How Does AI Change Labor Demand? Evidence from 41 Countries

Bharat Chandar*[1] and Bouke Klein Teeselink[2]

[1] *Stanford University*
[2] *King's College London, AI Objectives Institute, and King's Institute for Artificial Intelligence*

September 20, 2026

Abstract

New technologies can raise demand for some workers while lowering it for others, but knowing which workers gain and which lose is challenging. We study this question for generative AI. Our model decomposes changes in labor demand following firm AI adoption into a firm-wide productivity effect and worker-type-specific changes in task content. We take the model to the data, analyzing a sample of 1.25 billion job postings and 154 million employment records across 41 countries. We infer AI adoption from job advertisements that involve generative AI use. An instrumented event study shows that foreign affiliates of AI-adopting companies reduce the junior share of their workforce relative to comparable control affiliates. The decline primarily comes from growth in senior employment rather than from a fall in junior employment, with suggestive evidence of modest overall employment growth. Senior employment shifts toward AI-exposed occupations, while our point estimates suggest a shift away from these occupations among juniors. The model-based decomposition provides suggestive evidence of modest productivity gains from AI adoption.

Keywords: Generative Artificial Intelligence; Substitution; Complementarity; Technology adoption; Labor Demand

JEL: O33; J23; J24; J63

*Contact Bharat Chandar: chandarb@stanford.edu, Contact Bouke Klein Teeselink: bouke.klein_teeselink@kcl.ac.uk. We are thankful to Daron Acemoglu, Lukas Althoff, Erik Brynjolfsson, Daniel Carey, Anders Humlum, Chad Jones, Benny Kleinman, Ioana Marinescu, Neil Thompson, Chris Tonetti, and Lee Tucker, as well as seminar participants at the Stanford Digital Economy Lab, the University of Chicago, the Bank of England, the Barcelona Summer Forum, and the International Labor Organization for valuable comments and suggestions. We are grateful to Sangeeta, our Hindi teacher, whose class brought us together and made this collaboration possible. Bouke Klein Teeselink acknowledges generous support from the S&P Global Research Fellowship.

1 Introduction

New technologies can have different effects across workers within a firm. Between 1920 and 1940, AT&T replaced telephone operators with mechanical switching across much of its network, one city at a time. Feigenbaum and Gross (2024) show that when a city's exchange converted to dial, the local telephone workforce lost more than a third of its operators, with younger operators seeing the biggest declines. Yet the same conversion raised employment of clerks, mechanics, electrical engineers, and managers, who took on tasks the operators had performed and new tasks the automatic equipment required. The same technology, in the same firm, lowered demand for some workers and raised it for others.

We develop a model and an empirical strategy to identify how generative AI adoption changes demand for different groups of workers within firms. First, AI has a scale effect, whereby the technology lowers production costs and allows firms to expand, increasing demand for all workers across the firm. Second, AI changes the tasks that particular groups of workers perform. When AI fully takes over certain tasks, it reduces demand for the workers who used to perform those tasks. When AI creates new tasks, this increases demand for workers who perform them. Productivity gains on existing tasks can go in either direction, as firms now require fewer workers to complete a given number of tasks, but may increasingly rely on these tasks in the production process. We call the net worker-specific effect labor saving if it reduces demand, and labor expanding if it increases demand. The scale effect can offset labor saving, so employment could grow even for groups whose tasks are partially displaced.

These firm-level responses need not translate into similar changes in aggregate employment. Our general equilibrium model captures this distinction by accounting for adjustments in wages, prices, and labor supply. We derive a first-order relationship between firm-level employment effects and aggregate employment changes, showing that they can have opposite signs. To understand how, consider that employment at adopters can grow relative to non-adopters because adopters lower their prices and increase their market share. This redistribution of sales need not, however, increase aggregate spending on labor. In a benchmark case where firms have the same production functions and adoption raises productivity uniformly across occupations, productivity gains from AI capital raise real output but reduce the labor share. If the value of non-employment also rises in real terms, this may induce some workers to leave employment altogether. This causes aggregate employment to fall, even though adopters grow compared to comparable non-adopters.

We then take the model to the data by examining how company AI adoption changes affiliate employment relative to comparable control affiliates. We use Revelio Labs data, containing 1.25 billion job postings and 154 million employment records across 41 countries. We identify AI adoption from job advertisements that involve generative AI, and construct a monthly affiliate-level employment panel from

January 2021 through March 2026. We track workforce size and composition across occupations and seniority levels to distinguish the scale effect from worker-specific changes in labor demand. We identify the scale effect from employment changes in occupations that are not directly affected by AI between treated and control affiliates. We then use changes in employment shares across different worker groups to identify whether the worker-specific effects of AI are labor saving or labor expanding for each group. Combining these worker-specific effects with the scale effect identifies whether AI adoption ultimately increases or decreases each group's overall employment relative to non-adopters.

