liu, Wang, and Yu, 2025; C. Li, 2025). Whether these findings generalize to developing countries remains an open and economically important question.

There are several reasons to expect the labor market effects of generative AI to vary systematically across levels of economic development. First, sectoral and occupational composition differs sharply across countries. High-income economies are more service-oriented and intensive in cognitive, non-routine tasks, which prior work highlights as particularly exposed to recent waves of automation and task reallocation (D. H. Autor, Levy, and Murnane, 2003; Acemoglu and D. Autor, 2011; D. H. Autor, 2015). In contrast, employment in developing countries remains more concentrated in agriculture and routine manufacturing, where the direct applicability of current generative AI tools may be more limited. As a result, the balance between task substitution and task creation may differ across contexts (Acemoglu and Restrepo, 2019).

Second, cross-country differences in AI exposure and adoption may further shape heterogeneity in impacts. Recent work documents substantial variation in potential exposure to AI technologies across countries, driven by differences in occupational structure and task content (Felten, Raj, and Seamans, 2021; Pizzinelli et al., 2023). Language and localization frictions may also matter, as early development and deployment of generative AI systems has been concentrated in English, potentially slowing diffusion in some settings (Gmyrek, Berg, and Bescond, 2023).

At the same time, generative AI may have particularly large effects in developing economies through task expansion and productivity gains. By lowering the skill, language and experience thresholds required to perform a range of cognitive tasks, generative AI may relax binding human-capital and input constraints faced by firms, enabling workers to undertake tasks that were previously infeasible or prohibitively costly. Consistent with this mechanism, recent studies emphasize that AI can act as a complement to human skills and organizational capacity, potentially driving firm growth rather than labor displacement in some contexts (Brynjolfsson, D. Li, and Raymond, 2025; Noy and Zhang, 2023; Mäkelä and Stephany, 2024; Schubert, 2025).

Taken together, the theoretical effects of generative AI on labor demand across countries are ambiguous and potentially heterogeneous in sign and magnitude. This heterogeneity underscores the importance of examining AI's labor market impacts across both developed and developing economies using systematic empirical evidence.

To examine how generative AI affects labor demand across advanced and developing economies, we assemble a novel global panel of online job postings. The balanced event-window estimation sample comprises 77.8 million postings from 37,395 firm-country pairs across 186 countries over the period 2021Q2–2025Q3, collected directly from employers' career websites. We map postings to standardized occupations and industries and aggregate the data to a firm-country-quarter panel. Our baseline analysis uses a balanced event window spanning six quarters before and eleven quarters after the public release of ChatGPT in 2022Q4. The estimation sample is restricted to firm-country pairs with valid 2019 exposure and at least one observed pre-event and post-event quarter, after which missing quarters within the event window are filled with zero postings. This window allows us to study both the immediate and medium-run response of hiring demand while retaining a pre-period that is less contaminated by the most acute phase of the COVID-19 shock.

To measure firms' exposure to AI-driven labor substitution, we use the complementarity-adjusted occupational exposure index developed by Pizzinelli et al. (2023), which builds on Felten, Raj, and Seamans (2021). This measure distinguishes more clearly between substitution and complementarity at the occupation level and is therefore well suited to studying the labor-demand effects of generative AI. We construct firm-country exposure as the posting-share-weighted average of occupation-level AI exposure using 2019 postings only. Fixing exposure in the pre-period ensures that it is predetermined with respect to post-2022 hiring dynamics and reduces concerns that measured exposure is itself shaped by firms' post-treatment adjustments in occupational composition. Throughout, we refer to this measure as predetermined AI exposure, and we define high-substitution occupations as occupations in the top quartile of the occupation-level exposure distribution.

Our baseline empirical strategy is a continuous difference-in-differences design that exploits cross-firm-country variation in predetermined exposure to generative AI. We estimate how job postings evolve after 2022Q4 as a function of baseline exposure, controlling for firm-country fixed effects, country-quarter fixed effects, and industry-by-quarter fixed effects. This specification compares more and less exposed firm-country units within tightly defined country-time and industry-time environments while absorbing broader time-varying shocks at the country-quarter and industry-quarter levels. We report results separately for advanced and developing economies and complement the baseline specification with event-study estimates that assess pre-trends and trace the dynamics of adjustment.

Our analysis uses online job postings as a proxy for labor demand. While postings capture recruiting activity rather than realized employment, they are widely interpreted as forward-looking measures of firms' hiring intentions. In addition, coverage is limited to firms and occupations that recruit through online postings, which tend to be more formal and digitally intensive and may vary across countries. Accordingly, our estimates should be interpreted as changes in recruiting demand within the segment of the labor market observed in the data, rather than as direct measures of employment. The empirical design mitigates these concerns by focusing on within firm-country changes over time and by absorbing common shocks through rich fixed effects, but the interpretation remains one of hiring demand rather than realized labor market outcomes.

We find that the clearest post-2022 response operates through vacancy composition rather than broad-based changes in overall posting volume. Across specifications, more exposed firm-country units reduce the share of postings in high-substitution occupations and lower the posting-weighted AI substitution intensity of their vacancy mix. Event-study estimates indicate that the timing of this adjustment differs across development groups. In advanced economies, the response is more gradual, with exposure-related differences emerging over several quarters after the release of ChatGPT. In developing economies, the adjustment appears earlier in the event window, particularly in posting outcomes, although the most robust evidence across specifications remains compositional rather than aggregate. Reduced-form estimates show stronger aggregate exposure gradients in advanced economies, while occupation-level estimates indicate sharper within-firm reallocation away from more AI-exposed roles in developing economies. Taken together, the results point to a heterogeneous global adjustment in which generative AI reshapes the composition of firms' recruiting demand and, in some settings, the level of posting activity as well, with especially clear evidence of reallocation away from the most AI-substitutable roles.

The paper makes three contributions. First, it provides firm-side evidence on hiring demand using a global dataset of employer-site postings, extending a literature that has so far been concentrated largely in advanced economies. Second, by constructing exposure at the firm-country level from pre-treatment occupational composition, the design isolates differential post-2022 changes within narrow country-time and industry-time cells and reduces concerns that exposure is endogenously determined in the post-treatment period. Third, by emphasizing vacancy composition—including the share of high-substitution postings, the posting-weighted AI substitution mix, and occupation-level posting responses—the analysis speaks directly to recent work distinguishing substitution from complementarity in the measurement of AI exposure. The results suggest that the early labor-market impact of generative AI is visible less in broad-based job disappearance than in a systematic reallocation of hiring demand away from the most AI-substitutable roles.

Our paper contributes to a growing literature on the labor-market consequences of generative AI and related automation technologies. A foundational strand studies how new technologies reallocate tasks between labor and capital and thereby reshape labor demand (D. H. Autor, Levy, and Murnane, 2003; Acemoglu and D. Autor, 2011; Acemoglu and Restrepo, 2019). A second strand develops occupation- and task-based measures of AI exposure, emphasizing that exposure is not equivalent to displacement and that complementarity is central to interpretation (Felten, Raj, and Seamans, 2021; Pizzinelli et al., 2023; Eloundou et al., 2024). A third, emerging strand uses vacancies and other high-frequency firm-side data to study early labor-demand responses to generative AI, typically in a small number of advanced economies (Acemoglu, D. Autor, et al., 2022; Y. Liu, Wang, and Yu, 2025; Klein Teeselink, 2025; C. Li, 2025).[1] We contribute by providing global evidence and by focusing on compositional adjustments in hiring demand, a margin that is central to recent measurement work on AI exposure.

The remainder of the paper proceeds as follows. Section 2 describes the LinkUp postings data, occupational crosswalks, and the construction of predetermined firm-country exposure. Section 3 outlines the empirical strategy. Section 4 presents the main results, mechanism tests, and robustness checks. Section 5 concludes.

2 Data

We construct a novel panel dataset of firm–country job postings using data from LinkUp, accessed through the Dewey Data platform (LinkUp, 2025). LinkUp compiles job listings by

[1] Related work also studies how firms adjust skill demands and organizational practices in response to AI and automation (Hosseini Maasoum and Lichtinger, 2025; Klein Teeselink and Carey, 2026; Ji et al., 2023; J. Liu, 2023)

continuously indexing vacancies posted directly on employers' career websites rather than on third-party job boards or aggregators. As of March 2026, LinkUp indexes more than 330 million job postings from over 80,000 employers across approximately 195 countries, with global coverage updated daily and historical data extending back to 2007.[2] This direct-from-source collection process improves the timeliness and accuracy of observed vacancies and reduces duplicate or expired listings.

