

# Across the Atlantic: Early Labor Market Responses Following the Introduction of AI in the United States and the European Union

Nicolò Piergiovanni Bagnoli<sup>a</sup>, Pierre Bouquet<sup>a</sup>, Amin Kaboli<sup>b</sup>, Yossi Sheffi<sup>a</sup>

<sup>a</sup>Massachusetts Institute of Technology (MIT), Center for Transportation & Logistics, 1 Amherst Street, Cambridge, MA, USA; <sup>b</sup>Institute of Mechanical Engineering, Swiss Federal Institute of Technology in Lausanne (EPFL);

## ABSTRACT

We examine early labor market adjustments in the period associated with generative AI (Artificial Intelligence) introduction by comparing employment dynamics in the U.S. (United States) and the E.U. (European Union). Using large-scale workforce data linked to task-based AI exposure measures, we estimate within-firm employment reallocation by seniority and occupation while absorbing firm-level shocks. Across both regions, early-career employment contracts after 2022, with systematically larger relative declines in higher-exposure groups. Once firm-level shocks are controlled for, the magnitude of the declines is similar across the U.S. and E.U. Moving beyond exposure-quintile aggregation to an occupation-level framework reveals substantial heterogeneity that aggregate analyses obscure. Within the same occupation, employment responses vary across seniority levels. Correlation analyses show that exposure-linked contraction in Europe is more closely associated with changes in entry-level classes, while in the United States it is concentrated in large, high-exposure occupational groups.

## KEYWORDS

Generative AI; Large language models; Employment reallocation; Firm-level analysis; Occupation-level dynamics; Cross-Atlantic labor markets

## 1. Introduction

Generative AI has diffused unusually quickly into knowledge work, accelerating long-running debates about whether advanced automation will primarily augment workers or displace them, and on what margins adjustment will occur (Bick, Blandin, and Deming 2024; Hampole et al. 2025). For labor markets, a key near-term challenge is empirical: measured “exposure” to AI capabilities is not itself an outcome. Exposure scores can help create hypotheses about where adjustment pressure should be greatest, but realizing those pressures in employment data depends on adoption pathways, complementary investments, institutions, and firms’ internal labor market design (Eloundou et al. 2024; Manning 2025). We thus focus on connecting exposure measures to observed employment adjustments in a way that separates within-firm reallocation from broad firm-level shocks, and we do so in a cross-Atlantic setting where differences in institutions and regulatory environments plausibly shape diffusion and organizational response.

Classic approaches map occupational task requirements to computerisation risk (Frey and Osborne 2017) and later work links evolving AI capabilities to patents and tasks to produce exposure measures that can change over time (Webb 2019; Felten, Raj, and Seamans 2021). For large language models, Eloundou et al. (2024) provide a widely used rubric that maps O\*NET (Occupational Information Network) tasks to LLM (Large Language Models) feasibility under alternative interface assumptions, and Labaschin et al. (2025) extends this logic to firms by combining occupation-level exposure with workforce composition. Field and experimental studies show substantial productivity gains from AI assistance, often with heterogeneity by experience, consistent with AI affecting the effective task frontier and the relative productivity of junior workers (Brynjolfsson, Li, and Raymond 2025; Noy and Zhang 2023; Dell’Acqua et al. 2023). While these studies try to gauge how AI can change task productivity, they do not identify how firms adjust headcount composition across occupations and career stages at scale. A central contribution on this is given by Brynjolfsson, Chandar, and Chen (2025), documenting relative employment declines for more exposed early-career groups in the U.S. Beyond the United States, evidence is emerging but remains comparatively thin and heterogeneous across contexts (Kauhanen, Maliranta, and Nurmi 2024; Guarascio, Stoellinger et al. 2025; Klein Teeselink 2025). This motivates a cross-Atlantic analysis that is harmonized in data construction, exposure measurement, and identification strategy.

This work targets two research questions. R1 asks how to compare employment adjustments following the introduction of AI between the United States and the European Union under a like-for-like design. To do so, we construct large-firm workforce panels from Revelio Labs microdata and implement a firm-by-exposure-quintile event-study specification in the spirit of Brynjolfsson, Chandar, and Chen (2025). R2 asks how employment reallocation differs across occupations in the period surrounding the release of generative AI tools, and how these dynamics compare between the United States and the European Union. The results reveal substantial heterogeneity in occupational responses on both sides of the Atlantic: Labor adjustment in the AI era is characterized less by uniform employment changes and more by shifts in the internal composition of work, including across seniority groups. These occupation-level patterns provide a direct view of labor reallocation and complement the aggregate and quintile-based evidence by showing where, within bigger groups, employment expands

or contracts.

## 2. Related Literature

When investigating the impact of AI on the workforce, a first branch of research focuses on measuring job exposure to AI through task content. Canonical “computerisation”-style approaches map occupational task requirements into automation risk or susceptibility scores (Frey and Osborne 2017). Subsequent work refines this idea by linking AI capabilities to the content of patents and tasks, yielding measures of exposure that vary across occupations and time as the frontier of AI advances (Webb 2019; Felten, Raj, and Seamans 2021). For LLMs specifically, Eloundou et al. (2024) propose a rubric that matches O\*NET tasks to LLM capabilities and interface assumptions (LLM only versus LLM plus complementary software), providing occupation-level exposure scores that have become a common input to recent empirical designs. Complementary cross-country evidence suggests that measured exposure varies systematically with development stages, reflecting both task composition and skill distributions (Lewandowski, Madoñ, and Park 2025). A last line of work emphasizes that such exposure exercises should be interpreted as forward-looking “automation evaluations” with important limitations, and argues for empirically grounded validation using observed adoption and labor market outcomes (Manning 2024).

A second branch of the literature estimates realized impacts of generative AI in specific workplaces and task environments. Field and quasi-experimental evidence generally find sizable productivity gains from AI assistance, with strong heterogeneity by skill and experience. In a customer support setting, the introduction of a generative AI assistant raises productivity and disproportionately benefits novice workers, consistent with AI acting as a “skill equalizer” within the occupation (Brynjolfsson, Li, and Raymond 2025). Controlled experiments on writing-intensive tasks similarly show large reductions in completion time and improvements in output quality when workers can use ChatGPT (Noy and Zhang 2023). However, evidence also emphasizes that benefits depend on whether tasks lie within the model’s effective capability frontier; when tasks exceed that frontier, AI access can reduce accuracy and degrade performance (Dell’Acqua et al. 2023). Beyond office settings, industrial settings where AI is embedded in larger logistic, robotic, or manufacturing systems are starting to see wide adoption (Boysen and de Koster 2024). These micro settings are informative about mechanisms, but they do not directly identify economy-wide employment effects.

A third and rapidly expanding branch studies early labor market outcomes using high-frequency administrative data, online labor markets, job postings, and representative surveys. Using employee data and designs that absorb firm-level shocks, Brynjolfsson, Chandar, and Chen (2025) document sharp relative employment declines among U.S. early career workers in more exposed occupations, with adjustments occurring primarily through employment rather than compensation. Related work combines employer-level panels with sectoral variation in exposure and finds evidence consistent with wage and employment gains in more exposed sectors during earlier workplace AI tool rollouts, suggesting that the sign of short run effects may depend on adoption timing and whether AI is used for augmentation versus automation (Johnston and Makridis 2025). Job posting evidence points in the same

direction: generative AI reduces demand in automation-prone occupations while increasing demand in augmentation-prone ones, and it reshapes skill requirements within postings (Chen, Srinivasan, and Zakerinia 2025). On online labor platforms, multiple studies show reallocation rather than uniform contraction, with declining demand for some substitutable tasks and rising demand for complementary skills (Demirci 2025; Teutloff et al. 2025). Survey-based measurement of adoption indicates that diffusion has been unusually rapid relative to past general-purpose technologies, which increases the importance of near real-time monitoring of employment outcomes (Bick, Blandin, and Deming 2024; Hartley et al. 2024). In parallel, policy-focused synthesis work argues for combining such monitoring with institutional “adaptive capacity” to manage potentially rapid transition dynamics (Manning 2025).

Evidence outside the United States remains thinner but is emerging quickly. Administrative population data from Finland finds economically small and statistically indistinguishable effects on wages and employment between more and less exposed occupations in the first two years after ChatGPT’s release (Kauhanen, Maliranta, and Nurmi 2024). On the other hand, for the United Kingdom, employment reductions concentrated in junior positions and were driven by curtailed hiring, aligning with the early career channel emphasized in U.S. payroll data (Klein Teeselink 2025). At a broader regional level, studies on European regions document substantial heterogeneity in AI exposure driven by structural factors such as sectoral specialisation, innovation capacity, productivity, and workforce skills. Using a cluster-based typology of regions, this work argues that areas with strong innovation systems may experience employment gains through complementarity with existing production structures, while peripheral regions face structural constraints that could, over time, widen regional disparities in the European Union (Guarascio, Stoellinger et al. 2025).

Evidence from prior technological transitions motivates several complementary empirical needs that this thesis targets. First, the computer and internet era literature shows that technology-driven restructuring is not well captured by aggregate firm or economy-wide trends alone (Machin and Van Reenen 1998; Michaels, Natraj, and Van Reenen 2014). Second, comparative work on the ICT (Information and Communication Tech) productivity era highlights that the United States and Europe can experience meaningfully different propagation of the same general purpose technology because of differences in sectoral composition, competitive environment, and complementary intangible investments, implying that cross-Atlantic contrasts are central for interpretation rather than a robustness afterthought (van Ark, O’Mahony, and Timmer 2008). Finally, historical accounts of the Industrial Revolution emphasize that the timing of productivity gains versus wage and employment adjustment can be asynchronous, with potential “pauses” in broad-based worker gains even as output rises, motivating the need to understand different time-based implications across regions (Allen 2009). By leveraging large-scale worker and firm microdata and harmonized exposure measures, we provide a unified U.S.-E.U. view of employment reallocation dynamics by career stage, with a design that is explicitly built to detect the within-firm composition shifts and cross-region differences that have characterized past technological revolutions.

### 3. Data

This section describes the three primary data components used in the study. First, we introduce O\*NET as the task and occupation backbone used throughout the paper. Second, we describe the exposure-score layers mapped to O\*NET tasks and occupations, including a share-weighted extension developed in this work. Third, we document the Revelio Labs data used to construct the firm-by-month employment panels for the U.S. and European Union samples.

#### 3.1. *O\*NET Occupational Data*

The Occupational Information Network is the leading repository of standardized occupational descriptors in the United States. It serves as a foundational taxonomy for empirical work on task composition, skills, and technology exposure (National Center for O\*NET Development 2024). O\*NET contains on the order of tens of thousands of job titles consolidated into about one thousand detailed occupational profiles. Each occupational profile decomposes work into hierarchical elements across multiple levels of granularity. At the finest task level, O\*NET provides a large set of task statements describing specific activities workers perform, which can be aggregated into higher-level groupings such as Detailed Work Activities (DWAs). A significant input for the subsequent methodology is the "task ratings" component of O\*NET, which provides survey-based ratings used as a proxy to infer how work share is distributed across tasks within occupations. These ratings are used in Subsection 3.2.3 to construct task-share allocations. Additional details on O\*NET survey protocols, table extracts used, and the mapping logic employed in this paper are documented in Appendix A.