To estimate the effects of adoption, we must separate treatment from selection. A positive shock to a firm, for example, may make it both more likely to adopt AI and more likely to hire. Our design therefore proceeds in two steps. We first match each country affiliate in an adopting company to observably similar non-adopter affiliates in the same country and sector. We then use the multinational structure of the data, instrumenting for AI adoption with the adoption rate among other multinational parents headquartered in the same metropolitan area as the (foreign) parent. The main idea is that firms are more likely to adopt when the firms around them adopt (Cai and Szeidl, 2018; Kalyani et al., 2025), and corporate practices flow from parents to their foreign affiliates (Bloom et al., 2012; Guadalupe et al., 2012; Bilir and Morales, 2020). Together, these two processes give the instrument its relevance. Identification then requires that peer adoption around the parent's headquarters affects the affiliate's workforce only through the multinational's AI adoption. This is ex ante plausible because the peers operate in a different country from the affiliate and in different product and labor markets. We probe this assumption via pre-trends and from several other angles in the robustness section.

The results show that AI adoption changes which workers an affiliate employs more than how many. The junior share of employment at treated affiliates begins to fall after the introduction of ChatGPT, and it is 1.9 percentage points lower than in control affiliates by March 2026. This decline primarily reflects growth in senior employment rather than a fall in junior employment: senior employment rises by 6.7 percent, while junior employment is 2.5 percent lower, an estimate that is not statistically significant. Adoption shifts labor demand from the least-exposed seniors to the most-exposed seniors by 5 percent. It shifts labor demand away from the most AI-exposed juniors towards the least-exposed juniors by 2 percent. Total employment goes up by a marginally significant 3.3 percent. Occupation shares move much less: in 21 of the 22 occupation groups, the change is smaller than half a percentage point. The exception is computer and mathematical occupations, the most AI-exposed of the 22 groups, whose employment share grows by 0.8 percentage points while its junior share falls by 3.3 points. Through the lens of the model, generative AI has so far led to modest productivity increases that are seniority-biased. AI is labor saving for junior workers and labor expanding for seniors in exposed occupations. AI is labor expanding for computer and mathematical occupations, but these also see the largest seniority bias in their employment response from AI.

The shift from junior toward senior employment appears across most of the world. We re-estimate the design separately for each of the 31 countries where the instrument is strong enough to support a country-specific estimate, and combine the results in a random-effects meta-analysis. The junior share decreases in 23 of the 31 countries, with statistically significant declines in seven, including the United States, Brazil, and Saudi Arabia. The average decline across countries is 1.1 percentage points. The other margins mirror the pooled results. Senior employment rises by a marginally significant 3.4 percent on average, and no country shows a significant decline in total employment. Junior employment losses run somewhat deeper in richer and more digitized economies, and are shallower where labor regulation is more rigid, although most of these correlations are statistically insignificant. Overall, the shift away from junior workers thus shows initial evidence of generalization to developing and middle-income countries.

This paper contributes to three literatures. The first is the task framework of Autor et al. (2003); Acemoglu and Autor (2011); Acemoglu and Restrepo (2018, 2019, 2022), in which a technology displaces workers from some tasks, reinstates them in others, and raises their productivity in the tasks they keep, while the adopting unit's lower cost raises its demand for all labor (for an overview, see Acemoglu, Kong and Restrepo, 2025).[1] Recent work brings this framework to the firm, modeling adopters who pay a fixed cost to automate. In Humlum (2019), adopters expand output, shed production workers, and hire technical workers. Hubmer and Restrepo (2026) use a firm-dynamics model to explain how the aggregate labor share can fall while the typical firm's labor share rises. Closest to our approach, Hampole et al. (2025) relate changes in an occupation's employment at a firm to the average exposure of its tasks to AI, the concentration of exposure across those tasks, and firm-level growth. Concentrated exposure allows workers to reallocate effort toward less-affected tasks, potentially increasing demand for the occupation. They estimate these effects for U.S. public firms. Our model differs in two respects. First, we allow adoption to raise or lower labor demand in each occupation-by-seniority cell. Exposure to AI, as measured by scores such as those of Eloundou et al. (2024), can therefore imply lower demand for one group of workers and higher demand for another. Our empirical results corroborate that AI exposure is associated with different employment responses across worker groups. Second, we derive a closed-form, first-order relationship between firm-level difference-in-differences estimates and general equilibrium employment changes, and establish conditions under which they have opposite signs. Evidence on earlier technologies shows that adopters grow relative to non-adopters (Acemoglu et al., 2020; Koch et al., 2021; Aghion et al.,

[1]Recent extensions consider how task-level automation changes jobs' expertise requirements (Autor and Thompson, 2025), raises the importance of surviving tasks (Freund and Mann, 2025), and amends comparative advantage (Althoff and Reichardt, 2026).