The unit of observation in the raw data is an individual job posting. Each posting is identified by a unique `job_hash`, defined as an MD5 hash of the job listing URL. Because postings are indexed at the job–location level, a single vacancy posted in multiple locations appears as multiple observations. We aggregate these postings to construct measures of recruiting demand at the firm–country–occupation–month level and, subsequently, at the firm–country–quarter and firm–occupation–country–quarter levels. Throughout, our measures capture posted vacancies rather than realized hires or employment stocks. Online job postings have become a widely used source of high-frequency firm-level labor demand data in economics research. Hershbein and Kahn (2018) show that vacancy postings track cyclical shifts in skill demand, while Deming and Kahn (2018) document systematic variation in skill requirements across firms and labor markets. These features make job postings well suited to our setting, where the occupational composition of firms' vacancies is central to measuring AI exposure.

Each posting is assigned an occupation using the O*NET taxonomy provided by LinkUp.[3] These occupation codes are generated using a natural-language-processing system that maps job titles and descriptions to standardized O*NET occupations. We then convert O*NET occupations to ISCO-08 codes using an official crosswalk chain linking O*NET-SOC 2019, SOC 2018, SOC 2010, and ISCO-08. Because this chain is not always one-to-one, we construct a unique analysis crosswalk that assigns each O*NET occupation to a single ISCO-08 code. Occupation-level AI exposure measures are then merged at the ISCO-08 level, assigning a time-invariant exposure score to each mapped posting.

For the baseline analysis, we aggregate postings to the firm–country–quarter level and focus on the period from 2021Q2 to 2025Q3, spanning six quarters before and eleven quarters after the public release of ChatGPT in 2022Q4. We construct the main regression sample in

[2] By 2018, LinkUp covered about 70% of U.S. public firms.

[3] A subset of job postings does not have an assigned O*NET occupation code. Because the extract provided through the ARCTIC high-performance computing environment contains only preprocessed and matched records, we are unable to directly observe the total universe of postings and therefore cannot quantify the share of missing occupation assignments in the current draft. We are in the process of obtaining this information and will report the corresponding statistics in a revised version prior to the conference. All analyses in this paper are conducted on the subset of postings with valid occupation mappings and are internally consistent within this matched sample.

three steps. First, we restrict attention to firm–country pairs with positive posting activity in 2019 so that pre-determined exposure can be measured using the pre-treatment occupational composition of postings. Second, we retain firm–country pairs observed at least once before and at least once after the event within the estimation window. Third, we construct a balanced firm–country–quarter panel by adding missing quarters within the event window and setting postings to zero in those cells. The final regression sample therefore consists of a stable set of firm–country pairs with valid pre-treatment exposure and comparable pre/post observability. This sample contains 77.8 million job postings spanning 37,395 firm–country pairs in 186 countries.

Several limitations of the data are worth noting. Because LinkUp is built by scraping vacancies posted on firms' official websites, it does not capture all forms of hiring activity. In particular, informal recruitment, offline hiring, and vacancies filled through alternative channels—such as personal networks, public employment services, or small local job boards—are not observed. These limitations are likely to be more pronounced in developing economies, where informal labor markets are larger and firms are less likely to maintain comprehensive online recruiting infrastructure (Gmyrek, Berg, and Bescond, 2023). As a result, the data may underrepresent hiring demand among smaller, informal, or less digitally connected firms and occupations, particularly at the lower end of the skill distribution.

In addition, the universe of firms scraped by LinkUp is not directly observable to researchers. Consequently, changes in the number of firms appearing in the data over time may reflect either genuine firm entry or changes in scraping coverage, for example if additional firms' websites are incorporated into the indexing system. Our restriction to firm–country pairs with positive 2019 posting activity mitigates this concern by focusing the analysis on units that were already present in the data prior to the widespread diffusion of large language models and are therefore more likely to be observed consistently over time.

A further limitation is that job postings do not map one-to-one to the number of positions firms intend to fill. A single posting may correspond to multiple openings, while multiple postings may reflect repeated advertising for a single position. Accordingly, our measures should be interpreted as proxies for recruiting intensity and stated labor demand rather than exact counts of hires or vacancies. More broadly, our results capture changes in *posted* labor demand and recruiting behavior among a stable set of firms that actively recruit online, rather than changes in aggregate hiring or employment. The empirical design therefore emphasizes within-firm–country variation over time, which helps mitigate concerns related to differential coverage, shifts in scraping scope, and persistent cross-sectional heterogeneity.

2.1 Construction of Key Variables

Our main explanatory variable is an occupation-level measure of exposure to AI that is assigned to job postings based on their mapped occupation. Our baseline measure is the complementarity-adjusted AI occupational exposure index, C_AIOE, from Pizzinelli et al. (2023), which builds on the AI occupational exposure measure of Felten, Raj, and Seamans (2021). Higher values of C_AIOE indicate occupations whose task content is more exposed to AI-driven substitution.

We construct firm–country exposure measures using pre-period (2019) posting distributions, ensuring that exposure is predetermined relative to the diffusion of large language models. Specifically, for each firm–country pair, we compute the posting-share-weighted average of occupation-level C_AIOE scores across occupations posted in 2019. Formally,

$$Exposure_{fc} = \sum_{o \in \Omega} s_{ofc}^{2019} \cdot C\_AIOE_o, \tag{1}$$

where s_{ofc}^{2019} denotes the share of firm f's 2019 postings in country c that belong to occupation o, and C_AIOE_o is the occupation-level AI exposure score. Exposure is therefore fixed prior to the post-2022 period and can vary across countries within the same multinational firm. In the main specifications, we use a standardized version of this measure across firm–country pairs. As a robustness exercise, we construct an alternative exposure measure using occupation-level Gmyrek AI exposure scores, aggregated in the same manner.

We construct several outcomes from aggregated job postings. First, we measure overall posting demand as the total number of job postings in a firm–country–quarter and as an indicator for whether a given firm–country pair posts any vacancy in a quarter. Second, we construct composition measures based on occupation-level classifications. The share of high-substitution postings and the share of high-skill postings are defined relative to postings that can be mapped to occupations, while the posting-weighted AI substitution index is constructed using postings with valid occupation-level exposure scores. Because these composition measures are defined only when mapped or scored postings are observed, they are missing in firm–country–quarters with zero relevant postings.

In addition to the firm–country–quarter panel, we construct a firm–occupation–country–quarter panel to study how hiring demand evolves within narrowly defined occupation cells after the release of ChatGPT. The unit of observation in this panel is a firm–country pair interacted with an occupation group in a given quarter, and the outcome is the count of postings in that occupation cell. This finer level of aggregation allows us to compare hiring changes across occupations with different levels of AI substitutability.

2.2 Descriptive Statistics

Table 1 reports summary statistics separately for advanced and developing economies. We classify economies into advanced and developing groups using a predefined country classification applied consistently throughout the analysis; the full list of economies in the regression sample is reported in the appendix. All results are reported separately for these two groups to capture heterogeneity in labor market structure, digital infrastructure, and technology adoption.

The data reveal substantial differences in both the scale and composition of job postings across the two groups. Average quarterly postings per firm–country pair are significantly higher in advanced economies, while developing economies exhibit higher shares of high-substitution and high-skill occupations in the observed posting sample. This pattern likely reflects differences in the composition of online postings across development groups rather than true differences in the skill intensity of the underlying labor force.

Two complementary mechanisms may account for this pattern. The first is a firm-selection effect. LinkUp indexes vacancies from employer career websites, and firms in developing economies are less likely to be observed unless they are large, internationally visible multinationals or major domestic employers. The developing-economy sample is therefore more concentrated among firms that are already formal, digitally sophisticated, and operating in professional or managerial sectors—precisely the firms most likely to post high-skill vacancies online. Smaller or more informal firms, which account for a large share of employment in developing economies and tend to hire at lower skill levels, are largely absent from the data. The second is a posting-selection effect. Even among firms that are present in the data, those in developing economies may be more selective about which vacancies they advertise online. In contexts where online job posting is less commonly used as a general recruitment channel, firms may reserve their career websites primarily for high-skill or managerial roles while filling lower-skill positions through informal networks, walk-in applications, or local intermediaries. Both mechanisms would generate the pattern observed in Table 1: a developing-economy sample that is compositionally skewed toward high-skill, high-substitution occupations relative to the true occupational distribution of employment in those economies. Accordingly, the compositional differences in Table 1 should be interpreted as features of the observable sample rather than as evidence about the broader labor market. Table 2 reports the distribution of postings across calendar years within the regression sample and shows that observations are well distributed across the pre- and post-periods. Additional appendix tables provide further detail on the geographic, industry, and occupational composition of postings.