#### 3.2. *Exposure Scores*

This section outlines the construction of occupation-level exposure measures, moving from rubric-based task exposure scores to sentiment and share-weighted extensions that capture variation in both task importance and exposure over time.

##### 3.2.1. *Rubric-Based LLM Exposure Scores*

The baseline LLM exposure measures used in the recent labor literature are by Eloundou et al. (2024). Their core contribution is an expert rubric that labels O\*NET tasks according to whether a large language model can substantially reduce the time required to complete the task at a fixed quality threshold. The rubric further distinguishes between exposure achievable through standalone chat-based use and exposure requiring complementary software integration. Task-level labels are aggregated to occupation-level exposure measures by averaging across tasks within an occupation. A key implicit assumption in this aggregation is that tasks within an occupation are equally important, which can be restrictive given substantial heterogeneity in time allocation across tasks. To address this, we construct an exposure metric that is weighted by task-share using allocation weights  $\pi_{t,o}$  inferred from O\*NET task frequency data following the procedure in Bouquet, Bagnoli, and Sheffi (2025). O\*NET characterizes task execution through a two-stage survey in which incumbents first report task relevance and, conditional on relevance, report execution frequency on an ordinal scale. We convert these frequencies into expected

annual occurrence counts using a fixed annualization mapping and normalize within occupations to obtain weights  $\pi_{t,o}$  that sum to one.

Each task  $t$  is assigned an exposure category. Let  $E0$  denote no exposure,  $E1$  denote direct exposure (LLM via chat reduces task time by at least 50 percent at the quality threshold), and  $E2$  denote LLM plus exposure (additional LLM powered software could reduce task time by at least 50 percent). Following Eloundou et al. (2024), we construct three task-level measures, which bracket exposure:

$$\alpha_t = \begin{cases} 1 & \text{if } t \in E1 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

$$\beta_t = \begin{cases} 1 & \text{if } t \in E1 \\ 0.5 & \text{if } t \in E2 \\ 0 & \text{if } t \in E0 \end{cases} \quad (2)$$

$$\gamma_t = \begin{cases} 1 & \text{if } t \in E1 \text{ or } t \in E2 \\ 0 & \text{if } t \in E0 \end{cases} \quad (3)$$

Here,  $\alpha$  is a lower bound (direct exposure only),  $\beta$  is a mid-range measure (partial weight on LLM plus), and  $\gamma$  is an upper bound (full weight on both  $E1$  and  $E2$ ).

Let  $\mathcal{T}(o)$  be the set of tasks in occupation  $o$ . The original approach aggregates by an unweighted mean across tasks:

$$\alpha_o = \frac{1}{|\mathcal{T}(o)|} \sum_{t \in \mathcal{T}(o)} \alpha_t, \quad \beta_o = \frac{1}{|\mathcal{T}(o)|} \sum_{t \in \mathcal{T}(o)} \beta_t, \quad \gamma_o = \frac{1}{|\mathcal{T}(o)|} \sum_{t \in \mathcal{T}(o)} \gamma_t. \quad (4)$$

We instead weight each task by the share of the workday it occupies, computed as in Bouquet, Bagnoli, and Sheffi (2025). Let  $\pi_{t,o} \in [0, 1]$  denote the share of total work in occupation  $o$  allocated to task  $t$ , with  $\sum_{t \in \mathcal{T}(o)} \pi_{t,o} = 1$ . Our task-share rubric exposure measures are:

$$\alpha_o^{\text{TW}} = \sum_{t \in \mathcal{T}(o)} \pi_{t,o} \alpha_t, \quad \beta_o^{\text{TW}} = \sum_{t \in \mathcal{T}(o)} \pi_{t,o} \beta_t, \quad \gamma_o^{\text{TW}} = \sum_{t \in \mathcal{T}(o)} \pi_{t,o} \gamma_t. \quad (5)$$

Intuitively,  $\alpha_o^{\text{TW}}$  represents the share-weighted exposure of occupation  $o$ . Rather than treating all tasks within an occupation as equally important,  $\alpha_o^{\text{TW}}$  aggregates task-level exposure scores  $\alpha_t$  using the share of the workday  $\pi_{t,o}$  that workers in occupation  $o$  spend on each task  $t$ . Figure 1 compares the unweighted rubric scores to their share-weighted counterparts. The correlations are high (0.930 for  $\alpha$ , 0.971 for  $\beta$ , 0.980 for  $\gamma$ ), indicating that introducing  $\pi_{t,o}$  substantially changes which tasks drive exposure within a given occupation, but does not, in aggregate, overturn the overall occupation level exposure patterns or ranking.




Figure 1. Unweighted rubric based exposure scores from Eloundou et al. (2024) versus share-weighted rubric scores constructed using O\*NET task-share allocation weights  $\pi_{t,o}$  (Bouquet, Bagnoli, and Sheffi 2025). Each point is an occupation.

Task-share weighting addresses the within-occupation allocation problem, but rubric-based measures remain a largely static snapshot of technical feasibility at a given model vintage, locked at a specific point in time. Rubric labels are discrete and expert-assigned, which can miss gradual shifts in capability, heterogeneity in real-world implementation, and changes in demand side emphasis that are visible in higher frequency text, such as news articles. For these reasons, we decided to maintain the share-dependent task approach, but look elsewhere for the exposure metric.

### 3.2.2. Sentiment-Based Exposure Scores

To address the static nature of rubric-based measures, we consider sentiment-based exposure scores derived from the tone of news coverage about tasks and jobs. The key idea is to treat news sentiment regarding AI-enabled substitutions or augmentations as a dynamic proxy for evolving automation pressure, enabling historical trend analysis and near-real-time monitoring (Bouquet, Sheffi, and Kaboli 2026).

### 3.2.3. Share-Weighted Sentiment Exposure (Novel Approach)

Finally, we construct a sentiment-share exposure measure that re-weights task exposure by inferred shares, rather than treating each task within an occupation as equally important. Let an occupation  $o$  consist of tasks  $t \in \mathcal{T}(o)$ . We define the share-weighted occupation exposure as:

$$X_o^{\text{TW}} = \sum_{t \in \mathcal{T}(o)} \pi_{t,o} x_{t,o}. \quad (6)$$

Let  $x_{t,o}$  denote a task-level sentiment exposure metric (tasks are occupation-specific), and  $\pi_{t,o} \in [0, 1]$  denote the share of total work in occupation  $o$  allocated to task  $t$ , with  $\sum_{t \in \mathcal{T}(o)} \pi_{t,o} = 1$ . Empirically, we successfully construct share-weighted sentiment exposure scores for 796 job titles. As a validation check, the resulting share-weighted sentiment exposure is positively correlated with the benchmark rubric-based  $\beta$  exposure measure from Eloundou et al. (2024), with a Spearman correlation of 0.576, indicating meaningful agreement in broad exposure patterns while allowing sentiment

and share allocation to introduce additional variation in exposure intensity. As a descriptive summary of the resulting job-level distribution, Appendix B reports basic statistics and the ten most and least exposed job titles.

### 3.3. *Revelio Labs Workforce Microdata*

Revelio Labs compiles and harmonizes online résumé data and related sources to produce measures of employment stocks and flows across firms, geographies, and occupations. This data source has been used in related work linking LLM exposure measures to firm-level employment structures, which provides external support for the suitability of Revelio as a backbone for firm-by-occupation measurement (Labaschin et al. 2025).

#### 3.3.1. *Firm Universe*

The empirical analysis focuses on large firms, motivated by two considerations. First, large firms provide more stable month-to-month occupation counts, which is important for the fixed-effect specifications used later in the paper. Second, downstream estimation requires filtering on minimum occupation-by-firm cell sizes to ensure numerical stability and convergence, implying that an initial restriction to larger firms is coherent with the final estimation sample. Operationally, we start from a Dow Jones-based firm list (30 constituents) as a small, comprehensive (firms from most industries), and well-understood universe that enables rapid validation of the data pipeline (information on the Dow dataset is available in Appendix C). We then extend the logic to the comprehensive U.S. and E.U. datasets (Appendix D).

#### 3.3.2. *Geographic Scope, Timeframe, and Seniority Definition*

Figure 2 displays the geographic distribution of observed workforce counts in the constructed sample across the United States and separate E.U. countries.




Figure 2. Geographic distribution of the workforce in the constructed sample across the United States and individual E.U. countries.

The observation window spans from January 2021 through October 2025, providing a sufficiently long pre-period to characterize baseline labor demand and workforce composition, and a multi-year post-period to capture medium-run adjustment. Following the empirical timing used in the recent literature on generative AI and labor

market outcomes (Brynjolfsson, Chandar, and Chen 2025), we treat October 2022 (ChatGPT was released to the public in November 2022) as the last pre-shock month and use the subsequent period to examine whether differential dynamics by exposure emerge once frontier LLM capabilities begin to diffuse. Operationally, we therefore index outcomes to October 2022 in descriptive plots and interpret deviations thereafter as post-shock dynamics, while acknowledging that adoption timing is heterogeneous across firms and occupations and that any realized effects may unfold gradually. We stratify workers by seniority using Revelio Labs' model-based seniority measure, which maps a continuous score into seven ordinal categories that we re-label and aggregate into six seniority groups for use throughout the analysis.

Table 1. Mapping between Revelio Labs seniority levels and analysis seniority groups

Seniority	Analysis seniority group	Revelio Labs seniority level(s)
1	<i>Early Career 1</i>	Entry Level
2	<i>Early Career 2</i>	Junior Level
3	<i>Developing</i>	Associate Level
4	<i>Mid Career 1</i>	Manager Level
5	<i>Mid Career 2</i>	Director Level
6	<i>Senior</i>	Executive Level; Senior Executive Level

### 3.3.3. U.S. and E.U. Sample Construction

We construct the final E.U. and U.S. datasets from Revelio in three steps. First, we define a large-firm universe by selecting firms above a minimum scale threshold: we compute total employment by firm for a reference month and retain firms with more than 500 positions (a position is different from a unique employee, where a single employee can occupy more positions in a company, at different times). Second, within this large-firm universe, we measure each firm's geographic footprint by checking whether it exceeds 300 positions in the E.U., the U.S., or both regions. Third, we build two region-specific position panels by appending individual-level position records from qualifying firms. The dataset includes occupation codes, inferred seniority, job start and end dates, and individual-level attributes. The E.U. dataset includes all position records whose country is in the E.U. list for firms with more than 300 employees in the E.U. In comparison, the U.S. dataset includes all position records whose country is the United States for firms with more than 300 positions in the U.S. Importantly, firms that exceed the 300-position threshold in both regions contribute records to both datasets: the firm therefore appears in both geographies, and its position records are split by worker location, reflecting multinational employment structures.