2026), shift their workforces toward engineers and managers (Bonfiglioli et al., 2024), and separate from incumbent workers (Bessen et al., 2025).[2]

We also contribute to the rapidly growing empirical literature on the labor market effects of generative AI. There are two primary strands of this literature. One strand scores occupations by how much of their work generative AI can perform (Felten et al., 2023; Gmyrek et al., 2023; Eloundou et al., 2024; Handa et al., 2025; Tomlinson et al., 2025; Gmyrek et al., 2025), and examines the labor market trajectories of more versus less exposed occupations and firms (Brynjolfsson et al., 2026; Klein Teeselink, 2025; Azar et al., 2025; Klein Teeselink and Carey, 2026; Lambert and Schindler, 2026; Orr et al., 2026). Many of these papers find evidence against widespread displacement but find employment declines in more exposed occupations for junior workers, with some evidence of pre-trends before the launch of ChatGPT (Frank et al., 2026; Tucker, 2026).[3] A second strand matching the empirical setting in this paper measures which firms adopt the technology and compares their outcomes with those of non-adopters. Conclusions in the literature range from declines in junior employment (Hosseini Maasoum and Lichtinger, 2025) to null effects for earnings and hours (Humlum and Vestergaard, 2025; Chen and Stratton, 2026; Aldasoro et al., 2026; Yotzov et al., 2026; Bonney et al., 2026) to increases in employment (Kharazian et al., 2026). Work on earlier workplace AI uses similar designs and finds declines in labor demand in exposed jobs alongside growth of adopting firms (Acemoglu et al., 2022; Babina et al., 2024; Hampole et al., 2025). The existing research on the employment effects of generative AI lacks a framework for reconciling discrepancies between the first strand of literature that studies AI exposure and the second strand that studies AI adoption. In our model, an adopter–non-adopter contrast reflects the balance of labor saving, labor expanding, and scale effects for the group of workers in question, and the closed-form relation above states when that contrast and the market-level change in an exposed occupation should differ.[4] Our main empirical contribution is to infer the balance of these forces from affiliate workforce changes alone, without making any ex-ante assumptions about the sign

[2]See also Acemoglu, Anderson, Beede, Buffington, Childress, Dinlersoz, Foster, Goldschlag, Haltiwanger, Kroff, Restrepo and Zolas (2025) on which firms adopt advanced technologies, Dixon et al. (2021) on robots and managerial employment, Hirvonen et al. (2022) on technology subsidies and skill demand in Finland, and Lindner et al. (2022) on firm-level skill-biased technological change. For earlier workplace AI, establishments with AI-suitable tasks reduce non-AI hiring (Acemoglu et al., 2022), and AI-investing firms grow (Babina et al., 2024).

[3]See also Chandar (2025); Eckhardt and Goldschlag (2025); Gimbel et al. (2025); Massenkoff and McCrory (2026).

[4]Yotzov et al. (2026) and Bonney et al. (2026) directly ask firms if AI has caused them to increase or decrease employment, with most firms reporting that it has not changed employment meaningfully. Our model may also help with interpreting these results. Firms that do not use AI may report that AI has not affected their employment, while firms that do use AI may grow relative to competitors and even in absolute terms due to the scale effect in our model. These firm-level findings can still be consistent with shifts in general equilibrium suggested by trends in occupational AI exposure studies.

of impacts.

Our analysis also provides some of the first international estimates for the employment effects of AI adoption across countries. Most existing evidence on generative AI and employment, including the work that uses adoption measures closest to ours, comes from single high-income countries: the United States (Brynjolfsson et al., 2026; Hosseini Maasoum and Lichtinger, 2025; Azar et al., 2025), the UK (Klein Teeselink, 2025), Denmark (Humlum and Vestergaard, 2025), Finland (Kauhanen and Rouvinen, 2025), Sweden (Engberg et al., 2025), and Norway (Hernæs and Kostøl, 2026). The exceptions are Yotzov et al. (2026), who compare firm surveys in four countries; Lambert and Schindler (2026), Klein Teeselink and Carey (2026), and Huang et al. (2026) who study exposure in postings data across a number of countries; and three studies of developing economies, Aguirre and Meloni (2026) in Brazil, Ravanilla (2026) in the Philippines, and Hanney and Kim (2026) in India, each of which finds evidence of slowing labor demand in certain AI-exposed jobs. Because these studies differ in data, adoption or exposure measure, outcomes, and design, their estimates are hard to compare: within the United States, and within Denmark, different approaches yield declines in junior employment in some studies and null effects in others (Hosseini Maasoum and Lichtinger, 2025; Chen and Stratton, 2026; Humlum and Vestergaard, 2025; Bonin et al., 2026). We use the same data source, outcomes, adoption definition, and instrument in all 41 countries, so that differences in estimates reflect differences in effects rather than methodological differences. The comparison shows that the decline in the junior share is widespread and appears in the majority of countries, with statistically significant effects in economies as different as the United States, Brazil, and Saudi Arabia. We also examine whether income, digitization, and labor market regulation explain why the effects are larger in some countries than in others. We caution that even while using a consistent methodology, the composition of firms who appear in the Revelio data may vary across countries, so the estimates should be understood as relevant to digitally attached multinational firms.