3 Empirical Strategy

Our empirical analysis studies how firms and occupations with greater pre-existing exposure to AI-related task substitution changed their posting behavior after the public release of ChatGPT. We define the post period as 2023Q1 onward and construct exposure measures from the occupational composition of 2019 postings, so that exposure is predetermined relative to the shock.

3.1 Reduced-form specification

We begin with a reduced-form specification estimated separately for advanced and developing economies:

$$Y_{fct} = \beta_1 Post_t + \beta_2 \left(Exposure_{fc} \times Post_t\right) + \alpha_{fc} + \varepsilon_{fct}, \tag{2}$$

where Y_{fct} is an outcome for firm f in country c and quarter t, $Post_t$ is an indicator equal to one from 2023Q1 onward, $Exposure_{fc}$ is the predetermined firm-country AI exposure measure constructed from 2019 postings, and α_{fc} denotes firm-country fixed effects. Standard errors are clustered by firm and country.

We estimate this specification separately for advanced and developing economies and consider five outcomes: an indicator for whether the firm-country pair posts any vacancy in a quarter, the total number of postings, the share of postings in high-substitution occupations, the posting-weighted AI substitution mix, and the share of postings in high-skill occupations. The total-postings specification is estimated using Poisson pseudo-maximum likelihood (PPML), while the other outcomes are estimated using ordinary least squares.

This specification is intended to summarize differential post-period changes by exposure within development groups. Because it includes firm-country fixed effects but does not absorb quarter-level shocks, we interpret these estimates as reduced-form evidence rather than as our most saturated specification.

3.2 Occupation-level specification

To sharpen identification at the occupation margin, we estimate an occupation-level specification using firm–occupation–country posting counts:

$$\begin{aligned} Y_{foct} = {} & \gamma_1 \left(Post_t \times Developing_c\right) + \gamma_2 \left(Post_t \times OccExposure_o\right) \\ & + \gamma_3 \left(OccExposure_o \times Developing_c\right) + \gamma_4 \left(Post_t \times OccExposure_o \times Developing_c\right) \\ & + \alpha_f + \alpha_o + \alpha_c + \lambda_t + \varepsilon_{foct}, \end{aligned} \tag{3}$$

where Y_{foct} is the number of postings for firm f, occupation o, country c, and quarter t; $OccExposure_o$ is the occupation-level AI exposure score; $Developing_c$ is an indicator for developing economies; and α_f, α_o, α_c, and λ_t denote firm, occupation, country, and quarter fixed effects, respectively. This specification is estimated using PPML with standard errors clustered by firm and country.

The coefficient $gamma_3$ on $Post_t \times OccExposure_o$ captures the post-ChatGPT change in number of postings for more exposed occupations in advanced economies, while the triple interaction allows that relationship to differ in developing economies. The coefficient γ_4 captures whether occupation-level AI exposure translates more or less strongly into posting reductions in developing economies. On one hand, lower within-occupation skill buffers and weaker adjustment costs may amplify the substitution effect in developing countries. On the other hand, slower technology diffusion, language frictions, and demand expansion effects may dampen realized hiring reductions even in occupations that are highly exposed in principle. The sign of γ_4 is therefore an empirical question.

3.3 Event-study specification

To examine the dynamic evolution of exposure effects, we estimate event-study specifications separately for advanced and developing economies. Specifically, we interact the predetermined firm–country exposure measure with a full set of relative-time indicators around the public release of ChatGPT in 2022Q4:

$$Y_{fct} = \sum_{\substack{k=-4 \\ k \neq -1}}^{11} \beta_k \left(Exposure_{fc} \times 1\{t = 2022\text{Q}4 + k\} \right) + \alpha_{fc} + \lambda_{ct} + \lambda_{nt} + \varepsilon_{fct}, \tag{4}$$

where k indexes event time in quarters relative to 2022Q4, $Exposure_{fc}$ is the predetermined firm–country exposure measure constructed from 2019 postings, α_{fc} denotes firm-country fixed effects, λ_{ct} denotes country-quarter fixed effects, and λ_{nt} denotes NAICS-quarter fixed effects. The window spans $k \in \{-4, \ldots, -2, 0, \ldots, 11\}$, corresponding to 2021Q3 through 2025Q3. The omitted reference period is $k = -1$, corresponding to 2022Q3, the quarter immediately preceding the release of ChatGPT. Coefficients β_k for $k < 0$ serve as pre-trend diagnostics, while coefficients for $k \geq 0$ trace the dynamic hiring response following the public release of ChatGPT in 2022Q4.

We report event-study estimates for four outcomes: an indicator for whether a firm–country pair posts any vacancy in a quarter, the share of postings in high-substitution occupations, the posting-weighted AI substitution mix, and total postings estimated using PPML. For the composition outcomes, the regressions are estimated on the subset of firm–country–quarters with usable occupation information and weighted by total postings. Standard errors are clustered by firm and country.

3.4 Mechanism tests

We next examine whether the exposure gradient is stronger in settings where firms may face greater incentives to substitute away from exposed jobs. For the high-substitution-share outcome, we estimate:

$$Y_{fct} = \delta_1 \left(Exposure_{fc} \times Post_t \right) + \delta_2 \left(Exposure_{fc} \times Post_t \times Z \right) + \alpha_{fc} + \lambda_{ct} + \lambda_{nt} + \varepsilon_{fct}, \quad (5)$$

where Z denotes, in turn, an indicator for manufacturing, low English proficiency, or low internet usage. The fixed effects α_{fc}, λ_{ct}, and λ_{nt} denote firm-country, country-quarter, and NAICS-quarter fixed effects, respectively. These specifications are estimated on the pooled sample with standard errors clustered by firm and country.

The country-level mechanism variables are constructed from external 2021 measures. Low internet is an indicator for countries below the sample median of World Bank internet usage, measured as individuals using the Internet as a share of the population. Low English is an indicator for countries below the sample median of the 2021 EF English Proficiency Index.

3.5 Richer baseline and robustness checks

Our appendix reports a richer baseline specification of the form:

$$Y_{fct} = \theta \left(Exposure_{fc} \times Post_t \right) + \alpha_{fc} + \lambda_{ct} + \lambda_{nt} + \varepsilon_{fct}, \quad (6)$$

where α_{fc}, λ_{ct}, and λ_{nt} denote firm-country, country-quarter, and NAICS-quarter fixed effects. Because the post-period indicator is collinear with these time-varying fixed effects, this specification identifies only the differential exposure effect. We report these estimates in Appendix Table A4.

We assess robustness in several ways. First, we re-estimate the reduced-form specification clustering at the country level only (Appendix Table A5). Second, we examine task-composition outcomes built from occupation-level task measures (Appendix Table A6). Third, we test sensitivity to alternative thresholds for defining highly exposed occupations (Appendix Table A7). Fourth, we compare high-AI and low-AI occupations directly in an occupation-level PPML specification with firm-occupation-country and quarter fixed effects (Appendix Table A8). Fifth, we re-estimate the reduced-form specification using an alternative firm–country exposure measure constructed from occupation-level Gmyrek AI exposure scores (Appendix Table A9).

Our identification strategy exploits cross-firm-country differences in predetermined occupational exposure to AI-related substitution risk and compares how those differences translate into post-2022 changes in posting behavior. The reduced-form estimates summarize differential post-period responses within development groups, while the occupation-level and appendix specifications provide more saturated evidence by absorbing broader time-varying shocks and focusing on within-occupation adjustment. Throughout, we interpret the estimates as changes in posted vacancies rather than realized hiring or employment.

4 Results

4.1 Reduced-form evidence

Table 3 reports reduced-form estimates separately for advanced and developing economies. These specifications include firm-country fixed effects and compare post-period changes across firms with different levels of predetermined exposure.

In advanced economies (Panel A), posting activity declines sharply in the post period. The probability that a firm-country pair posts any vacancy falls by 10.4 percentage points (pp), and total postings also decline substantially in the PPML specification. More importantly, these declines are larger for more exposed firms. The interaction between exposure and the post period is negative and statistically significant for both the extensive margin and total postings, indicating that firms with higher pre-determined exposure experience a larger post-ChatGPT contraction in recruiting demand.

The composition results point in the same direction. The share of postings in high-substitution occupations and the posting-weighted AI exposure mix both decline in the post period, and both outcomes fall further for more exposed firms. By contrast, the interaction effect for the high-skill share is small and statistically insignificant. Taken together, the advanced-economy results suggest that the main response operates through lower vacancy posting overall and a reallocation away from more substitution-prone jobs.

In developing economies (Panel B), the level effects are weaker and less precisely estimated. Average post-period posting activity is lower for some outcomes, but the interaction between exposure and the post period is generally not statistically significant for vacancy levels. The clearest composition effect appears in the AI exposure mix, where higher exposure is associated with a statistically significant post-period decline. This indicates that in developing economies adjustment occurs primarily through the composition of posted vacancies rather than through a broad contraction in posting levels.