This procedure yields the following large-company panels between 2015 and 2025 (we take a wider time-frame than the observation window to study workforce dynamics before filtering to January 2021 - October 2025). The E.U. dataset contains 21,369,981 position records, 12,808,411 unique employees, and 2,150 companies. The U.S. dataset contains 47,645,193 position records, 28,235,297 unique employees, and 3,907 companies. More information on the regional datasets and company descriptive statistics can be found in Appendix D. By construction, these panels describe employment dynamics within large firms rather than the full workforce in each region. For context, U.S. total employment is on the order of 161 million workers (U.S. Bureau of

Labor Statistics 2025), while E.U. employment is close to 200 million in recent quarterly labor market statistics (Eurostat 2025). The large-firm focus can be particularly restrictive in economies with a substantial small and medium enterprise sector, such as Italy, where SMEs (Small and Medium Enterprises) are an important object of official statistical measurement and a meaningful share of economic activity (ISTAT 2025). The implications of this sampling frame for external validity are discussed in the discussion section.

## 4. Methods

This section presents methods used to measure post-2022 changes in employment composition and relate them to independently constructed AI exposure measures. We implement two complementary designs. We estimate a firm by exposure quintile event study that follows the core specification in Brynjolfsson, Chandar, and Chen (2025). This design asks whether, within the same firm and month, employment in more exposed occupational groups evolves differently from employment in less exposed groups after the onset of the generative AI era. We then extend our study to an occupation-based analysis. This design produces a job-specific post-2022 reallocation measure that does not use exposure scores in the regression itself. Exposure measures are used only for benchmarking in a second step, which helps separate the identification of employment reallocation from assumptions embedded in any specific exposure metric.

### 4.1. *Exposure Quintiles*

Let  $x_o$  denote an occupation level AI exposure score for occupation  $o$ . We use the sensitivity task-share scores to construct exposure quintiles. Operationally, we assign each occupation to  $q(o) \in \{1, 2, 3, 4, 5\}$ , where  $q(o) = 5$  indicates the most exposed quintile and  $q(o) = 1$  the least exposed quintile. We then map these pre-defined quintile assignments to the full E.U./U.S. datasets via the O\*NET SOC (Standard Occupation Classification) code. The Revelio dataset already provides for each worker the SOC code for both the E.U./U.S. datasets. All worker position records with the same O\*NET code inherit the same quintile label. Because quintile boundaries are defined at the occupation type level, each quintile contains approximately one fifth of occupations, but the resulting distribution of employees across quintiles is generally not uniform, since employment is concentrated in some occupations and sparse in others.

### 4.2. *Firm by Quintile Event Study with Firm-Time Effects*

For each firm  $f$ , exposure quintile  $q$ , and calendar month  $t$ , define the outcome  $y_{f,q,t}$  as the number of employees in firm  $f$  at time  $t$  whose occupation maps to exposure quintile  $q$ . The specification models the conditional mean of this count using a Poisson log link, which accommodates zero outcomes without taking logs of realized counts (Chen and Roth 2024). Following Brynjolfsson, Chandar, and Chen (2025), we estimate a Poisson event-study regression. Let  $t$  index calendar months and let  $\tau \equiv t - t_0$  denote event time in months relative to the reference date  $t_0$ , defined as the final pre-period month (October 2022). Thus,  $\tau = 0$  corresponds to the first post-period month and

$\tau = -1$  is the omitted reference period. The estimating equation is:

$$\log \mathbb{E}[y_{f,q,t} | X] = \alpha_{f,q} + \beta_{f,t} + \sum_{q' \neq 1} \sum_{\tau \neq -1} \gamma_{q',\tau} \cdot 1\{q = q'\} 1\{t - t_0 = \tau\}. \quad (7)$$

The components of equation (7) have a direct interpretation.

- (1)  $\alpha_{f,q}$  are firm-by-quintile fixed effects. They absorb time-invariant differences in how a firm staffs different exposure quintiles, such as persistently higher employment in low-exposure roles for some firms and persistently higher employment in high-exposure roles for others.
- (2)  $\beta_{f,t}$  are firm-by-month fixed effects. They absorb all firm-wide shocks at each calendar date, including demand fluctuations, reorganizations, and macroeconomic conditions that affect employment across all exposure quintiles within the firm in that month.
- (3)  $\gamma_{q,\tau}$  are the coefficients of interest. They measure differential changes in expected employment for exposure quintile  $q$  at event time  $\tau$ , relative to the omitted reference group, which is exposure quintile 1 in the reference month  $\tau = -1$ .

Because equation (7) includes both firm-by-quintile and firm-by-month fixed effects, identification comes entirely from within-firm, within-month variation. Intuitively, the coefficients compare how employment in more exposed quintiles evolves relative to less exposed quintiles inside the same firm at the same point in time, before and after the reference date. This structure removes confounding from firm-level shocks and isolates differential employment dynamics across exposure groups within firms. At the same time, the coefficients should not be interpreted mechanically as causal effects of generative AI. A causal interpretation would require a parallel-trends-type condition: absent the post-2022 technology shock, employment in different exposure quintiles within the same firm would have evolved similarly after conditioning on  $\alpha_{f,q}$  and  $\beta_{f,t}$ . We therefore interpret  $\gamma_{q,\tau}$  as differential employment dynamics across exposure groups in a post-AI-introduction era. Equation (7) is estimated by Poisson pseudo-maximum likelihood (PPML). PPML is consistent under correct specification of the conditional mean even when the conditional variance departs from the Poisson assumption, and it is well suited to settings with multiplicative structure and heteroskedasticity (Santos Silva and Tenreyro 2006). Standard errors are clustered at the firm level. The PPML regression sample is constructed separately for each seniority group and includes only firms that satisfy three conditions: (i) a minimum of 10 positions in every month of the analysis window, (ii) at least 100 cumulative positions in each AI exposure quintile over the full period, and (iii) complete support across all five exposure quintiles. These restrictions ensure stable identification of firm-month and firm-quintile fixed effects and prevent estimates from being driven by firms with sparse or highly unbalanced exposure composition.

#### 4.3. Occupation Level Within Firm Reallocation With an Offset

To study within-firm reallocation at a more granular level, we estimate a second specification at the occupation level that conditions on total firm employment. Let  $y_{f,o,t}$  denote employment in occupation  $o$  at firm  $f$  in month  $t$ , and let total firm employment

be  $Y_{f,t} = \sum_o y_{f,o,t}$ . We estimate the following PPML model:

$$\log \mathbb{E}[y_{f,o,t} | X] = \log Y_{f,t} + \alpha_{f,o} + \sum_{k \neq -1} \delta_{o,k} \cdot 1\{t \in k\}. \quad (8)$$

where  $\log Y_{f,t}$  enters as an offset with a coefficient fixed at one, and  $\alpha_{f,o}$  are firm-by-occupation fixed effects. Each component of equation (8) captures a distinct object. The offset  $\log Y_{f,t}$  conditions on total firm employment in a given month, so that identification comes from how a firm allocates a fixed employment mass across occupations rather than from overall firm expansion or contraction. The fixed effects  $\alpha_{f,o}$  absorb time-invariant differences in how firms staff particular occupations. Finally, the event-time coefficients  $\delta_{o,k}$  capture changes in the within-firm employment share of occupation  $o$  at event-time bin  $k$ , relative to the omitted pre-period reference bin.

This specification is designed to isolate changes in occupational composition within firms after 2022, while remaining agnostic about the correct measure of AI exposure in the first-stage regression. Because the model conditions on total firm employment but does not include firm-by-month fixed effects, it should be interpreted as controlling for firm-level scale rather than fully absorbing all firm-time shocks that may affect occupational composition. We therefore view the resulting  $\delta_{o,k}$  as descriptive measures of within-firm reallocation patterns over time, which can subsequently be related to independently constructed exposure measures. For this analysis, we restrict the sample to (i) firms with at least 10 positions in every month of the analysis window, (ii) occupation-seniority pairs with at least 10 total positions across all firms in every month, and (iii) occupation-seniority pairs with at least 100 firm-month observations available for estimation. These restrictions ensure stable estimation and limit the influence of sparsely populated occupation cells.

## 5. Results

We start by comparing the alternative exposure metrics within the Dow Jones firm sample, documenting how rubric, rubric-share, and sentiment-share measures align across seniority levels. After this, we scale the analysis to a cross-region comparison between the United States and the European Union, grouping the workforce into exposure quintiles to study how employment dynamics differ across the exposure distribution. Finally, we move down to the occupation level to identify which job categories account for the aggregate patterns, and to extract the most granular insights on where exposure is translating into differential labor market adjustments.

### 5.1. Share Weighted Exposure Metrics

To characterize how LLM shock translates into occupational labor market adjustment, we investigate two complementary share-weighted exposure measures. Figure 3 illustrates the resulting employment dynamics in the DOW for the *Early Career 2* (EC2) group, which we highlight because it combines large employment mass with meaningful heterogeneity in exposure. The full seniority panel set is reported in Appendix E. In EC2, the share-weighted sentiment score generates an exposure

ordering that is organically aligned with what one would expect: higher exposure quintiles display systematically different post-shock trajectories relative to lower exposure quintiles.




Figure 3. EC2 employment dynamics by exposure quintile under three exposure constructions. Headcount is normalized to October 2022. The vertical marker denotes the ChatGPT release window.

Across the three exposure constructions, employment declines for Early Career 2 workers are broadly similar in magnitude, but differ in how clearly they stratify across exposure quintiles. In particular, the share-sentiment measure yields a monotonic ordering of employment shifts across all five quintiles (Table 2), with progressively larger declines at higher exposure levels. The patterns suggest that combining share-weighting with sentiment information produces exposure groupings whose relative dynamics align more closely with the intended exposure ranking. We therefore use the share-sentiment exposure as the primary metric in the comparative U.S.-E.U. analysis that follows.

Table 2. Early Career 2 employment shifts by AI exposure quintile

Exposure measure	Exposure Quintile				
	Q1	Q2	Q3	Q4	Q5
<i>Panel A: Last-period shift (Oct 2025)</i>					
Eloundou (unweighted)	-0.0537	-0.0686	-0.0654	-0.0907	-0.0923
$\pi$ -weighted rating	-0.0694	-0.0920	-0.0983	-0.0936	-0.1025
$\pi$ -weighted sentiment	-0.0192	-0.0725	-0.0916	-0.1174	-0.1238
<i>Panel B: Average post-baseline shift (Oct 2022–Oct 2025)</i>					
Eloundou (unweighted)	-0.0262	-0.0209	-0.0360	-0.0489	-0.0554
$\pi$ -weighted rating	-0.0349	-0.0330	-0.0595	-0.0498	-0.0623
$\pi$ -weighted sentiment	-0.0127	-0.0397	-0.0579	-0.0598	-0.0684

## 5.2. E.U. vs U.S. - Quintile-Based

### 5.2.1. Sample construction and coverage

After restricting the raw position records to the common analysis window (January 2021 to October 2025) and applying firm cleaning described in Section 4.2, the U.S. dataset contains 21,318,095 position records. This sample already reflects a non-trivial coverage constraint on exposure measurement: approximately 22% of SOC titles are excluded because they lack sufficient task and occupation metadata to compute the exposure score used in this work. For the E.U. panel, the final sample comprises 9,462,370 position records.