Our final contribution is to the measurement of AI adoption itself. Following the job-posting approach of Hosseini Maasoum and Lichtinger (2025), we identify AI adoption from job advertisements that involve generative AI. We apply this measure to 1.25 billion postings across 41 countries.[5]

The remainder of the paper proceeds as follows. Section 2 develops the model. Section 3 describes the Revelio Labs data and the adoption measure, and documents how adoption varies across countries and affiliates. Section 4 sets out the matching procedure and the cross-border instrument. Section 5 reports the estimates, compares them across countries, and uses the model to decompose the estimated

[5]Alternative measures of firms' AI activity, such as surveys, patents, investor disclosures, and software usage logs, capture different objects: whether firms invent AI or use it, whether they build or buy it, and whether the data record what firms do or what investors believe (Babina, 2026).

employment changes. Section 6 concludes.

2 Theoretical Framework

The theoretical framework developed in this section characterizes how firms adjust their workforce in response to AI adoption. In Proposition 1 we derive how employment changes in adopting firms relative to comparable non-adopting firms. We show how this may vary across groups of workers, such as by occupation and seniority.

The effect decomposes between three channels. *Labor saving* occurs when AI replaces a worker's tasks, depressing the firm's demand for their services. When tasks are complements in production, task-level productivity increases also reduce labor demand. *Labor expanding* occurs when AI enables workers to take on new tasks like prompt engineering or AI auditing. Finally, increases in a firm's productivity lead to a *scale effect*, with the firm expanding its inputs to meet rising demand. The first two channels reshape *which* workers and factor inputs the firm employs; the third reshapes *how many*. The balance of the first two decides whether AI is a technology biased towards certain groups of workers. This section formalizes the channels and microfounds regression methods to identify for which occupation–seniority cells AI is labor saving and for which it is labor expanding. Derivations and proofs are in Appendix A1; Appendices A2 and A3 close the model and relate the firm-level comparison to general-equilibrium employment changes.

2.1 Environment

There is a continuum of firms $f \in \mathcal{F} = [0, 1]$, each producing a differentiated variety. A representative consumer has CES preferences over the varieties:

$$Y = \left[\int_{\mathcal{F}} \xi_f^{1/\sigma} \, y_f^{(\sigma-1)/\sigma} \, df \right]^{\sigma/(\sigma-1)}, \tag{1}$$

where $\sigma > 1$ is the elasticity of substitution across firms and ξ_f is a firm-specific demand shifter. Maximizing (1) subject to the budget $\int p_f y_f \, df \leq E$ yields

$$y_f \;=\; \xi_f \, p_f^{-\sigma} \, P^{\sigma-1} \, E, \tag{2}$$

where P is the CES price index of the varieties. We take E as given until Section 2.5, which closes the model with a household side and an income identity.

Workers have an exogenous seniority $s \in \mathcal{S}$ and an occupation $o \in \mathcal{O}$. A *worker type* is the pair $t = (o, s) \in \mathcal{T} \equiv \mathcal{O} \times \mathcal{S}$, and w_t is its wage. Besides labor, firms can use AI capital k, rented at rate r per unit, but only after paying a firm-specific fixed cost of adoption $\mathcal{F}_f \geq 0$.

Output of firm f is a CES aggregate of its factor inputs,

$$y_f \;=\; \left[\sum_g \left(\Psi_f^g \, x_f^g\right)^{(\eta-1)/\eta}\right]^{\eta/(\eta-1)}, \tag{3}$$

where g runs over the labor types and, at adopters, AI capital; x_f^g is the quantity of factor g, written ℓ_f^t for labor type t and k_f for AI capital; and $\Psi_f^g > 0$ is the factor's *effective productivity* at firm f. The elasticity of substitution across factors, $\eta > 0$ with $\eta \neq 1$, is maintained throughout to be below the elasticity across firms, $\eta < \sigma$.[6]

Appendix A1.1 derives (3) from a task model. Firm output is a CES aggregate, with elasticity η, over a continuum of tasks, and the firm assigns each task to the factor that performs it at the lowest unit cost. The quantity $(\Psi_f^g)^{\eta-1}$ is then the factor's productivity, raised to $\eta - 1$, integrated over the tasks assigned to it; we call it the factor's *effective task content*. It rises when the factor's set of tasks expands and declines when the set of tasks shrinks. A productivity gain on the tasks the factor already performs raises it when tasks are substitutes ($\eta > 1$) and lowers it when they are complements ($\eta < 1$), since less of a more productive factor is then needed per unit of output. Equation (3) may alternatively be taken as primitive.