Overall, Table 3 shows that the relationship between AI exposure and post-ChatGPT hiring adjustments is much stronger in advanced economies than in developing economies. The basic pattern is robust to alternative exposure definitions. Appendix Table A9 re-estimates the reduced-form specification using the Gmyrek exposure index and yields qualitatively similar conclusions, especially for composition outcomes. Appendix Table A11 shows that the main AI exposure measure is positively but not perfectly correlated with alternative occupation-level indices, suggesting that the robustness results are informative rather than mechanical.

4.2 Occupation-level evidence

Table 4 reports the occupation-level specification, which absorbs firm, occupation, country, and quarter fixed effects. This specification asks whether, within the same firm and country, postings decline more in occupations with higher pre-ChatGPT AI substitution exposure. Occupation-level exposure (C_AIOE) is demeaned prior to constructing interaction terms, so all coefficients are evaluated at mean exposure.

Three findings emerge. First, the coefficient on $Post \times Developing$ is positive and significant ($\hat{\gamma}_1 = 0.22$), indicating that in the post-ChatGPT period, developing-country firms post more in occupations with average AI exposure relative to advanced economies. This likely reflects the sample composition: developing-country firms in our data are concentrated among large, formal employers recruiting primarily in professional and managerial occupations, where the post-period trend in posting activity is positive at mean exposure levels.

Second, the coefficient on $Post \times OccExposure$ is negative and significant, indicating that in advanced economies, an increase in occupation-level AI exposure is associated with a decline in postings after ChatGPT's release. This confirms that the compositional shift documented in the reduced-form results reflects reallocation across occupations within firms, rather than changes in firm composition alone.

Third, the triple interaction $Post \times Occ\widetilde{Exposure} \times Developing$ is also negative and significant ($\hat{\gamma}_4 = -0.11$), implying that the exposure gradient is steeper in developing economies. The total post-period exposure effect in developing economies is $\hat{\gamma}_2 + \hat{\gamma}_4 = -0.21$, approximately twice as large as in advanced economies. This finding is consistent with the mechanisms discussed in Section **??**: developing-country operations tend to concentrate AI-exposed roles in more routine tasks, and the cost-benefit calculation for substituting these positions is more favorable at lower prevailing wages. Taken together, the occupation-level estimates indicate that generative AI is associated with a reallocation of hiring away from more AI-substitutable occupations within firms, and that this reallocation is sharper in developing economies than in advanced economies.

4.3 Event Study Results

Figures 1 and 2 present event-study estimates for posting and composition outcomes. The coefficients trace the evolution of exposure-related differences over time relative to the omitted pre-period. Pre-period coefficients are generally small and imprecisely estimated, providing no strong evidence of systematic differential pre-trends across exposure levels.

The dynamic patterns differ sharply between advanced and developing economies, both in the timing and in the sequencing of quantity and compositional adjustments.

In advanced economies, the adjustment is gradual and unfolds in two distinct stages. Overall posting levels begin declining only around three to four quarters after ChatGPT's release, consistent with the delayed quantity response documented in the U.S. literature. Y. Liu, Wang, and Yu (2025), using *Lightcast* postings, find that quantity effects in high-substitution occupations intensified gradually over time — rising from roughly 6% in the first year after ChatGPT's launch to 18% by the third year — and Acemoglu, D. Autor, et al. (2022) similarly document that aggregate employment effects of AI adoption are difficult to detect in the short run even when firm-level adjustments are underway. Share of high-substitution postings and posting-weighted AI substitution decline gradually, and they become statistically different from zero in quarters seven through eleven. This sequencing suggests that in advanced economies, firms first reduce overall recruiting activity among more exposed units before restructuring the occupational composition of their remaining vacancies. Institutional factors such as longer hiring cycles, employment protection legislation, and more complex organizational planning processes may contribute to this drawn-out adjustment. In developing economies, the same two-stage pattern unfolds faster and with sharper magnitudes. Posting levels among more-exposed firms decline immediately upon ChatGPT's release, beginning in the first post-period quarter and persisting through approximately the fifth quarter. Compositional changes follow sooner than in advanced economies: the share of high-substitution postings and the posting-weighted AI substitution mix both become statistically significant by the third quarter after ChatGPT's release. The earlier onset of both stages relative to advanced economies points to a different underlying mechanism. One plausible explanation is centralized workforce planning by multinational firms, whose headquarters-level decisions to pause hiring in AI-exposed roles propagate quickly to subsidiary operations in developing countries, generating an immediate quantity response that precedes any deliberate occupational restructuring. Under this interpretation, the reversion of the quantity coefficient toward zero after approximately four to five quarters reflects the resolution of an initial "wait and see" period rather than a recovery in posting levels *per se*, after which firms proceed to restructure the composition of their remaining vacancies. An alternative explanation is that substitution occurs faster in developing-country operations, where AI-exposed roles tend to be more routine and the cost-benefit calculation for replacing workers with AI tools is more favorable given lower prevailing wages (Acemoglu and Restrepo, 2019). Distinguishing between these mechanisms would require firm-level data on headquarter location and subsidiary structure, which we leave for future work.

4.4 Mechanism analysis

Table 5 examines whether the exposure gradient varies across settings that may affect firms' ability to substitute tasks. The dependent variable in all three columns is the share of postings in high-substitution occupations.

Across specifications, the interaction between exposure and the post period remains negative and statistically significant, indicating that more exposed firms reduce the share of postings in highly exposed occupations after the introduction of ChatGPT. The negative interaction with manufacturing suggests that this relationship is stronger in manufacturing sectors, consistent with the view that substitution may be easier in production environments with more standardized tasks.

We also find some evidence that the exposure effect is attenuated in countries with lower English proficiency, as indicated by the positive interaction with the low-English indicator. By contrast, the interaction with low internet usage is positive but not statistically significant.

Additional appendix evidence points toward task reallocation rather than a purely mechanical reduction in hiring. Appendix Table A6 shows that in advanced economies, higher exposure is associated with increases in communication and responsibility content in posted vacancies, as well as an increase in the non-routine task index, while the coefficient on the skills component is small and statistically insignificant. This pattern is consistent with firms shifting the composition of vacancies toward task bundles that are less routine and less directly exposed to AI-driven substitution.

4.5 Robustness tests

Table 6 presents robustness checks for the high-substitution-share outcome. The baseline estimate remains negative and statistically significant when excluding U.S. firms. The same qualitative result also survives when exposure is measured using alternative raw and decile-based definitions.

The main composition result is also robust to alternative definitions of high-exposure occupations. Appendix Table A7 shows that the negative relationship between exposure and the share of highly exposed postings remains present when high-exposure occupations are defined using terciles, quartiles, quintiles, or deciles of the occupation-level exposure distribution. In advanced economies the coefficients are consistently negative and statistically significant across all thresholds, while in developing economies the same pattern is present but somewhat less precise, especially under the strictest decile cutoff. Appendix Table A5 further shows that the reduced-form results are similar when standard errors are clustered at the country level rather than by firm and country. As an additional falsification exercise, Appendix Table A10 assigns a placebo shock date of 2022Q2 and re-estimates the composition specifications using only pre-ChatGPT periods. The resulting coefficients are small and statistically insignificant across both outcomes and both development groups, which is consistent with the absence of strong differential pre-trends before the actual post period.

5 Conclusion

This paper studies how the public release of ChatGPT in November 2022 affected firm hiring demand across advanced and developing economies using a novel global panel of online job postings. Leveraging predetermined firm-country exposure to AI based on pre-2022 occupational composition, we estimate how more and less exposed firms adjust their recruiting behavior following the diffusion of generative AI.

Our findings point to a heterogeneous but structured global adjustment. Across specifications, the most consistent evidence indicates that generative AI primarily affects the composition of hiring demand rather than leading to broad-based reductions in posting volume. More exposed firms reduce the share of vacancies in high-substitution occupations and lower the posting-weighted AI substitution intensity of their vacancy mix. These compositional changes are robust across alternative specifications, exposure measures, and sample restrictions.

Event-study estimates suggest that the timing of adjustment differs across development groups. In advanced economies, exposure-related differences in hiring outcomes emerge gradually over several quarters following the release of ChatGPT. In developing economies, adjustment appears earlier in the event window, particularly in posting outcomes, although evidence on posting volumes is less precise and less consistent across specifications than for compositional measures. Complementary occupation-level estimates indicate sharper within-firm reallocation away from more AI-exposed roles in developing economies, even as aggregate reduced-form exposure gradients are stronger in advanced economies.