### 5.2.2. Raw Employment Dynamics

Figure F1 and Figure F2 report raw employment dynamics by seniority group and AI exposure quintile, shown separately for the European Union and the United States. In both regions, the sharpest post-period contraction occurs among early-career workers. Employment for *Early Career 1* declines steadily after late 2022, reaching cumulative losses of 8.1 percent in the E.U. and 8.6 percent in the United States by October 2025 relative to the October 2022 baseline. A similar pattern emerges for *Early Career 2*, with employment falling by approximately 6.4 percent in the E.U. and 7.2 percent in the United States over the same period. In this group, the contraction is clearly stratified by exposure, with higher AI exposure quintiles exhibiting faster and more persistent declines throughout most of the post period. In contrast, *Developing* workers experience pronounced expansion in both regions. After a steady pre-period increase, employment accelerates following October 2022, resulting in cumulative growth of approximately 17.2 percent in the E.U. and 10.3 percent in the United States by the end of the sample. The remaining seniority groups display more muted dynamics. Employment for *Mid Career 1* and *Mid Career 2* rises modestly through 2023 before flattening or slightly declining thereafter. For *Senior* workers, raw employment growth is more pronounced, though less tightly linked to exposure than in earlier career stages.

### 5.2.3. Regression-Adjusted Effects

Figure 4 and Figure 5 compare, in the E.U. and the U.S., respectively, raw normalized employment differences by quintile and (ii) regression-adjusted quintile coefficients with 95% confidence intervals, for each seniority group. Tables reporting the underlying numeric estimates are provided in Appendix F.

A consistent pattern emerges in both regions for *Early Career 1* and *Early Career 2*: higher exposure groups experience a statistically meaningful relative decline versus Quintile 1 by the end of the sample. In the E.U., for EC2, the last period (October 2025) regression adjusted effects are approximately -5.01% for Quintile 4 and -4.04% for Quintile 5 (both statistically significant). In the U.S., for the same seniority, the corresponding last period effects are approximately -5.06% for Quintile 4 and -5.28% for Quintile 5 (both statistically significant). For *Early Career 1*, both regions show smaller but still negative end-of-sample effects for the most exposed groups, with statistical significance concentrated in Quintiles 4 and 5. In the E.U., the last period effects are approximately -2.80% (Quintile 4) and -3.20% (Quintile 5), both statistically significant. In the U.S., the last period effects are approximately -2.78% (Quintile 4) and -3.12% (Quintile 5), both statistically significant.

For *Developing*, raw series show large positive growth that is strongly stratified by exposure, but regression-adjusted effects are substantially smaller and often change sign. This is one of the most interesting findings we observe, and a motivation to extend to the occupation level analysis to gather more understanding of this case. In the E.U., only Quintile 3 exhibits a positive and statistically significant end-of-sample effect (approximately +3.88%), while Quintile 4 is negative and statistically significant at the end of the sample (approximately -2.94%). In the U.S., Quintiles 2 and 3 are positive and statistically significant at the end of the sample (approximately +1.86% and +8.44%), while Quintiles 4 and 5 are modestly negative but not statistically distinguishable from zero at conventional levels.

For *Mid Career* groups, the E.U. and U.S. diverge more in inference, but share a common theme: relative effects are generally modest and sensitive to firm controls. In the E.U., *Mid Career 1* shows positive and statistically significant end-of-sample effects for Quintiles 3 and 4 (approximately +5.55% and +3.16%), while *Mid Career 2* shows statistically significant end-of-sample effects across Quintiles 3 to 5 (approximately +4.26%, +2.76%, +3.95%). In the U.S., *Mid Career 1* shows positive and statistically significant end-of-sample effects for Quintiles 2 to 4, while *Mid Career 2* effects are close to zero and not statistically significant at the end of the sample. Finally, for *Senior*, both regions show a pronounced gap between raw trends and regression-adjusted effects. Raw series suggest strong relative growth for highly exposed senior roles, but after controlling for firm shocks, the estimated effects are small and not statistically different from zero at the end of the sample.




Figure 4. E.U.: raw normalized differences versus regression-adjusted PPML effects by seniority and exposure quintile.




Figure 5. U.S.: raw normalized differences versus regression-adjusted PPML effects by seniority and exposure quintile.

#### 5.2.4. Robustness Checks

We assess the robustness of our empirical strategy by re-estimating the U.S.-E.U. specifications on three systematically restricted subsamples: gender, ethnicity, and technology-oriented occupations. The objective is to verify whether the qualitative patterns of AI-exposure-related employment dynamics are sensitive to workforce composition or task environment. Across all subsamples, the estimated dynamics remain highly consistent with the baseline results. As illustrated in Appendix G, splitting the sample by gender (male vs. female, separately for the E.U. and the U.S.) yields virtually identical trajectories, with wider confidence intervals reflecting reduced sample size. Analogous stability is observed for all subsets.

### 5.3. E.U. vs U.S. – Occupation Based

For the occupation based analysis, we proceed by working with two subsets of the main datasets described earlier. First, we analyze a sample of 800 firms randomly drawn in the European Union and the United States. Second, we focus on a concentrated sample consisting of the 300 largest firms by baseline employment in each region. This allows us to distinguish patterns that are representative of the broader firm population from those that are driven by large, systemically important employers. By conditioning on narrowly defined O\*NET occupations, the occupation-based framework allows us to study heterogeneous employment adjustments within a single job category, abstracting from compositional shifts across occupations. In particular, it enables us to trace how employment responses in the age of AI differ across seniority levels and regions within the same occupation. To illustrate the value of this approach, we focus on the O\*NET occupation *Customer Service Representatives* (43-4051.00), which is especially well suited for within-occupation analysis given its large employment base, clear seniority structure, and exposure to large language model capabilities.

The event-study results for the randomly sampled set of 800 firms are broadly similar across regions (visible in Figure H3 and Figure H4). In both the European Union and the United States, employment responses for *Early Career 1* and *Early Career 2* follow comparable post-period trajectories. However, a systematic regional divergence emerges for the *Developing* seniority group. Across both the 800-firm panel and the top-300 firms by size, the European Union exhibits statistically significant post-period employment growth within *Customer Service Representatives*. In contrast, the corresponding estimates for the United States are not statistically significant in either panel. While the raw dynamics in the U.S. suggest weaker or declining trends, the estimated effects do not support a clear post-period expansion. As a result, the positive developing-level adjustment appears specific to the European Union and robust to alternative firm-size samples.

A complementary and compact way to summarize these patterns is to examine the correlation between occupation-specific employment share changes and AI exposure scores. Figure 6 reports Pearson correlations between occupation–seniority employment share changes and sentiment-share AI exposure scores, shown separately for the United States and the European Union. For each seniority level, we report both an unweighted correlation, where each occupation–seniority cell receives equal weight, and an employment-weighted correlation, where cells are weighted by the number of workers they represent.




Figure 6. The x-axis in each chart is seniority level from 1 to 6, and the y-axis is Pearson correlation  $r$ . For each seniority level, there are two bars: an unweighted correlation and an employment-weighted correlation, shown with diagonal hatching. 95% confidence intervals are described in Appendix I. The panel comprises 800 randomly sampled companies.

Focusing first on the U.S. panel based on 800 randomly sampled firms in Figure 6, the unweighted correlations are generally small and statistically indistinguishable from zero in both regions. A similar pattern holds in the European Union, where unweighted correlations are modest at low and mid seniority levels and increase monotonically with seniority, reaching roughly 0.19 at seniority level 6. Taken in isolation, these unweighted statistics suggest limited systematic alignment between AI exposure and employment share changes when each occupation-seniority cell is treated symmetrically. Once correlations are weighted by employment, however, a much clearer pattern emerges. In both regions, employment-weighted correlations at the lowest seniority levels are larger in magnitude and statistically significant. In the European Union, the employment-weighted correlation is approximately  $-0.22$  at seniority level 1 and  $-0.18$  at seniority level 2, with confidence intervals that exclude zero. In the United States, the corresponding correlations are smaller in absolute value but remain negative at early seniority levels, around  $-0.09$  at seniority 1 and  $-0.08$  at seniority 2. These results indicate that exposure-linked employment contractions are concentrated in occupation-seniority segments that account for a large share of total employment. A negative Pearson correlation implies that occupations with higher AI exposure scores tend to experience larger declines in employment shares.

Beyond early-career roles, the European Union exhibits a distinct and systematic pattern. Employment-weighted correlations turn positive and rise steadily with seniority, becoming statistically significant at the highest seniority level. This indicates that, in the E.U., higher AI-exposed occupations at more advanced seniority levels are associated with relative employment growth. The United States does not display a comparable pattern in the randomly sampled firm panel: employment-weighted correlations at mid and higher seniority levels remain close to zero and are not statistically distinguishable from zero. Notably, the European pattern observed in the 800-firm

sample closely mirrors the dynamics found among the top 300 firms by size, whereas in the United States the large-firm sample exhibits negative correlations at higher seniority levels, indicating exposure-linked employment contraction even among more senior roles (Figure J1). These results support a “canaries in the coal mine” interpretation for early-career workers in both regions, while highlighting a distinct European adjustment pattern in which higher-seniority, high-exposure roles appear comparatively insulated or even expanding. In contrast, workforce adjustments in large U.S. firms appear to operate along the exposure margin across a broader segment of the seniority distribution.

## 6. Discussion

In setting out to contribute empirical evidence on how labor demand is adjusting in the early diffusion phase of generative AI, two observations stand out. First, the high-level picture is remarkably aligned across regions once we net out firm-level shocks, particularly for early career groups. Second, meaningful heterogeneity emerges when we move from exposure quintiles to occupation-specific dynamics, suggesting that aggregation can conceal economically relevant reallocations that matter for practitioners and policymakers.

### 6.1. *Comparing the United States and Europe Under a Harmonized Design.*

A central objective was to provide a like-for-like comparison across the United States and the European Union using the same data source family, harmonized exposure measures, and a specification that absorbs contemporaneous firm-level shocks. The raw dynamics already reveal a common pattern: early-career employment contracts after late 2022 in both regions, with declines of approximately 8.1 percent in the E.U. and 8.6 percent in the United States for *Early Career 1*, and about 6.4 percent in the E.U. and 7.2 percent in the United States for *Early Career 2*. In contrast, *Developing* roles expand in both regions over the same window, with larger raw growth in the E.U. (approximately 17.2 percent) than in the U.S. (approximately 10.3 percent). The regression-adjusted results show that, for both *Early Career 1* and *Early Career 2*, higher exposure groups experience statistically meaningful relative declines versus the low exposure reference. For *Early Career 2*, the end of sample effects for high exposure quintiles are tightly similar across regions, on the order of roughly minus 4 to minus 5 percent relative to Quintile 1 by October 2025. For *Early Career 1*, the end of sample gaps are smaller but still negative and statistically concentrated among the most exposed groups. While the trend is similar across regions, the differences are not large enough under this design to infer that one region is experiencing substantially greater labor market disruption than the other.