2.2 AI adoption and the labor-demand shifter

When firm f adopts AI, tasks may be reassigned across factors and productivity may change within tasks. Three things can change relative to the pre-adoption configuration:

(i) **Displacement.** Some tasks formerly performed by labor type t are reassigned to AI capital because AI's unit cost on them falls below labor's: the set of tasks performed by the type *shrinks*.

(ii) **Reinstatement.** New tasks are created (e.g., AI auditing, prompt engineering, AI workflow design), and the set of tasks performed by the type *expands*.

(iii) **Per-task productivity gains.** Within tasks a worker continues to perform, AI assistance may raise productivity (e.g., copilots, summarization tools).

In addition, AI capital performs a set of tasks of positive measure after adoption. Let $\Psi_f^{t,0}$ and $\Psi_f^{t,A}$ denote the type's effective productivity before and after adoption. The effect of adoption on labor type t at firm f is summarized by the log change in the type's effective task content,

$$\alpha^{(o,s,f)} \;\equiv\; (\eta - 1)\,\ln\!\left(\frac{\Psi_f^{t,A}}{\Psi_f^{t,0}}\right). \tag{4}$$

[6]The case $\eta = 1$ is the Cobb-Douglas limit, where the outer exponent $\eta/(\eta-1)$ of (3) is singular.

The shifter decomposes additively into the contributions of channels (i)–(iii):

$$\alpha^{(o,s,f)} \;=\; \alpha^{(o,s,f)}_{\text{reinst}} \;-\; \alpha^{(o,s,f)}_{\text{disp}} \;+\; \alpha^{(o,s,f)}_{\text{prod}}. \tag{5}$$

Each component is a log ratio of effective task content on a different set of tasks: $\alpha_{\text{disp}} \geq 0$ is the log task content of the tasks the type no longer performs, $\alpha_{\text{reinst}} \geq 0$ the log task content gained from new tasks, and α_{prod} the log change in task content on the tasks the type retains. A per-task productivity gain raises α_{prod} when $\eta > 1$ and lowers it when $\eta < 1$. Appendix A1.2 gives the formulas under the task microfoundation and proves (5).

We call $\alpha^{(o,s,f)}$ the *labor-demand shifter* for type (o, s) at firm f. AI is *labor-saving* for a type if its shifter is negative and *labor-expanding* if it is positive. Statements about two types compare their shifters (we write $\alpha^{(t,f)}$ for $\alpha^{(o,s,f)}$, and drop f when a single firm is understood): AI adoption is *biased toward* type t relative to t' if $\alpha^{(t)} > \alpha^{(t')}$, and *biased against* it otherwise.

2.3 Firm behavior

We now derive factor demand and the adoption decision at given factor prices. Write w_g for the price of factor g: the wage $w_t \equiv w_{o,s}$ for labor type $t = (o, s)$ and $w_k \equiv r$ for AI capital. Cost minimization under (3) gives the factor demands and the unit cost

$$x^g_f \;=\; \left(\Psi^g_f\right)^{\eta-1} w_g^{-\eta}\, c_f^{\eta}\, y_f, \tag{6}$$

$$c_f \;=\; \left[\sum_g \left(w_g/\Psi^g_f\right)^{1-\eta}\right]^{1/(1-\eta)}, \tag{7}$$

with the sums over the labor types and, at adopters, AI capital.[7] We write c^A_f for the unit cost of firm f if it adopts, evaluated at the post-adoption productivities $\Psi^{t,A}_f$ and including AI capital with productivity $\Psi^{k,A}_f > 0$, and c^{NA}_f for its unit cost if it does not, evaluated at $\Psi^{t,0}_f$ over the labor types only. By (4), an adopter's demand for labor type t is $\ell^t_f = e^{\alpha^{(o,s,f)}}(\Psi^{t,0}_f)^{\eta-1} w_t^{-\eta} (c^A_f)^{\eta} y_f$.