Taken together, these findings suggest that the early labor-market impact of generative AI is best understood as a reallocation of hiring demand across tasks and occupations rather than as a uniform decline in recruiting activity. Firms appear to adjust what they hire for more than how much they hire, with especially clear reductions in demand for the most AI-substitutable roles.

These results have several implications. First, they highlight the importance of distinguishing between changes in the level and the composition of labor demand when assessing the impact of new technologies. Second, they suggest that evidence from advanced economies may not fully capture the timing and nature of adjustment in developing economies, where responses may emerge earlier along certain margins. Finally, the findings underscore the value of high-frequency, firm-side data in detecting early labor-market responses to rapidly diffusing technologies such as generative AI.

Several limitations should be noted. Our analysis relies on online job postings, which capture recruiting activity rather than realized employment and may overrepresent formal and digitally intensive sectors. In addition, while the use of predetermined exposure mitigates concerns about endogenous adjustment, our estimates should be interpreted as reduced-form responses rather than fully structural causal effects. Future work could extend this analysis using linked employer-employee data, alternative measures of AI adoption, and longer time horizons to assess the persistence of these effects.

Overall, our results provide early global evidence that generative AI is reshaping firm hiring behavior, with effects that vary across countries and operate primarily through the composition of labor demand.

6 Tables and Figures

Table 1: Summary statistics by development group

	Advanced	Developing	Difference
Job postings (quarterly)	143.774	26.663	117.111***
	(1274.248)	(252.589)	(3.136)
Share high-sub occupations	0.310	0.372	-0.061***
	(0.288)	(0.330)	(0.001)
Share high-skill occupations	0.764	0.875	-0.111***
	(0.292)	(0.228)	(0.001)
Posting-weighted AI substitution mix	4.639	4.735	-0.096***
	(0.303)	(0.291)	(0.001)
Predetermined exposure (2019, raw)	4.657	4.727	-0.070***
	(0.282)	(0.271)	(0.001)
Predetermined exposure (2019, standardized)	-0.052	0.176	-0.228***
	(0.926)	(0.887)	(0.003)
Total postings	73,406,868	4,459,629	
Firm-country pairs	28,365	9,292	
Firms	14,635	1,040	
Countries	51	138	
Firm-country-quarter observations	510,570	167,256	

Notes: The unit of observation is a firm-country-quarter in the balanced regression sample. Entries in columns (1) and (2) report means, with standard deviations in parentheses. Column (3) reports the difference in means between advanced and developing economies (advanced minus developing), with standard errors in parentheses. Job postings are quarterly posting counts. Share high-sub occupations is the share of mapped postings in occupations in the top quartile of the occupation-level AI exposure distribution. Share high-skill occupations is the share of mapped postings in ISCO major groups 1–3. Posting-weighted AI substitution mix is the average occupation-level AI exposure score across postings with valid occupation mappings and scores. The predetermined exposure measures are constructed from 2019 posting distributions. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table 2: Distribution of postings across calendar years

Year	Advanced (%)	Developing (%)
2021	20.373	17.427
2022	26.423	25.767
2023	22.329	22.232
2024	18.191	19.975
2025	12.684	14.598

Notes: Entries report the share of total postings accounted for by each calendar year within the balanced regression sample, separately for advanced and developing economies. The event window runs from 2021Q2 to 2025Q3, so 2021 and 2025 are only partially observed.

Table 3: Reduced-form estimates by development group

	(1) Any posting	(2) Postings	(3) High-sub	(4) AI mix	(5) High-skill
Panel A. Advanced countries					
Post period	-0.104***	-0.341***	-0.012***	-0.011**	0.005
	(0.007)	(0.025)	(0.003)	(0.005)	(0.004)
Exposure × post	-0.011***	-0.045***	-0.002*	-0.002*	-0.001
	(0.004)	(0.014)	(0.001)	(0.001)	(0.001)
Observations	510,570	510,570	403,577	403,577	403,577
SE clustering	Firm, country				
Firm-country FE	Yes	Yes	Yes	Yes	Yes
Panel B. Developing countries					
Post period	-0.082***	-0.146	0.012***	0.009***	0.010***
	(0.009)	(0.149)	(0.003)	(0.003)	(0.003)
Exposure × post	-0.003	-0.121	-0.005	-0.009***	0.000
	(0.005)	(0.103)	(0.003)	(0.003)	(0.002)
Observations	167,256	167,256	118,246	118,246	118,246
SE clustering	Firm, country				
Firm-country FE	Yes	Yes	Yes	Yes	Yes

Notes: Entries report reduced-form estimates by development group. Standard errors clustered by firm and country are reported in parentheses. All specifications include firm-country fixed effects only. Columns (1) and (2) are estimated on the full sample. Column (2) is estimated using Poisson pseudo-maximum likelihood (PPML). Columns (3), (4), and (5) are conditional on having positive mapped postings, so the number of observations is smaller in those columns. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4: Occupation-level estimates

	(1) # Postings
Post × developing	0.233***
	(0.059)
Post × occupation exposure	-0.101***
	(0.016)
Occupation exposure × developing	0.423***
	(0.124)
Post × occupation exposure × developing	-0.111**
	(0.055)
Observations	6,264,707
SE clustering	Firm, country
Firm FE	Yes
Occupation FE	Yes
Country FE	Yes
Quarter FE	Yes

Notes: Entries report occupation-level PPML estimates. Standard errors clustered by firm and country are reported in parentheses. Specifications absorb firm, occupation, country, and quarter fixed effects. Occupation exposure (C_AIOE) is demeaned to its sample mean prior to constructing interaction terms. Coefficients on lower-order interaction terms are therefore evaluated at mean occupation-level AI exposure. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table 5: Mechanism tests: high-substitution share

	(1) Manufacturing	(2) Low English	(3) Low Internet
Exposure × post	-0.007***	-0.009***	-0.009***
	(0.002)	(0.001)	(0.002)
Exposure × post × manufacturing	-0.005*		
	(0.003)		
Exposure × post × low English		0.008*	
		(0.005)	
Exposure × post × low internet			0.005
			(0.006)
Observations	521,435	521,435	521,435
SE clustering		Firm, country	
Firm-country FE	Yes	Yes	Yes
Country×quarter FE	Yes	Yes	Yes
NAICS×quarter FE	Yes	Yes	Yes

Notes: Entries report mechanism-test estimates for the high-substitution-share outcome. For this analysis, we merge two country-level indicators. The first is a low-internet indicator constructed from World Bank data on individuals using the Internet (% of population). The second is a low-English indicator constructed from the 2021 EF English Proficiency Index (EF EPI), a country-level measure of English proficiency compiled by Education First. In each case, we define the indicator as equal to one for countries below the sample median of the corresponding 2021 measure. Standard errors clustered by firm and country are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table 6: Robustness checks: high-substitution share

	(1) Baseline	(2) No US	(3) Raw exp.	(4) Exp. decile
Exposure × post	-0.010***	-0.007***		
	(0.002)	(0.003)		
Raw exposure × post			-0.031***	
			(0.007)	
Exposure decile × post				-0.003***
				(0.001)
Observations	521,435	302,571	521,435	521,435
Firm-country FE	Yes	Yes	Yes	Yes
Country×quarter FE	Yes	Yes	Yes	Yes
NAICS×quarter FE	Yes	Yes	Yes	Yes

Notes: Entries report robustness estimates for the high-substitution-share outcome. Standard errors clustered by firm and country are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