At the same time, the results for higher seniority groups reinforce why firm-level controls are essential for interpretation. For *Developing* and beyond, raw series can look strongly stratified by exposure, but regression-adjusted effects are substantially smaller and can change sign. In the E.U., only one intermediate exposure group remains positive and significant at the end of the sample, while a higher exposure group is negative and significant. In the U.S., some mid-exposure groups are positive and significant, while the highest exposure groups are not distinguishable from zero. This

gap between raw and adjusted patterns is informative: it implies that without absorbing firm time variation, exposure gradients can partly reflect where employment is expanding or contracting.

### **6.2. Occupation-Based Framework for Granular Monitoring.**

The second objective was methodological and practical: to complement quintile-level analyses with an occupation-based design that allows employment dynamics to be examined occupation by occupation, revealing more granular and informative patterns of adjustment. This focus is motivated by the need to understand how reallocation unfolds across specific roles rather than only across exposure bins, while also maintaining a framework that can be readily updated as the technology frontier and adoption environment continue to evolve. When conditioning on a narrow O\*NET occupation, the analysis can reveal within-occupation reallocations across seniority that are invisible in quintiles. The *Customer Service Representatives* example is particularly instructive because it is large, plausibly exposed to LLM-relevant tasks, and features a clear seniority structure. Within this single occupation, early career groups contract in both regions, but the *Developing* category diverges. From an applied perspective, this is not merely a statistical nuance. Firms make workforce and people sustainability decisions at the job family level, not at the abstract exposure quintile level. A logistics firm will care about planners, customer service, warehouse supervisors, and procurement analysts, while a manufacturer will care about quality roles, maintenance, and process engineers. The occupation-based framework supports that decision-making logic directly: practitioners can focus on the occupations that are common in their industry, examine whether adjustments are concentrated in hiring pipelines or in mid-career roles, and benchmark against both regions.

A complementary summary examines the correlation between occupation by seniority employment changes and AI exposure scores. This last panel helps picture a more complex story. In the E.U., both the unweighted and weighted patterns highlight early-career roles as the leading edge of adjustment, consistent with the share-sensitivity “canaries” mechanism. In the U.S., the employment-weighted correlations indicate that exposure-linked contraction is strongly present in large occupation groups, implying that adjustment may be concentrated in scale-intensive job families. This implies that, rather than reflecting only marginal adjustment at the hiring pipeline, the U.S. pattern is also consistent with scale-intensive contraction in large, exposed occupation groups, extending beyond the earliest career stages (especially in the top 300 firms subset). Put differently, the E.U. pattern points to earlier adjustment at the margin of entry, whereas the U.S. pattern suggests that exposure-related contraction is more visible in the employment mass of exposed roles, even when the quintile event-study panels appear comparatively muted and not statistically significant.

### **6.3. Limitations and Avenues for Extension.**

Several limitations should temper interpretation and motivate follow-on work. The underlying job data, while rich, may be skewed and not fully representative of the regional picture. Selecting companies with more than 300 employees limits the study to a “large-firm” context that is blind to micro-shifts in employment structures, which could potentially shift results (selection bias). We use the sentiment-share AI

exposure scores, which are not comprehensive of the whole O\*NET code universe, as some of the tasks used to create the scores lack data. Furthermore, the empirical design cannot fully capture all macro and geopolitical forces that shape labor demand contemporaneously, including pandemic aftereffects, wars, energy price shocks, and trade and tariff dynamics. While firm time fixed effects help absorb many firm-specific shocks, broad region-wide or sector-wide shocks that differentially coincide with exposure may still confound interpretation.

As studies suggest, exposure measures are imperfect proxies for adoption and effective capability. An occupation can be highly exposed yet see little adoption due to regulation, data access constraints, or organizational inertia. Conversely, moderate exposure occupations may adapt rapidly if the workflow integration is straightforward. Finally, the current results focus on employment levels. A more comprehensive picture would require working-hours, wages, task composition, and mobility (particularly whether early career contraction reflects reduced hiring, faster exits, or slower promotions). Extending the occupation-based approach into a dynamic mobility framework would be a natural next step, and it would connect directly to the occupation-specific granularity that practitioners need when designing reskilling and internal transition programs.

## 7. Conclusion

This paper provides cross-regional evidence on early labor market adjustments in the AI age, using an exposure quintile framework in the spirit of the “Canaries in the Coal Mine” approach. Our first contribution is a harmonized comparison between the United States and the European Union based on large-scale worker and firm microdata and an empirical design that absorbs firm-level shocks. Under this framework, both regions exhibit a consistent pattern: Employment among early career workers declines after late 2022, and the contraction is systematically larger in more exposed groups. Once firm-level shocks are controlled for, the magnitude of the exposure-adjusted gaps along the early career pipeline is broadly similar across regions, suggesting that differences in raw employment trends do not translate into a clear divergence in exposure-driven risk between the United States and Europe.

Our second contribution is analytical. We complement the quintile-based design with an occupation-level sensitivity analysis that estimates employment changes by occupation and seniority and then relates these objects to external exposure measures. This perspective reveals heterogeneity that is mechanically averaged out by coarse exposure bins and can therefore remain hidden in standard event study panels. In practice, the occupation-based view shows that exposure effects can be concentrated in specific job families, differ sharply across seniority levels, and even reverse sign across regions within the same occupation. The case of *Customer Service Representatives* illustrates this: While early career groups contract in both regions, the *Developing* group, while expanding in Europe, does not show a clear direction in the U.S.

Crucially, the correlation analysis confirms the “canaries in the coal mine” interpretation at early career stages, with exposure-linked employment contractions concentrated among lower-seniority roles. At the same time, it reveals sharply different

adjustment patterns across regions at higher seniority levels. In the European Union, employment-weighted correlations suggest relative expansion even within occupations classified as highly exposed to AI, particularly at advanced seniority levels. On the other hand, in the United States, high-seniority roles, especially within larger firms, exhibit correlations consistent with employment contraction in high-exposure occupations. These patterns suggest that AI exposure alone is unlikely to fully account for observed employment adjustments. Instead, the way AI adoption is reflected in workforce reallocation may vary with firm characteristics and institutional context, highlighting the importance of accounting for these dimensions when interpreting exposure-employment relationships.

### **Acknowledgements and Disclosure**

The authors would like to thank Prof. Michel Bierlaire for his support and insightful feedback, as well as Prof. Amin Kaboli and Prof. Yossi Sheffi for their guidance and constructive comments throughout the development of this work. The authors are especially grateful to Pierre Thomas R. Bouquet for his continuous intellectual support and constant motivation. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The authors declare no financial relationships or other potential conflicts of interest that could have influenced the research presented in this paper.

### **Use of AI**

During the revision process, the authors used ChatGPT 5.2 to improve language clarity and readability. The authors subsequently reviewed and edited the content as needed and take full responsibility for the final manuscript.

### **Data Availability Statement**

The authors confirm that the data supporting the findings of this study are available. All experiments were conducted in a fully reproducible manner, and the source code is publicly available at <https://github.com/nicolobagnoli/Across-the-Atlantic>. Company-level information can be provided by the authors upon request.

### **References**

Allen, Robert C. 2009. *The British Industrial Revolution in Global Perspective*. Cambridge: Cambridge University Press.

Bick, Alexander, Adam Blandin, and David J. Deming. 2024. *The Rapid Adoption of Generative AI*. Working Paper 32966. National Bureau of Economic Research. <https://www.nber.org/papers/w32966>.

Bouquet, Pierre, Nicolo Bagnoli, and Yossi Sheffi. 2025. “Estimating the Task Content of Work: Workforce Design for AI-Driven Human–Robot Collaboration in Intralogistics.” Working paper, compiled December 2025.

Bouquet, Pierre, Yossi Sheffi, and Amin Kaboli. 2026. *News Sentiment as a Dynamic Predictor of Job Automation Risk*. Research Paper 2026/002. MIT Center for Transportation & Logistics. Available at SSRN.

Boysen, Nils, and René de Koster. 2024. “50 Years of Warehousing Research—An Operations Research Perspective.” *European Journal of Operational Research* Survey article.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*. Technical Report. Stanford Digital Economy Lab. <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/>.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. “Generative AI at Work.” *The Quarterly Journal of Economics* 140 (2): 889–942.

Chen, Jiafeng, and Jonathan Roth. 2024. “Logs with Zeros? Some Problems and Solutions.” *The Quarterly Journal of Economics* 139 (2): 891–936.

Chen, W. X., S. Srinivasan, and S. Zakerinia. 2025. *Displacement or Complementarity? The Labor Market Impact of Generative AI*. Working Paper 25-039. Harvard Business School. <https://www.hbs.edu/ris/Publication%20Files/25-039-05fbec84-1f23-459b-8410-e3cd7ab6c88a.pdf>.

Dell’Acqua, Fabrizio, Edward McFowland, Ethan Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Suraj Raj, et al. 2023. *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality*. Working Paper 24-013. Harvard Business School. [https://www.hbs.edu/ris/Publication%20Files/24-013\\_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdf](https://www.hbs.edu/ris/Publication%20Files/24-013_d9b45b68-9e74-42d6-a1c6-c72fb70c7282.pdf).

Demirci, Ozge. 2025. “Who Is AI Replacing? The Impact of Generative AI on Online Freelance Platforms.” *Management Science* Advance online publication.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. “GPTs are GPTs: Labor Market Impact Potential of Large Language Models.” *Science* 384 (6702): 1306–1308.

Eurostat. 2025. “EU labour market: quarterly statistics.” Accessed 2025-12-21. [https://ec.europa.eu/eurostat/statistics-explained/index.php/EU\\_labour\\_market\\_-\\_quarterly\\_statistics](https://ec.europa.eu/eurostat/statistics-explained/index.php/EU_labour_market_-_quarterly_statistics).

Felten, Edward, Manav Raj, and Robert Seamans. 2021. “How Susceptible Is the U.S. Workforce to Artificial Intelligence?” *Strategic Management Journal* 42 (10): 1840–1868.

Frey, Carl Benedikt, and Michael A. Osborne. 2017. “The Future of Employment: How Susceptible Are Jobs to Computerisation?” *Technological Forecasting and Social Change* 114: 254–280.

Guarascio, Dario, Roman Stoellinger, et al. 2025. “AI Exposure and Employment Across European Regions.” *Structural Change and Economic Dynamics* .