Each firm has zero mass in the sectoral aggregates, so it takes the price index P and expenditure E as given, and sets the standard CES markup $p_f = \mu c_f$ with $\mu \equiv \sigma/(\sigma - 1)$. Combining with (2), output and variable profit are

$$y_f \;=\; \xi_f\, \mu^{-\sigma}\, c_f^{-\sigma}\, P^{\sigma-1}\, E, \tag{8}$$

$$\pi_f \;=\; \frac{\mu^{1-\sigma}}{\sigma}\, \xi_f\, c_f^{1-\sigma}\, P^{\sigma-1}\, E, \tag{9}$$

[7] The production function is a CES in the effective inputs $\Psi^g_f x^g_f$ with prices w_g/Ψ^g_f, and (6)–(7) are the standard CES formulas; see Appendix A1.3.

each evaluated at c_f^A or c_f^{NA} according to the firm's adoption status. A firm adopts iff the profit gain covers the fixed cost, $\pi_f^A - \pi_f^{NA} \geq \mathcal{F}_f$, that is, iff

$$\xi_f \Big[\big(c_f^A\big)^{1-\sigma} - \big(c_f^{NA}\big)^{1-\sigma} \Big] \; \geq \; \frac{\sigma\, \mathcal{F}_f}{\mu^{1-\sigma}\, P^{\sigma-1}\, E}. \tag{10}$$

Adoption is more attractive for larger firms (greater ξ_f), for firms with a larger cost reduction from AI, and for firms with a lower fixed cost. Because $\mathcal{F}_f \geq 0$ and $\sigma > 1$, every adopter has $c_f^A \leq c_f^{NA}$, strictly when $\mathcal{F}_f > 0$.

2.4 Comparing employment at adopters and non-adopters

Call two firms f, f' *twins* if they share all production-side fundamentals—the demand shifter $\xi_f = \xi_{f'}$, the pre-adoption effective productivity $\Psi_f^{t,0} = \Psi_{f'}^{t,0}$ for every type t, and the potential adoption technology $(\{\alpha^{(t,f)}\}_t, \Psi_f^{k,A})$—but differ in their adoption fixed cost: $\mathcal{F}_f$ lies below the cutoff (10) and $\mathcal{F}_{f'}$ above it, so that f adopts and f' does not. The next proposition compares an adopting twin to its non-adopting twin at the same factor prices. The matched event study estimator of Section 4 approximates this empirically: pre-period matching proxies for twin status, and restricting to firms in a common labor market delivers common factor prices.

Proposition 1 (Adopter–non-adopter type-level employment ratio)**.** *Let f be an adopter and f' a non-adopting twin facing the same factor prices $\{w_t\}_{t\in\mathcal{T}}$ and r, the same sectoral price index P, and the same nominal expenditure E. For every type $t = (o, s)$,*

$$\frac{\ell_f^{t,A}}{\ell_{f'}^{t,NA}} \; = \; e^{\alpha^{(o,s,f)}} \cdot \left(\frac{c_{f'}^{NA}}{c_f^A} \right)^{\sigma-\eta}. \tag{11}$$

Taking logs:

$$\ln \ell_f^{t,A} - \ln \ell_{f'}^{t,NA} \; = \; \underbrace{\alpha^{(o,s,f)}}_{\textit{type-specific shifter}} \; + \; \underbrace{(\sigma - \eta)\, \ln\!\big(c_{f'}^{NA}/c_f^A\big)}_{\textit{firm-specific scale term}}. \tag{12}$$

The proof is in Appendix A1.3. The decomposition in (12) is the central object we take to the data. The first term is *type-specific*: it varies across cells within a firm via $\alpha^{(o,s,f)}$. The second term is *firm-specific*: it is identical across cells within a firm (depending only on the firm's overall cost reduction).

The scale term is positive for every adopter with $\mathcal{F}_f > 0$: twins share $\Psi^{t,0}$, so $c_{f'}^{NA} = c_f^{NA}$, the log cost reduction $\ln(c_{f'}^{NA}/c_f^A)$ is then positive by (10), and $\sigma > \eta$. Its coefficient combines two responses to a one-percent reduction in unit cost. Holding wages fixed, labor per unit of output falls by η percent, because each type has become relatively more expensive than the firm's input bundle. Output rises by σ

percent, because the cheaper variety gains market share. Since $\sigma > \eta$, the output response dominates, so cost-reducing adoption raises the adopter's demand for every type, including types for which AI is labor-saving. A type can lose tasks to AI and still grow at the adopting firm, provided the firm grows by enough. The overall effect of adoption on a type's employment is the sum of the type-specific shifter and the firm-specific scale term.