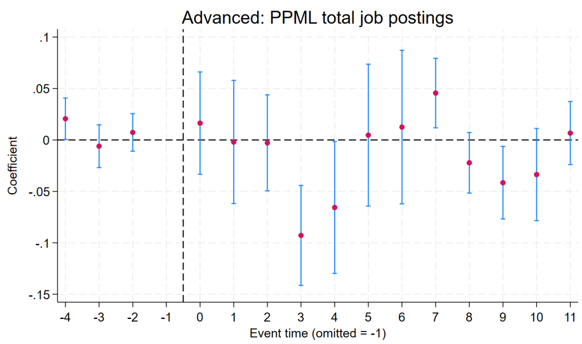

(a) PPML, advanced economies

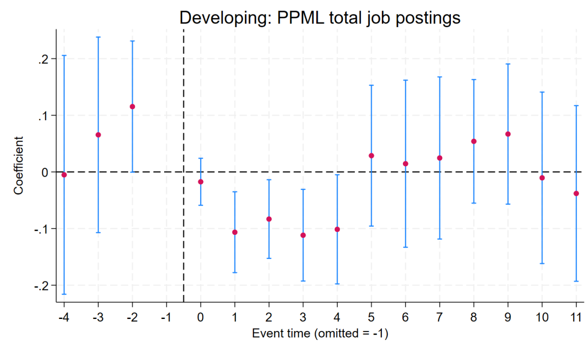

(b) PPML, developing economies

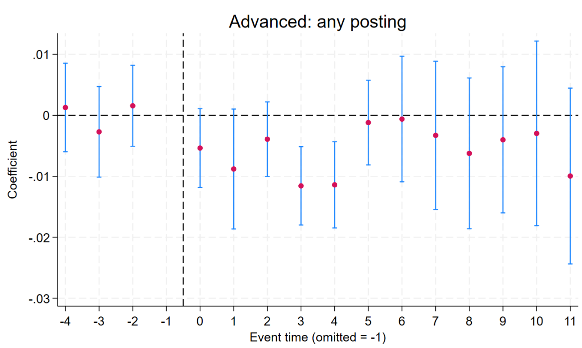

(c) Any posting, advanced economies

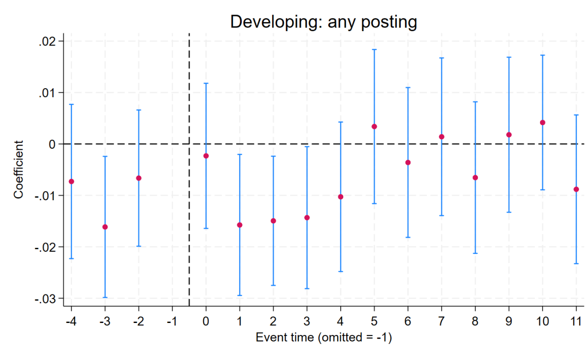

(d) Any posting, developing economies

Figure 1: Event-study estimates for posting outcomes by development group

Notes: The figure reports event-study coefficients for posting outcomes by development group. Panels (a) and (b) show PPML estimates for total postings in advanced and developing economies, respectively. Panels (c) and (d) show estimates for the extensive-margin outcome indicating whether a firm-country pair posts any vacancy in a quarter. Coefficients are plotted relative to the omitted pre-period. Vertical bars denote 95% confidence intervals.

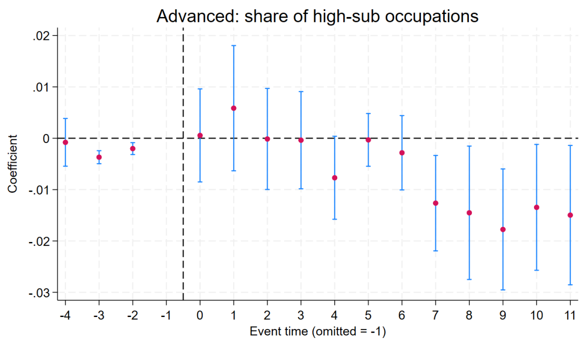

(a) High-substitution share, advanced economies

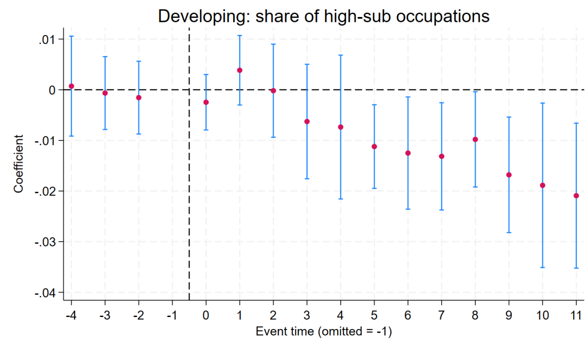

(b) High-substitution share, developing economies

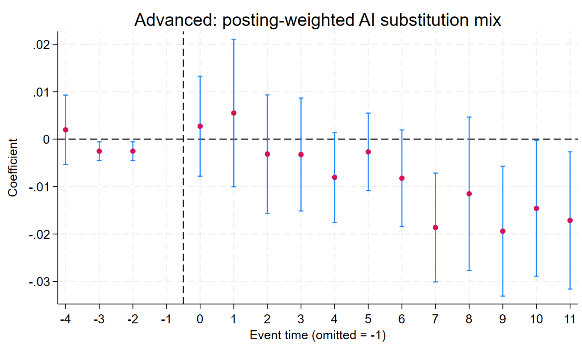

(c) AI exposure mix, advanced economies

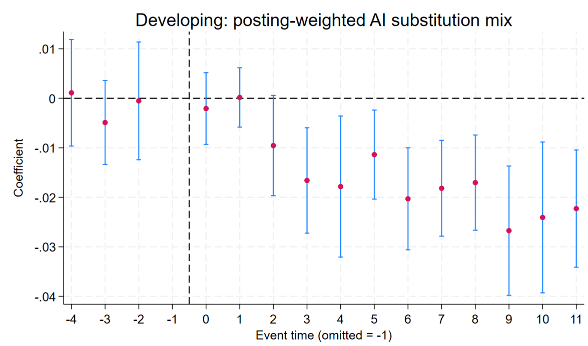

(d) AI exposure mix, developing economies

Figure 2: Event-study estimates for composition outcomes by development group

Notes: The figure reports event-study coefficients for composition outcomes by development group. Panels (a) and (b) show estimates for the share of postings in high-substitution occupations in advanced and developing economies, respectively. Panels (c) and (d) show estimates for the posting-weighted AI exposure mix. Coefficients are plotted relative to the omitted pre-period. Vertical bars denote 95% confidence intervals.

References

Acemoglu, Daron and David Autor (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings". In: *Handbook of Labor Economics*. Ed. by Orley Ashenfelter and David Card. Vol. 4B. Amsterdam: Elsevier, pp. 1043–1171.

Acemoglu, Daron, David Autor, et al. (2022). "Artificial Intelligence and Jobs: Evidence from Online Vacancies". In: *Journal of Labor Economics* 40.S1, S293–S340.

Acemoglu, Daron and Pascual Restrepo (2019). "Automation and New Tasks: The Implications of the Task Content of Production for Labor Demand". In: *Journal of Economic Perspectives* 33.2, pp. 3–30.

Autor, David H. (2015). "Why Are There Still So Many Jobs? The History and Future of Workplace Automation". In: *Journal of Economic Perspectives* 29.3, pp. 3–30.

Autor, David H., Frank Levy, and Richard J. Murnane (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration". In: *The Quarterly Journal of Economics* 118.4, pp. 1279–1333.

Babina, Tania et al. (2025). "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition". In: *Technology, Productivity, and Economic Growth*. University of Chicago Press, pp. 75–118.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (2025). "Generative AI at Work". In: *The Quarterly Journal of Economics* 140.2, pp. 889–942.

Deming, David J. and Lisa B. Kahn (2018). "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals". In: *Journal of Labor Economics* 36.S1, S337–S369.

Eloundou, Tyna et al. (2024). *GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models*. arXiv preprint. arXiv: 2303.10130.

Felten, Edward W., Manav Raj, and Robert Seamans (2021). "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses". In: *Strategic Management Journal* 42.12, pp. 2195–2217. DOI: 10.1002/smj.3286.

Gmyrek, Paweł, Janine Berg, and David Bescond (2023). *Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality*. ILO Working Paper 96. Geneva: International Labour Office. DOI: 10.54394/FHEM8239.

Hershbein, Brad and Lisa B. Kahn (2018). "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings". In: *American Economic Review* 108.7, pp. 1737–1772.

Hosseini Maasoum, Seyed M. and Guy Lichtinger (2025). *Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data*. Tech. rep. Working paper. SSRN 5425555. SSRN. URL: https://ssrn.com/abstract=5425555.

Ji, Yufei et al. (2023). *Generative AI's Labor-Replacing Impacts on Occupations Also Foster Short-Run Job Opportunities for Early Adopters*. Tech. rep. Working paper. SSRN 4862800. SSRN. URL: https://ssrn.com/abstract=4862800.

Klein Teeselink, Bouke (2025). *Generative AI and Labor Market Outcomes: Evidence from the United Kingdom*. Tech. rep. Working paper. SSRN 5516798. SSRN. URL: https://ssrn.com/abstract=5516798.

Klein Teeselink, Bouke and Daniel Carey (2026). *AI, Automation, and Expertise*. Tech. rep. Working paper. SSRN 6134506. SSRN. URL: https://ssrn.com/abstract=6134506.

Li, Chen (2025). *Artificial Intelligence in Recruitment: Examine the Effects on Hiring and Firm Outcomes*. Tech. rep. Working paper. SSRN 5455354. SSRN. URL: https://ssrn.com/abstract=5455354.

LinkUp (2025). *Job Records*. Version 2025. Dewey Data. DOI: 10.82551/V091-HN53. URL: https://doi.org/10.82551/V091-HN53.