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D.W. Schmidt, and Bryan Seegmiller. 2025. *Artificial Intelligence and the Labor Market*. Working Paper 33509. National Bureau of Economic Research. <https://www.nber.org/papers/w33509>.

Hartley, Jonathan, Filip Jolevski, Vitor Melo, and Brendan Moore. 2024. “The Labor Market Effects of Generative Artificial Intelligence.” SSRN working paper, [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=5136877](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5136877).

ISTAT. 2025. “Informazioni sulla rilevazione: Rilevazione sulle piccole e medie imprese e sull’esercizio di arti e professioni (PMI).” Accessed 2025-12-21. <https://www.istat.it/informazioni-sulla-rilevazione/pmi/>.

Johnston, Andrew, and Christos Makridis. 2025. “The Labor Market Effects of Generative AI: A Difference-in-Differences Analysis Using a Comprehensive Census of U.S. Employers.” SSRN working paper, [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=5375017](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5375017).

Kauhanen, Antti, Mika Maliranta, and Simo Nurmi. 2024. *Assessing Early Labor Market Effects of Generative AI: Evidence from Population Data*. Working Paper 121. ETLA Economic Research. <https://pub.etla.fi/ETLA-Working-Papers-121.pdf>.

Klein Teeselink, Bouke. 2025. “Generative AI and Labor Market Outcomes: Evidence from the United Kingdom.” SSRN working paper, <https://papers.ssrn.com/sol3/papers>.

cfm?abstract\_id=5516798.

Labaschin, Benjamin, Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. 2025. *Extending ‘GPTs Are GPTs’ to Firms*. Technical Report. Centre for the Governance of AI (GovAI). <https://www.governance.ai/research-paper/extending-gpts-are-gpts-to-firms>.

Lewandowski, Piotr, Karol Madoń, and Albert Park. 2025. *Workers’ Exposure to AI Across Development Stages*. Discussion Paper 18235. IZA Institute of Labor Economics. <https://docs.iza.org/dp18235.pdf>.

Machin, Stephen, and John Van Reenen. 1998. “Technology and Changes in Skill Structure: Evidence from Seven OECD Countries.” *The Quarterly Journal of Economics* 113 (4): 1215–1244.

Manning, Sam. 2024. “Predicting AI’s Impact on Work.” GovAI Analysis, Jun. <https://www.governance.ai/analysis/predicting-ais-impact-on-work>.

Manning, Sam. 2025. *Addressing the U.S. Labor Market Impacts of Advanced AI*. Research paper. Centre for the Governance of AI (GovAI). <https://www.governance.ai/research-paper/addressing-the-u-s-labor-market-impacts-of-advanced-ai>.

Michaels, Guy, Ashwini Natraj, and John Van Reenen. 2014. “Has ICT Polarized Skill Demand? Evidence from Eleven Countries over 25 Years.” *Review of Economics and Statistics* 96 (1): 60–77.

National Center for O\*NET Development. 2024. “O\*NET Content Model and Database (Release 28.3).” U.S. Department of Labor/Employment and Training Administration.

Noy, Shakked, and Whitney Zhang. 2023. *Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence*. Technical Report. MIT Department of Economics. Working paper, [https://economics.mit.edu/sites/default/files/inline-files/Noy\\_Zhang\\_1.pdf](https://economics.mit.edu/sites/default/files/inline-files/Noy_Zhang_1.pdf).

Santos Silva, João M. C., and Silvana Tenreyro. 2006. “The Log of Gravity.” *The Review of Economics and Statistics* 88 (4): 641–658.

Teutloff, Ole, Johanna Einsiedler, Otto Kässi, Fabian Braesemann, Pamela Mishkin, and R. Maria del Rio-Chanona. 2025. “Winners and Losers of Generative AI: Early Evidence of Shifts in Freelancer Demand.” *Journal of Economic Behavior and Organization* .

U.S. Bureau of Labor Statistics. 2025. “Household Data Annual Averages: 11b. Employed persons by detailed occupation and age.” Accessed 2025-12-21. <https://www.bls.gov/cps/cpsaat11b.htm>.

van Ark, Bart, Mary O’Mahony, and Marcel P. Timmer. 2008. “The Productivity Gap between Europe and the United States: Trends and Causes.” *Journal of Economic Perspectives* 22 (1): 25–44.

Webb, Michael. 2019. *The Impact of Artificial Intelligence on the Labor Market*. Technical Report. Stanford University. Working paper, <https://ssrn.com/abstract=3482150>.

## Appendix A. O\*NET Data Examples

This appendix provides concrete examples of the O\*NET tables used in the paper and clarifies how O\*NET organizes task statements into higher-level work elements. The example tables shown below are reproduced from Bouquet, Bagnoli, and Sheffi (2025).

### A.1. O\*NET Structure and Task Hierarchy

O\*NET represents work content using a hierarchical structure that links very specific task statements to broader work-activity groupings. Figure A1 summarizes the hierarchy.

At the most granular level, *Tasks* are specific statements that are tied to an occupation. These tasks can be linked to multiple *Detailed Work Activities* (*DWAs*), reflecting that a single task may contribute to more than one broader activity category. *DWAs* are shared across occupations and serve as a standardized intermediate layer.

*DWAs* are then mapped to *Intermediate Work Activities* (*IWAs*), and each *IWA* is mapped to a *Generalized Work Activity* (*GWA*). *IWAs* and *GWAs* are also shared across occupations and provide increasingly broad descriptions of work content. In practical terms, this hierarchy enables aggregation from occupation-specific tasks to comparable, cross-occupation work activities that can be used for measurement and analysis.




Figure A1. O\*NET task hierarchy used for aggregation. Tasks are occupation-specific statements. Tasks may map to multiple DWAs (many-to-many). Each DWA maps to a single IWA (many-to-one), and each IWA maps to a single GWA (many-to-one). DWAs, IWAs, and GWAs are standardized groupings shared across occupations.

### A.2. Survey Mode and Reported Task Ratings

Task information in O\*NET is primarily collected through structured surveys administered to incumbent workers and occupation analysts. For a given occupation, respondents first indicate whether a task is relevant using a binary relevance screen. Tasks marked as not relevant are not rated further, and the resulting relevance statistics provide a direct measure of whether a task is performed in that occupation.

For tasks deemed relevant, respondents report how frequently the task is performed using a seven-category frequency scale<sup>1</sup>. O\*NET reports, for each occupation-task

<sup>1</sup>Frequency-of-Task (FT) categories: 1 = Yearly or less, 2 = More than yearly, 3 = More than monthly, 4 =

pair, the share of respondents selecting each frequency category, along with sample sizes and design-adjusted standard errors.

O\*NET also provides normal-theory confidence intervals, motivated by its sampling design and standard asymptotic arguments. In the main text, we treat the published task-frequency estimates as approximately normally distributed for the purposes of uncertainty propagation. Throughout the paper, we treat O\*NET task measures and external employment counts as independent inputs, as they originate from distinct collection systems.

### **A.3. *Data Table Examples***

Tables A1 and A2 provide illustrative rows for the occupation **11-1011.00 Chief Executives**. These examples are intended to make the structure of the underlying O\*NET tables explicit. Both example tables are reproduced from Bouquet, Bagnoli, and Sheffi (2025).

---

More than weekly, 5 = Daily, 6 = Several times daily, 7 = Hourly or more.

Table A1. Example rows from the O\*NET Task Ratings dataset for 11-1011.00 Chief Executives. Reproduced from Bouquet, Bagnoli, and Sheff (2025).

O*NET-SOC	Title	Task ID	Task	Scale ID	Scale Name	Cat.	Data Value	N	SE	LCI	UCI	Supp.	Date	Source
11-1011.00	Chief Executives	8823	Direct or...	FT	Frequency (1-7)	1	5.92	76	4.27	1.35	22.44	N	08/2023	Incumbent
11-1011.00	Chief Executives	8823	Same as above	FT	Frequency (1-7)	2	15.98	76	5.60	7.65	30.40	N	08/2023	Incumbent
11-1011.00	Chief Executives	8823	Same as above	FT	Frequency (1-7)	3	29.68	76	9.53	14.52	51.18	N	08/2023	Incumbent
11-1011.00	Chief Executives	8823	Same as above	FT	Frequency (1-7)	4	21.18	76	8.39	8.99	42.23	N	08/2023	Incumbent
11-1011.00	Chief Executives	8823	Same as above	FT	Frequency (1-7)	5	19.71	76	7.37	8.85	38.29	N	08/2023	Incumbent
11-1011.00	Chief Executives	8823	Same as above	FT	Frequency (1-7)	6	4.91	76	2.67	1.63	13.92	N	08/2023	Incumbent
11-1011.00	Chief Executives	8823	Same as above	FT	Frequency (1-7)	7	2.63	76	2.22	0.48	13.21	N	08/2023	Incumbent
11-1011.00	Chief Executives	8823	Same as above	IM	Importance	—	4.52	75	0.11	4.31	4.74	N	08/2023	Incumbent

Table A2. Example rows from the Tasks-to-DWA crosswalk for 11-1011.00 Chief Executives. Reproduced from Bouquet, Bagnoli, and Sheffi (2025).

O*NET-SOC	Title	Task ID	Task	DWA ID	DWA Title	Date	Source
11-1011.00	Chief Executives	20461	Review and analyze legislation, laws, or public policy and recommend changes to promote or support interests of the general population or special groups.	4.A.2.a.4.I09.D03	Analyze impact of legal or regulatory changes.	07/2014	Analyst
11-1011.00	Chief Executives	20461	Same as above	4.A.4.b.6.I08.D04	Advise others on legal or regulatory compliance matters.	07/2014	Analyst
11-1011.00	Chief Executives	8823	Direct or coordinate an organization's financial or budget activities to fund operations, maximize investments, or increase efficiency.	4.A.4.b.4.I09.D02	Direct financial operations.	03/2014	Analyst
11-1011.00	Chief Executives	8824	Confer with board members, organization officials, or staff members to discuss issues, coordinate activities, or resolve problems.	4.A.4.a.2.I03.D14	Confer with organizational members to accomplish work activities.	03/2014	Analyst
11-1011.00	Chief Executives	8825	Analyze operations to evaluate performance of a company or its staff in meeting objectives or to determine areas of potential cost reduction, program improvement, or policy change.	4.A.2.a.4.I07.D09	Analyze data to assess operational or project effectiveness.	03/2014	Analyst
11-1011.00	Chief Executives	8826	Direct, plan, or implement policies, objectives, or activities of organizations or businesses to ensure continuing operations, to maximize returns on investments, or to increase productivity.	4.A.2.b.1.I09.D01	Implement organizational process or policy changes.	03/2014	Analyst

## Appendix B. Share-Weighted Sentiment Exposure Summary.

### B.1. *Distribution of Share-Weighted Sentiment Exposure*

Table B1. Share-weighted sentiment exposure: summary statistics and extreme job titles.