Remark 1 (Identification)**.** *For a type with no change in effective productivity ($\alpha^{(o,s,f)} = 0$), the log employment difference between an adopter and its non-adopting twin equals the scale term $(\sigma - \eta)\ln(c_{f'}^{NA}/c_f^A)$. The scale term can be identified from employment changes for types not directly affected by AI.*

The difference in employment changes between any two types $t = (o, s)$ and $\tilde{t} = (\tilde{o}, \tilde{s})$ equals

$$\left[\ln \ell_f^{t,A} - \ln \ell_{f'}^{t,NA}\right] - \left[\ln \ell_f^{\tilde{t},A} - \ln \ell_{f'}^{\tilde{t},NA}\right] \;=\; \alpha^{(o,s,f)} - \alpha^{(\tilde{o},\tilde{s},f)}. \tag{13}$$

Differences in employment changes across types t and $\tilde{t}$ difference out the scale term. They thus show whether AI is biased towards or against a type t relative to another type $\tilde{t}$. For example, under the model differences in employment changes between junior and senior workers within a firm recover whether AI is on average biased toward or against junior workers relative to senior workers. Within a seniority level, differences in employment changes between more and less exposed occupations recover whether AI is biased toward or against the more exposed occupation relative to the less exposed occupation.

In particular, the scale term and the labor-demand shifters may be separately identified. The scale term is identified from log employment changes for unexposed workers. The labor demand shifter for a type is identified by comparing its employment changes to the unexposed type. This method holds as long as the firm employs workers with $\alpha^{(o,s,f)} = 0$.

Remark 2 (CES assumption)**.** *Under CES firm production functions, the type-specific term depends only on the type's own shifter $\alpha^{(o,s,f)}$, and the common term equals $(\sigma - \eta)\ln(c_{f'}^{NA}/c_f^A)$. Appendix A4 extends Proposition 1 to nested CES, where a nest-specific cost term enters between the type-specific and the firm-level terms.*[8]

Relation to labor-share evidence. Firm-level studies of other technologies infer labor-saving technology from a fall in the labor share of costs at adopters (Acemoglu et al., 2020; Koch et al., 2021), or, in Humlum (2019), from occupations' wage bills

[8]Among homothetic, twice-differentiable technologies, constant pairwise elasticities of substitution characterize the CES family (Uzawa, 1962; McFadden, 1963). Absent this condition, the type-specific term in (12) depends on the vector of labor demand shifters and matrix of cross elasticities.

relative to intermediate inputs. We observe only headcount in the data, so we cannot measure changes in the labor share of inputs. Instead we estimate the relative employment effect of AI exposure across occupations to infer for which groups AI is labor saving and for which it is labor expanding. This approach requires the assumption of common markets, which allow us to measure effects in labor quantities while differencing out any wage effects. A key contribution of Proposition 1 is that the sign of the type-specific shifter is left free. A high exposure score such as Eloundou et al. (2024) can result in lower demand for one group of workers and higher demand for another. We can measure when AI is labor saving and when it is labor expanding in the data, without making any ex-ante assumptions about the sign of impacts.

2.5 Closing the model and equilibrium

Appendix A2 closes the model with a household labor supply block and an income identity. In the labor supply block, workers of each seniority choose an occupation or an outside option based on wages and a Fréchet taste shock. Firms and AI capital are owned by capitalists, and total spending equals total market income. Proposition 2 proves that an equilibrium exists.

Appendix A3 then studies the first-order effect of introducing AI into the economy. Starting from an equilibrium without AI, a set of firms gains access to a small AI technology and adopts. Proposition 3 shows that the general-equilibrium change in employment for an occupation differs from the matched adopter–non-adopter difference of Proposition 1 because of wage, price, and labor supply feedback. We show in simulation (Table A1) that all sign combinations are possible: the sign of the general equilibrium change in employment for an occupation can be the same as the sign of the matched difference, or it can be opposite. The results suggest nuance in comparing employment changes from AI estimated using different approaches. Findings in the literature that use adoption-based differences-in-differences methods estimate a different quantity than approaches that track aggregate employment changes.

3 Data and Descriptive Statistics

Our data come from Revelio Labs, a workforce-analytics provider that builds labor force datasets from publicly available online records. We use two of their products, namely online professional profiles predominantly from LinkedIn, and job postings. The LinkedIn profile data contains the near universe of all public LinkedIn profiles, containing information on each position any individual has held, including start and end dates, an estimated O*NET-SOC occupation code, seniority, and company identifiers.[9] For each occupation code, we measure AI exposure using Eloundou et al. (2024).[10]

3.1 Definitions of Entities

We use the multinational structure of the Revelio data for our empirical analysis. We define an entity as a subset of an organization in Revelio. The different types of entities are defined below.

Company A company is a record in Revelio's company database to which positions and job ads are attributed. An example of a company is Amazon Web Services.

Affiliate An affiliate is a company's operations in one country. An example is Amazon Web Services' operations in the United States.

Parent Each company record carries a pointer to a parent, which is itself a company record. A company with no separate owner is its own parent. An example is Amazon.com, Inc., which is the parent of Amazon Web Services.