Liu, Jiacheng (2023). *AI Exposure without Labor Data: Measuring AI's Impact through Firms*. Tech. rep. Working paper. SSRN 4668083. SSRN. URL: https://ssrn.com/abstract=4668083.

Liu, Yan, He Wang, and Shu Yu (2025). *Labor Demand in the Age of Generative AI: Early Evidence from the U.S. Job Posting Data*. Tech. rep. Working paper. SSRN 5504741. SSRN. URL: https://ssrn.com/abstract=5504741.

Mäkelä, Elina and Fabian Stephany (2024). *Complement or Substitute? How AI Increases the Demand for Human Skills*. arXiv preprint. arXiv: 2412.19754.

Noy, Shakked and Whitney Zhang (2023). "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence". In: *Science* 381.6654, pp. 187–192.

Pizzinelli, Carlo et al. (2023). *Labor Market Exposure to AI: Cross-Country Differences and Distributional Implications*. IMF Working Paper WP/23/216. International Monetary Fund.

Schubert, Gregor (2025). *Organizational Technology Ladders: Remote Work and Generative AI Adoption*. Tech. rep. Working paper. SSRN 5094265. SSRN. URL: https://ssrn.com/abstract=5094265.

Appendix

Table A1: Top 10 countries by postings, by development group

Rank	Country	Postings	Share (%)
Panel A. Advanced countries			
1	USA	64,075,006	87.29
2	GBR	3,519,135	4.79
3	CAN	2,028,356	2.76
4	AUS	904,675	1.23
5	DEU	691,575	0.94
6	FRA	361,946	0.49
7	SGP	188,081	0.26
8	NLD	183,968	0.25
9	NZL	181,374	0.25
10	IRL	148,286	0.20
Total postings in advanced sample		73,406,864	100.00
Panel B. Developing countries			
1	IND	1,775,821	39.82
2	CHN	532,803	11.95
3	MEX	382,730	8.58
4	POL	214,558	4.81
5	BRA	173,750	3.90
6	PHL	161,424	3.62
7	MYS	147,076	3.30
8	ARE	129,023	2.89
9	ZAF	81,314	1.82
10	ROU	74,932	1.68
Total postings in developing sample		4,459,629	100.00

Notes: Rankings are based on total postings summed within each development group. Shares report the percentage of group-level postings accounted for by each country.

Table A2: Top 5 industries by postings, by development group

Rank	Industry	Postings	Share (%)
Panel A. Advanced countries			
1	Retail Trade	15,199,838	20.71
2	Health Care and Social Assistance	11,732,312	15.98
3	Manufacturing	8,274,850	11.27
4	Accommodation and Food Services	7,315,741	9.97
5	Finance and Insurance	5,371,407	7.32
Total postings in advanced sample		73,406,864	100.00
Panel B. Developing countries			
1	Manufacturing	1,321,780	29.64
2	Finance and Insurance	1,020,552	22.88
3	Professional, Scientific, and Technical Services	648,236	14.54
4	Accommodation and Food Services	523,251	11.73
5	Information	302,374	6.78
Total postings in developing sample		4,459,629	100.00

Notes: Rankings are based on total postings summed within each development group. Shares report the percentage of group-level postings accounted for by each industry.

Table A3: Top 5 ISCO major groups by postings, by development group

Rank	ISCO major group	Postings	Share (%)
Panel A. Advanced countries			
1	Professionals	19,165,047	26.12
2	Technicians and Associate Professionals	16,493,034	22.48
3	Service and Sales Workers	15,898,337	21.67
4	Managers	7,529,033	10.26
5	Clerical Support Workers	4,751,791	6.48
Total postings in advanced sample		73,378,032	100.00
Panel B. Developing countries			
1	Professionals	2,014,146	45.18
2	Managers	936,859	21.01
3	Technicians and Associate Professionals	774,486	17.37
4	Clerical Support Workers	284,114	6.37
5	Service and Sales Workers	216,413	4.85
Total postings in developing sample		4,458,071	100.00

Notes: Rankings are based on total postings summed within each development group. Shares report the percentage of group-level postings accounted for by each ISCO major group.

Table A4: Main baseline estimates, by development group

	(1) Any posting	(2) Postings	(3) High-sub	(4) AI mix	(5) High-skill
Panel A. Advanced countries					
Exposure × post	-0.004	-0.002	-0.009***	-0.011***	0.001
	(0.003)	(0.022)	(0.002)	(0.002)	(0.001)
Observations	510,444	510,405	403,479	403,479	403,479
SE clustering			Firm, country		
Firm-country FE	Yes	Yes	Yes	Yes	Yes
Country×quarter FE	Yes	Yes	Yes	Yes	Yes
NAICS×quarter FE	Yes	Yes	Yes	Yes	Yes
Panel B. Developing countries					
Exposure × post	-0.001	-0.090	-0.007*	-0.012***	-0.001
	(0.004)	(0.092)	(0.004)	(0.003)	(0.002)
Observations	166,860	166,564	117,840	117,840	117,840
SE clustering			Firm, country		
Firm-country FE	Yes	Yes	Yes	Yes	Yes
Country×quarter FE	Yes	Yes	Yes	Yes	Yes
NAICS×quarter FE	Yes	Yes	Yes	Yes	Yes

Notes: Entries report estimates from the main baseline specification, separately for advanced and developing economies. Standard errors clustered by firm and country are reported in parentheses. Columns (1) and (2) are estimated on the full sample. Columns (3), (4), and (5) are conditional on having valid mapped or scored postings, so the number of observations is smaller in those columns. Column (2) is estimated using PPML. Columns (1), (3), (4), and (5) are estimated using linear fixed-effects models. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table A5: Country-clustered robustness, by development group

	(1) Any posting	(2) Postings	(3) High-sub	(4) AI mix	(5) High-skill
Panel A. Advanced countries					
Post period	-0.104***	-0.341***	-0.012***	-0.011**	0.005
	(0.006)	(0.025)	(0.003)	(0.005)	(0.003)
Exposure × post	-0.011***	-0.045***	-0.002**	-0.002*	-0.001
	(0.003)	(0.014)	(0.001)	(0.001)	(0.001)
Observations	510,570	510,570	403,577	403,577	403,577
SE clustering			Country		
Firm-country FE	Yes	Yes	Yes	Yes	Yes
Panel B. Developing countries					
Post period	-0.082***	-0.146	0.012***	0.009***	0.010***
	(0.007)	(0.142)	(0.003)	(0.003)	(0.003)
Exposure × post	-0.003	-0.121	-0.005*	-0.009***	0.000
	(0.004)	(0.092)	(0.003)	(0.003)	(0.002)
Observations	167,256	167,256	118,246	118,246	118,246
SE clustering			Country		
Firm-country FE	Yes	Yes	Yes	Yes	Yes

Notes: Entries report country-clustered robustness estimates by development group. Standard errors clustered by country are reported in parentheses. All specifications include firm-country fixed effects only. Columns (1) and (2) are estimated on the full sample. Columns (3), (4), and (5) are conditional on having positive mapped or scored postings, so the number of observations is smaller in those columns. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table A6: Task-component posting mixes: advanced countries

	(1) Communication	(2) Responsibility	(3) Skills	(4) Non-routine index
Exposure × post	0.010***	0.019***	0.001	0.010***
	(0.002)	(0.003)	(0.003)	(0.002)
Observations	403,479	403,479	403,479	403,479
SE clustering		Firm, country		
Firm-country FE	Yes	Yes	Yes	Yes
Country×quarter FE	Yes	Yes	Yes	Yes
NAICS×quarter FE	Yes	Yes	Yes	Yes

Notes: Entries report task-component posting-mix estimates for advanced countries. Standard errors clustered by firm and country are reported in parentheses. Task-content measures are posting-weighted averages constructed only from occupations with valid task scores. Task scores are standardized before aggregation. The non-routine index is defined as the average of the standardized communication, responsibility, and skills measures. Observations with no scored postings in a firm-country-quarter are excluded rather than set to zero. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table A7: Threshold sensitivity, by development group

	(1) Top tercile	(2) Top quartile	(3) Top quintile	(4) Top decile
Panel A. Advanced countries				
Exposure × post	-0.009***	-0.009***	-0.010***	-0.007***
	(0.002)	(0.002)	(0.002)	(0.001)
Observations	403,479	403,479	403,479	403,479
SE clustering	Firm, country			
Firm-country FE	Yes	Yes	Yes	Yes
Country×quarter FE	Yes	Yes	Yes	Yes
NAICS×quarter FE	Yes	Yes	Yes	Yes
Panel B. Developing countries				
Exposure × post	-0.008***	-0.007*	-0.009**	-0.005
	(0.003)	(0.004)	(0.004)	(0.003)
Observations	117,840	117,840	117,840	117,840
SE clustering	Firm, country			
Firm-country FE	Yes	Yes	Yes	Yes
Country×quarter FE	Yes	Yes	Yes	Yes
NAICS×quarter FE	Yes	Yes	Yes	Yes