Panel A: Job-level exposure score statistics (N=796)	
Mean	0.4411
Std. dev.	0.0711
Min	0.2694
Max	0.6921
Panel B: Top 10 most exposed job titles	
Administrative Law Judges, Adjudicators, and Hearing Officers	0.6921
Judges, Magistrate Judges, and Magistrates	0.6870
Court Reporters and Simultaneous Captioners	0.6779
Judicial Law Clerks	0.6664
Lawyers	0.6621
Title Examiners, Abstractors, and Searchers	0.6591
News Analysts, Reporters, and Journalists	0.6579
Arbitrators, Mediators, and Conciliators	0.6387
Historians	0.6209
Editors	0.6176
Panel C: Top 10 least exposed job titles	
Control and Valve Installers and Repairers, Except Mechanical Door	0.2694
Motorcycle Mechanics	0.2811
Paper Goods Machine Setters, Operators, and Tenders	0.2829
Farm Equipment Mechanics and Service Technicians	0.2837
Dishwashers	0.2839
Tree Trimmers and Pruners	0.2934
Maintenance Workers, Machinery	0.2957
Mobile Heavy Equipment Mechanics, Except Engines	0.2996
Agricultural Equipment Operators	0.3014
Watch and Clock Repairers	0.3019

### B.2. *Comparison Between Unweighted and Share-Weighted Exposure Metrics*

Figure B1 compares occupation-level unweighted sentiment exposure scores to their share-weighted counterparts for the  $\alpha$ ,  $\beta$ , and  $\gamma$  rubric dimensions. Each panel plots the unweighted score on the horizontal axis against the corresponding share-weighted score on the vertical axis, with the 45-degree line indicating perfect agreement. Reported correlations are Pearson correlations, capturing linear association in levels rather than rank similarity. Across dimensions, the correlations are positive but heterogeneous in magnitude. For  $\beta$  and especially  $\gamma$ , the correlations suggest closer alignment between unweighted and share-weighted measures along those dimensions. Taken to-

gether, these patterns indicate that share-weighting does not simply rescale exposure uniformly across occupations, but instead alters exposure in ways that depend on how concentrated exposure-relevant tasks are in the workday. While the distributions of the underlying measures differ, being more concentrated and approximately Gaussian for the sentiment scores, and more diffuse for the alternative exposure metric, this difference in marginal distributions does not drive the observed correlations. In particular, Spearman rank correlations (reported in the main text), which are invariant to monotonic transformations and depend only on ordinal rankings, yield values similar to those of Pearson. This indicates that the association reflects a stable alignment in occupational rankings rather than being mechanically induced by distributional shape.




Figure B1. Unweighted versus share-weighted sentiment exposure scores by rubric dimension. Notes: Each panel plots occupation-level unweighted exposure against time-weighted exposure. Dashed lines denote perfect agreement. Reported correlations are Pearson correlations.

### Appendix C. DOW Microdata Summary Statistics

Table C1 reports summary statistics for the DOW microdata used in the analysis.

Table C1. Microdata summary statistics (DOW microdata)

Variable	N	Unique	Top value	Top freq.	Missing
user_id	4,315,587	3,269,235	NA	45	0
position_id	4,315,587	4,315,587	NA	1	0
company_id	4,315,587	30	NA	712,437	0
seniority	4,315,587	7	1	1,247,088	0
salary	4,314,842	3,838,617	291,334.99	38	745
onet_code	4,314,870	994	15–1252	325,233	717
startdate	3,970,640	873	2022–01–01	70,646	344,947
enddate	2,025,089	69	2022–08–01	72,370	2,290,498
highest_degree	3,044,941	6	Bachelor	1,589,343	1,270,646
sex_predicted	4,315,587	3	M	2,521,163	0
eth_predicted	4,314,213	6	White	2,675,673	1,374

## Appendix D. Microdata Summary Statistics by Region

This appendix reports summary statistics for the E.U./U.S. microdata in the 2015–2025 period.

### D.1. European Union

Table D1. Microdata summary statistics: European Union

Variable	N	Unique	Top value	Top freq.	Missing
user_id	21,369,981	12,808,411	NA	NA	0
position_id	21,369,981	21,369,981	NA	NA	0
company_id	21,369,981	2,150	NA	376,229	0
seniority	21,369,981	7	2	6,229,896	0
country	21,369,981	27	France	4,726,710	0
salary	21,368,283	9,689,519	223,986.24	NA	1,698
onet_code	21,367,143	1,009	15–1299.00	898,873	2,838
startdate	20,132,666	964	2018–01–01	227,902	1,237,315
enddate	14,321,501	173	2022–09–01	177,808	7,048,480
weight	21,369,981	10,798,846	NA	67,787	0
highest_degree	12,439,355	6	Master	6,610,433	8,930,626
sex_predicted	21,369,981	3	M	12,227,555	0
eth_predicted	21,368,064	6	White	18,491,980	1,917
region	21,369,981	1	E.U.	2,136,9981	0

### D.2. United States

Table D2. Microdata summary statistics: United States

Variable	N	Unique	Top value	Top freq.	Missing
user_id	4,764,5193	28,235,297	NA	NA	0
position_id	4,764,5193	4,764,5193	NA	NA	0
company_id	4,764,5193	3,907	NA	921,971	0
seniority	4,764,5193	7	1	15,512,312	0
salary	4,764,1289	17,164,535	680,000.0	NA	3,904
onet_code	4,763,7433	1,010	15–1252.00	1,852,068	7,760
startdate	4,359,3759	958	2022–01–01	487,623	4,051,434
enddate	3,006,5397	186	2022–08–01	447,974	17,579,796
weight	4,764,5193	21,350,318	NA	619,129	0
highest_degree	3,340,8009	6	Bachelor	18,193,894	14,237,184
sex_predicted	4,764,5193	3	M	26,522,419	0
eth_predicted	4,763,1732	6	White	32,498,209	13,461
region	4,764,5193	1	U.S.	4,764,5193	0

### D.3. *Global Company Descriptive Statistics*

Table D3. Company-level employee count statistics in the constructed dataset.

Statistic	Value
Average positions per company	11,308.21
Median positions per company	2,873.00
Minimum positions in a company	373
Maximum positions in a company	771,223
Standard deviation	31,615.74
10th percentile	749
25th percentile	1,302
50th percentile	2,873
75th percentile	8,152
90th percentile	25,084
95th percentile	47,778
99th percentile	129,751

### Appendix E. Additional Figures: Exposure Scores




Figure E1. Eloundou beta rubric exposure quintiles across all seniority groups.




Figure E2. Share-weighted Eloundou beta rubric exposure quintiles across all seniority groups.




Figure E3. Share-weighted sentiment exposure quintiles across all seniority groups.

## Appendix F. E.U. vs U.S. Additional Figures and Tables

### F.1. *Raw trends panels (rotated for readability)*




Figure F1. E.U.: raw normalized employment by seniority and exposure quintile (rotated).




Figure F2. U.S.: raw normalized employment by seniority and exposure quintile (rotated).

**F.2. *Regression Tables: Raw Differences vs PPML Adjusted Effects***

Table F1. E.U.: Raw normalized employment changes and PPML-adjusted effects by seniority and AI exposure quintile

Sen.	Q	Raw (normalized)				PPML (firm fixed effects)				Raw – PPML	
		Avg. (log)	Avg. (%)	Last (log)	Last (%)	Avg. (log)	(SE)	Last (log)	(SE)	Avg. (log)	Last (log)
<i>Panel A: Early Career 1</i>											
1	2	-0.0155	-1.54	-0.0578	-5.61	-0.0002	(0.0055)	-0.0038	(0.0097)	-0.0153	-0.0540
1	3	-0.0247	-2.44	-0.0893	-8.54	0.0033	(0.0059)	-0.0094	(0.0100)	-0.0280	-0.0798
1	4	-0.0449	-4.39	-0.1015	-9.66	-0.0160***	(0.0056)	-0.0284***	(0.0055)	-0.0289	-0.0731
1	5	-0.0280	-2.77	-0.0862	-8.26	-0.0102*	(0.0061)	-0.0325***	(0.0109)	-0.0179	-0.0537
<i>Panel B: Early Career 2</i>											
2	2	-0.0165	-1.64	-0.0489	-4.77	-0.0067	(0.0043)	-0.0132**	(0.0066)	-0.0098	-0.0357
2	3	-0.0259	-2.56	-0.0607	-5.89	0.0005	(0.0050)	-0.0001	(0.0077)	-0.0264	-0.0606
2	4	-0.0406	-3.98	-0.0930	-8.88	-0.0267***	(0.0044)	-0.0513***	(0.0063)	-0.0139	-0.0417
2	5	-0.0317	-3.12	-0.0816	-7.83	-0.0184***	(0.0056)	-0.0413***	(0.0088)	-0.0134	-0.0403
<i>Panel C: Developing</i>											
3	2	0.0892	9.33	0.1438	15.46	-0.0137*	(0.0071)	-0.0174	(0.0108)	0.1029	0.1612
3	3	0.1154	12.23	0.1915	21.11	0.0194**	(0.0092)	0.0381**	(0.0148)	0.0960	0.1534
3	4	0.0857	8.94	0.1357	14.53	-0.0189**	(0.0090)	-0.0298**	(0.0140)	0.1046	0.1655
3	5	0.0974	10.23	0.1553	16.80	-0.0105	(0.0099)	-0.0170	(0.0150)	0.1079	0.1723
<i>Panel D: Mid Career 1</i>											
4	2	0.0319	3.24	0.0348	3.54	0.0040	(0.0071)	0.0123	(0.0111)	0.0279	0.0225
4	3	0.0538	5.53	0.0646	6.68	0.0328***	(0.0070)	0.0540***	(0.0106)	0.0210	0.0106
4	4	0.0417	4.26	0.0466	4.77	0.0181**	(0.0082)	0.0311**	(0.0144)	0.0236	0.0155
4	5	0.0412	4.21	0.0428	4.37	0.0116	(0.0088)	0.0137	(0.0132)	0.0296	0.0291
<i>Panel E: Mid Career 2</i>											
5	2	0.0145	1.46	0.0030	0.30	0.0024	(0.0071)	0.0168*	(0.0099)	0.0121	-0.0138
5	3	0.0333	3.39	0.0273	2.77	0.0213***	(0.0068)	0.0417***	(0.0095)	0.0120	-0.0144
5	4	0.0291	2.96	0.0153	1.54	0.0150*	(0.0077)	0.0272***	(0.0106)	0.0142	-0.0119
5	5	0.0467	4.78	0.0511	5.25	0.0205**	(0.0090)	0.0388***	(0.0135)	0.0262	0.0124
<i>Panel F: Senior</i>											
6	2	0.0642	6.63	0.0732	7.59	-0.0305*	(0.0182)	-0.0330	(0.0237)	0.0947	0.1062
6	3	0.0748	7.76	0.0998	10.49	-0.0268	(0.0177)	-0.0221	(0.0233)	0.1016	0.1219
6	4	0.0834	8.69	0.1102	11.65	-0.0122	(0.0190)	0.0018	(0.0235)	0.0955	0.1084
6	5	0.1290	13.76	0.1895	20.86	0.0062	(0.0208)	0.0285	(0.0275)	0.1227	0.1609