Notes: Entries report threshold-sensitivity estimates by development group. The dependent variable is the share of postings above the indicated occupation-level AI-exposure threshold. Standard errors clustered by firm and country are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table A8: High-AI versus low-AI occupations

	(1) PPML occupation postings
Post × high-AI occupation	-0.160***
	(0.011)
Post × high-AI occupation × developing	0.179***
	(0.037)
Observations	2,967,835
SE clustering	Firm, country
Firm-occupation-country FE	Yes
Quarter FE	Yes

Notes: Entries report PPML estimates from an occupation-level specification comparing high-AI and low-AI occupations. High-AI occupations are defined as occupations in the top quartile of the occupation-level AI exposure distribution, while low-AI occupations are defined as occupations in the bottom quartile. Standard errors clustered by firm and country are reported in parentheses. All specifications absorb firm-occupation-country fixed effects and quarter fixed effects. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table A9: Reduced-form estimates by development group (Gmyrek exposure)

	(1) Any posting	(2) Posting	(3) High-sub	(4) AI mix	(5) High-skill
Panel A. Advanced countries					
Post period	-0.106***	-0.359***	-0.013***	-0.012**	0.004
	(0.006)	(0.028)	(0.004)	(0.005)	(0.004)
Gmyrek exposure × post	-0.013***	-0.048***	-0.006***	-0.005	-0.004***
	(0.003)	(0.012)	(0.002)	(0.003)	(0.001)
Observations	510,570	510,570	403,577	403,577	403,577
SE clustering			Firm, country		
Firm-country FE	Yes	Yes	Yes	Yes	Yes
Panel B. Developing countries					
Post period	-0.082***	-0.166	0.015***	0.013***	0.011***
	(0.010)	(0.123)	(0.003)	(0.003)	(0.003)
Gmyrek exposure × post	-0.000	-0.081	-0.023***	-0.031***	-0.008*
	(0.009)	(0.078)	(0.004)	(0.004)	(0.004)
Observations	167,256	167,256	118,246	118,246	118,246
SE clustering			Firm, country		
Firm-country FE	Yes	Yes	Yes	Yes	Yes

Notes: Entries report reduced-form estimates by development group using the firm-country Gmyrek AI exposure index. Standard errors clustered by firm and country are reported in parentheses. All specifications include firm-country fixed effects only. Columns (1) and (2) are estimated on the full sample. Column (2) is estimated using Poisson pseudo-maximum likelihood (PPML). Columns (3), (4), and (5) are conditional on having positive mapped postings, so the number of observations is smaller in those columns. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A10: Placebo timing test: fake shock in 2022Q2

	(1) High-sub	(2) AI mix
Panel A. Advanced countries		
Exposure × placebo post	0.001	-0.001
	(0.002)	(0.001)
Observations	145,198	145,198
Firm-country FE	Yes	Yes
Country×quarter FE	Yes	Yes
NAICS×quarter FE	Yes	Yes
Panel B. Developing countries		
Exposure × placebo post	-0.005	-0.004
	(0.004)	(0.004)
Observations	41,775	41,775
Firm-country FE	Yes	Yes
Country×quarter FE	Yes	Yes
NAICS×quarter FE	Yes	Yes

Notes: Entries report placebo timing estimates using only pre-ChatGPT periods. The fake shock date is 2022Q2, implying a placebo post period beginning in 2022Q3. All specifications absorb firm-country, country-quarter, and NAICS-quarter fixed effects. Standard errors clustered by firm and country are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

Table A11: Correlation Between *AIexp* and Alternative Exposure Measures (ISCO-08)

	Pearson correlation	Spearman correlation
AIexp vs AIOE_lm (Felten et al. 2021)	0.4459	0.4259
AIexp vs Gmyrek exposure	0.6043	0.5391

Notes: Correlations are computed at the ISCO-08 4-digit occupation level for occupations with non-missing values in both measures after merging. Pearson correlations are linear correlations. Spearman correlations are rank correlations (monotonic association).

Table A12: Pizzinelli et al. (2023) Substitution Exposure and Complementarity Measures: Variable Definitions

Variable	**Definition / interpretation**
`isco08`	ISCO-08 occupation code (intended at the 4-digit level).
`isco08lab`	ISCO-08 occupation label (string).
`complementarity_theta`	Occupation-level "potential complementarity" index, denoted θ (higher values indicate greater complementarity with AI, as defined in the source data).
`aioe_lm`	AI Occupational Exposure (AIOE) from Felten, Raj, and Seamans (2021) for *language modeling* applications.
`aioe_img`	AIOE from Felten, Raj, and Seamans (2021) for *image generation* applications.
`aioe_all`	AIOE from Felten, Raj, and Seamans (2021) for *all AI applications*.
`C_AIOE`	Complementarity-adjusted AIOE (*AIexp*). In the source file, this measure is constructed using the original U.S. SOC 2010 classification and then translated to international classifications; as a result, it may not match exactly a mechanically recomputed transformation using ISCO-08, though the two are typically very similar.[a]
`caioe_w50`	*AIexp* variant using weight 0.50 on θ (as provided in the source file).
`median_theta`	Median value of θ reported in the source file (based on the original SOC 2010 construction).
`median_aioe_all`	Median value of `aioe_all` reported in the source file (based on the original SOC 2010 construction).

Notes: This table documents the occupation-level measures used to construct firm–country Substitution Exposure in the paper. Measures are provided by Pizzinelli et al. (2023) and build on Felten, Raj, and Seamans (2021).

[a] Source-file note: *AIexp* was computed using SOC 2010 and then translated to international classifications.

Table A13: Economies in the regression sample by development (Advanced vs. Developing)

Advanced		Developing	
Australia	Austria	Afghanistan	Albania
Belgium	Bermuda	Algeria	Angola
Canada	Cayman Islands	Argentina	Armenia
Croatia	Cyprus	Aruba	Azerbaijan
Czech Republic	Denmark	Bahamas	Bahrain
Estonia	Finland	Bangladesh	Barbados
France	Germany	Belarus	Belize
Gibraltar	Greece	Benin	Bhutan
Guam	Hong Kong	Bolivia	Bosnia and Herzegovina
Iceland	Ireland	Botswana	Brazil
Israel	Italy	Brunei Darussalam	Bulgaria
Japan	Latvia	Burkina Faso	Burundi
Lithuania	Luxembourg	Cambodia	Cameroon
Macau	Malta	Chile	China
Netherlands	New Zealand	Colombia	Costa Rica
Northern Mariana Islands	Norway	Cuba	Democratic Republic of the Congo
Portugal	Singapore	Djibouti	Dominican Republic
Slovakia	Slovenia	Ecuador	Egypt
Spain	Sweden	El Salvador	Equatorial Guinea
Switzerland	Taiwan	Eritrea	Ethiopia
United Kingdom	United States	Fiji	Gabon
		Georgia	Ghana
		Guatemala	Guinea
		Guyana	Haiti
		Honduras	Hungary
		India	Indonesia
		Iran	Iraq
		Jamaica	Jordan
		Kazakhstan	Kenya
		Kuwait	Kyrgyzstan
		Laos	Lebanon
		Lesotho	Liberia
		Liechtenstein	Madagascar
		Malawi	Malaysia
		Maldives	Mauritania
		Mauritius	Mexico
		Moldova	Mongolia
		Montenegro	Morocco
		Mozambique	Myanmar

Advanced	Developing	
	Namibia	Nepal
	Nicaragua	Niger
	Nigeria	Oman
	Pakistan	Panama
	Paraguay	Peru
	Philippines	Poland
	Qatar	Republic of the Congo
	Romania	Russia
	Rwanda	Saudi Arabia
	Senegal	Seychelles
	Sierra Leone	Somalia
	South Africa	South Sudan
	Sri Lanka	Suriname
	Eswatini	Tajikistan
	Tanzania	Thailand
	Trinidad and Tobago	Tunisia
	Turkey	Ukraine
	Uganda	United Arab Emirates
	Uruguay	Uzbekistan
	Venezuela	Vietnam
	Yemen	Zambia
	Zimbabwe	

Notes: This table lists the economies that appear in the final regression sample after the event-window restriction and the requirement that each firm×country cell has at least one pre-period and one post-period observation. We classify economies into advanced and developing groups using a standard country classification broadly aligned with the IMF World Economic Outlook grouping.