*Notes:* Raw trends report normalized employment changes relative to each quintile's own October 2022 baseline. PPML coefficients are estimated from Poisson regressions with firm-by-quintile and firm-by-month fixed effects; Quintile 1 is the reference category. Standard errors are reported in parentheses. Differences indicate how raw trends change once firm-level shocks are controlled for. *Significance levels:* \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

Table F2. US: Raw normalized employment changes and PPML-adjusted effects by seniority and AI exposure quintile

Sen.	Q	Raw (normalized)			PPML (firm fixed effects)			Raw – PPML	
		Avg. (log)	Avg. (%)	Last (log)	Last (%)	Avg. (log)	(SE)	Last (log)	(SE)
<i>Panel A: Early Career 1</i>									
1	2	-0.0175	-1.74	-0.0627	-6.08	0.0042	(0.0041)	-0.0002	(0.0072)
1	3	-0.0363	-3.56	-0.0933	-8.91	0.0070	(0.0046)	0.0046	(0.0073)
1	4	-0.0429	-4.20	-0.0921	-8.80	-0.0162***	(0.0040)	-0.0282***	(0.0067)
1	5	-0.0391	-3.83	-0.0896	-8.57	-0.0150***	(0.0042)	-0.0317***	(0.0069)
<i>Panel B: Early Career 2</i>									
2	2	-0.0271	-2.67	-0.0644	-6.24	-0.0111***	(0.0039)	-0.0222***	(0.0063)
2	3	-0.0353	-3.47	-0.0695	-6.72	-0.0011	(0.0050)	-0.0025	(0.0078)
2	4	-0.0433	-4.23	-0.0959	-9.15	-0.0251***	(0.0044)	-0.0519***	(0.0068)
2	5	-0.0408	-4.00	-0.0836	-8.02	-0.0294***	(0.0058)	-0.0542***	(0.0086)
<i>Panel C: Developing</i>									
3	2	0.0521	5.35	0.0836	8.72	0.0040	(0.0059)	0.0184**	(0.0092)
3	3	0.0790	8.22	0.1360	14.57	0.0411***	(0.0076)	0.0810***	(0.0124)
3	4	0.0392	4.00	0.0665	6.88	-0.0127*	(0.0068)	-0.0080	(0.0110)
3	5	0.0504	5.17	0.0808	8.41	-0.0164*	(0.0093)	-0.0135	(0.0151)
<i>Panel D: Mid Career 1</i>									
4	2	-0.0031	-0.31	-0.0229	-2.26	0.0045	(0.0050)	0.0199***	(0.0074)
4	3	0.0143	1.44	0.0000	0.00	0.0300***	(0.0054)	0.0539***	(0.0081)
4	4	-0.0004	-0.04	-0.0202	-2.00	0.0102*	(0.0057)	0.0254***	(0.0091)
4	5	-0.0101	-1.01	-0.0301	-2.97	-0.0081	(0.0072)	-0.0013	(0.0106)
<i>Panel E: Mid Career 2</i>									
5	2	-0.0002	-0.02	-0.0186	-1.84	-0.0014	(0.0097)	0.0088	(0.0125)
5	3	0.0002	0.02	-0.0218	-2.16	0.0010	(0.0097)	0.0060	(0.0122)
5	4	-0.0065	-0.65	-0.0337	-3.32	-0.0033	(0.0099)	-0.0014	(0.0127)
5	5	-0.0007	-0.07	-0.0224	-2.21	0.0037	(0.0103)	0.0139	(0.0133)
<i>Panel F: Senior</i>									
6	2	0.0638	6.58	0.0917	9.60	0.0050	(0.0109)	0.0183	(0.0144)
6	3	0.0547	5.62	0.0836	8.72	-0.0060	(0.0105)	0.0066	(0.0141)
6	4	0.0435	4.45	0.0599	6.18	-0.0186*	(0.0113)	-0.0212	(0.0158)
6	5	0.0799	8.32	0.1233	13.13	0.0091	(0.0118)	0.0252	(0.0165)

*Notes:* Raw trends report normalized employment changes relative to each quintile's own October 2022 baseline. PPML coefficients are estimated from Poisson regressions with firm-by-quintile and firm-by-month fixed effects; Quintile 1 is the reference category. Standard errors are reported in parentheses. Differences indicate how raw trends change once firm-level shocks are controlled for. *Significance levels:* \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

## Appendix G. Robustness Checks by Subsample

Below are robustness checks for the E.U.–U.S. analysis using alternative subsample definitions. For each case, the empirical specification is unchanged relative to the baseline; only the underlying workforce sample varies. The figures below document that the qualitative employment dynamics across AI exposure quintiles remain stable, with increased uncertainty driven solely by smaller sample sizes.




Figure G1. U.S.: employment dynamics by AI exposure, restricted to male workers.




Figure G2. U.S.: employment dynamics by AI exposure, restricted to female workers.




Figure G3. E.U.: employment dynamics by AI exposure, restricted to male workers.




Figure G4. E.U.: employment dynamics by AI exposure, restricted to female workers.

We test on Ethnicity and if a Job is considered a "tech-job" or not (filtering SOC titles starting with "15-1") and find no greater patterns to report. Results are on the git and omitted for reasons of length.

## Appendix H. Customer Service Representatives (O\*NET 43-4051.00)




Figure H1. United States: Occupation-level employment dynamics for Customer Service Representatives by seniority. Each panel reports Poisson event-study estimates relative to the October 2022 baseline. The panel is for the top 300 companies in size in the U.S..




Figure H2. European Union: Occupation-level employment dynamics for Customer Service Representatives by seniority. Each panel reports Poisson event-study estimates relative to the October 2022 baseline. The panel is for the top 300 companies in size in the E.U.




Figure H3. United States: Occupation-level employment dynamics for Customer Service Representatives by seniority. Each panel reports Poisson event-study estimates relative to the October 2022 baseline. The panel is for randomly sampled 800 companies in the U.S.




Figure H4. European Union: Occupation-level employment dynamics for Customer Service Representatives by seniority. Each panel reports Poisson event-study estimates relative to the October 2022 baseline. The panel is for randomly sampled 800 companies in the E.U.

### Appendix I. Monte Carlo Construction of Confidence Intervals for Correlations

The key reason for using a Monte Carlo (MC) procedure is that the object being correlated with exposure is not a directly observed outcome: it is an estimated PPML event-study coefficient (and a nonlinear transformation of it), which carries sampling uncertainty. The MC procedure propagates that coefficient-level uncertainty into uncertainty over the correlation itself. In the event-study PPML specification, the period-specific effects enter as coefficients on indicators (denoted  $\delta_{o,k}$  in the main text). For each occupation  $i$  within a given seniority group (and region), we observe a fixed exposure score  $x_i$  and we estimate a last-period event-study coefficient  $\hat{\delta}_i$  with standard error  $SE_i$ . Because PPML is estimated on the log scale,  $\hat{\delta}_i$  is interpreted as a log change relative to the omitted baseline period. For interpretability, we convert this estimated log effect into a percent change via the standard transformation

$$y_i = 100 \left( \exp(\hat{\delta}_i) - 1 \right).$$

We then compute the Pearson correlation across occupations within the seniority group between exposure and estimated impact, i.e.,  $\text{Corr}(x_i, y_i)$ . We compute both an unweighted correlation (each occupation contributes equally) and a weighted correlation that assigns each occupation a weight  $w_i$ . A naive confidence interval for a correlation typically assumes both variables are directly observed without estimation error. Here, however, the quantity  $y_i$  is itself a transformation of an estimated coefficient. This matters for two reasons. First,  $\hat{\delta}_i$  has sampling uncertainty summarized by  $SE_i$ , so

$y_i$  inherits uncertainty even if  $x_i$  is fixed. Second, the mapping  $\delta \mapsto 100(\exp(\delta) - 1)$  is nonlinear, so uncertainty that is approximately symmetric on the log scale can become asymmetric on the percent-change scale. In such settings, a percentile-based MC interval is a clean way to capture the induced distribution of the correlation without relying on linear approximations. The MC construction proceeds as follows. For each occupation  $i$ , we approximate the sampling distribution of the coefficient by a normal distribution centered at the estimate with variance equal to the squared standard error:

$$\delta_i^{(r)} \sim \mathcal{N}(\hat{\delta}_i, SE_i^2),$$

where  $r \in \{1, \dots, R\}$  indexes Monte Carlo draws. For each draw  $r$ , we transform the simulated coefficient into a simulated percent change

$$y_i^{(r)} = 100(\exp(\delta_i^{(r)}) - 1),$$

and then recompute the correlation across occupations using the fixed exposure values  $x_i$ . In the unweighted case, the simulated correlation is

$$\rho^{(r)} = \text{Corr}(x_i, y_i^{(r)}).$$

In the weighted case, we compute the weighted Pearson correlation, which can be written as

$$\rho_w^{(r)} = \frac{\sum_i w_i (x_i - \bar{x}_w) (y_i^{(r)} - \bar{y}_w^{(r)})}{\sqrt{\sum_i w_i (x_i - \bar{x}_w)^2} \sqrt{\sum_i w_i (y_i^{(r)} - \bar{y}_w^{(r)})^2}},$$

where  $\bar{x}_w = \frac{\sum_i w_i x_i}{\sum_i w_i}$  and  $\bar{y}_w^{(r)} = \frac{\sum_i w_i y_i^{(r)}}{\sum_i w_i}$ . Repeating this for  $R$  draws yields an empirical distribution  $\{\rho^{(1)}, \dots, \rho^{(R)}\}$  (or  $\{\rho_w^{(1)}, \dots, \rho_w^{(R)}\}$ ). We report a 95% confidence interval using the 2.5th and 97.5th percentiles of the simulated correlations:

$$CI_{95} = [\text{percentile}_{2.5}(\rho^{(r)}), \text{percentile}_{97.5}(\rho^{(r)})].$$

In this implementation, we do not resample occupations; the observed occupation set within each seniority bucket is treated as the analysis population. The only uncertainty propagated is the measurement uncertainty of  $\hat{\delta}_i$  captured by  $SE_i$ . This is appropriate when the goal is to reflect estimation uncertainty from the underlying PPML model in the correlation summary, while keeping the composition of occupations fixed.

## Appendix J. Top 300 companies by size correlation panel




Figure J1. The x-axis in each chart is seniority level from 1 to 6, and the y-axis is Pearson correlation  $r$ . For each seniority level, there are two bars: an unweighted correlation and an employment-weighted correlation, shown with diagonal hatching. 95% CI calculations are described in Appendix I. The panel comprises the top 300 companies by size in the dataset.