Corporate group All companies that share a parent form a corporate group. An example is Amazon.com, Inc., Amazon Web Services, and every other company Amazon owns.

Foreign Affiliate An affiliate is a foreign affiliate when its parent is headquartered in a different country than where the affiliate is located. An example is Amazon Web Services' operations in the United Kingdom, which is a foreign affiliate of Amazon.com, Inc., headquartered in the United States.

These definitions are summarized in Table 1. Below, we discuss how we construct the sample and define treatment and control. As a preview, a unit of observation is a foreign affiliate. In our main specification, AI adoption is defined at the company level using job postings. We instrument for company adoption using adoption rates among peers in the parent's metropolitan area.

[9] Revelio assigns each position one of seven seniority levels: entry, junior, associate, manager, director, executive, and senior executive. The levels come from an ensemble model that combines the position's title, company, and industry with the worker's career history and age.

[10] Eloundou et al. (2024) rate, for each of the 19,265 tasks in O*NET, whether a large language model by itself (alpha), or with additional software built on top (beta), could cut the time a worker needs to perform the task in half at equal quality. We use their occupation level beta scores.

Table 1: Entities in the data

Entity	Definition	Example
Company	A record in Revelio's company database. Positions and job postings are attributed to companies.	Amazon Web Services
Affiliate	A company's operations in one country. Our unit of observation.	Amazon Web Services in the United States
Parent	The company at the top of the ownership chain. A company with no owner is its own parent.	Amazon.com, Inc.
Corporate group	All companies with the same parent.	Amazon.com, Inc., Amazon Web Services, Whole Foods Market, and the rest of Amazon
Foreign affiliate	An affiliate whose parent is headquartered in another country. Our estimation sample.	Amazon Web Services in the United Kingdom

3.2 Sample Construction

The unit of observation is an affiliate. We first describe the restrictions that define the panel of affiliates, then the subset of foreign affiliates on which we estimate.

Affiliates We keep affiliates with at least 50 positions in October 2022, the month before ChatGPT's launch, so that occupation and seniority shares do not swing with a handful of hires. Second, the affiliate must have at least one active position in every month from January 2021 through March 2026, which rules out entry and exit. Third, its October 2022 positions must span at least three distinct occupations, so that composition has room to shift.

Companies The affiliate's company must have posted at least one job advertisement of any kind before November 2022 and at least one thereafter. This is to ensure that the company remains active on the platform.

Countries To ensure reasonable measures of AI adoption rates by country in the descriptive statistics, we keep countries with at least 50 distinct affiliates that posted a verified AI job in that country. This drops Kenya, Nigeria, Latvia, Estonia, Slovenia, Costa Rica, Ethiopia, Iceland, Russia, and Iran; in Russia and Iran, LinkedIn is officially blocked.

Foreign Affiliates After the above restrictions, the resulting panel contains 395,638 affiliates across 41 countries at monthly frequency from January 2021 through March 2026. This is the sample we use for descriptive statistics about AI adoption rates unless otherwise noted. The estimation sample used to estimate the effect of AI adoption on employment is the subset of 100,813 foreign affiliates, whose parent is headquartered in a different country. We filter to foreign affiliates to ensure that the instrument, which is based on the adoption rate of peers in a parent's metropolitan area, is well-defined and plausibly satisfies the exclusion restriction. The metropolitan area of the parent's headquarters must contain at least ten multinational parents, so that the instrument rests on at least nine peers. The instrument is discussed further in Section 4.

3.3 Measuring Adoption of Generative AI

We measure a company's adoption of generative AI from the text of its job postings in two stages. In the first stage, we flag all posts that contain generative-AI terms, including model and product names (ChatGPT, Claude, Gemini, Copilot, Llama, Mistral, Stable Diffusion, Midjourney, DALL-E), core vocabulary (large language model, generative AI, GPT, prompt engineering), and developer tooling (LangChain, retrieval-augmented generation, vector database, OpenAI API, Azure OpenAI).[11] This yields more than 2 million candidate postings between January 2021 and March 2026. In the second stage we pass every candidate through a large language model (Llama 3.1 8B Instruct) to judge whether the advertised job involves generative AI, whether by using tools such as ChatGPT, by building with large-language-model frameworks, or by listing experience with such tools as preferred. This filters out stray keyword matches (an airline's copilots, applicants named Claude, "LLM" as a law degree), postings that involve conventional machine learning rather than generative AI, and ads that mention the technology only to restrict it, for instance to bar the use of ChatGPT in applications. This leaves 1.36 million postings, 99 percent of them posted after ChatGPT's launch. We label a company an *adopter* if it has at least one AI posting after November 1, 2022. The adopter label therefore indicates that a company has brought generative AI into some part of its work. It does not indicate how this affects demand across workers, which is the focus of our estimates in Section 5.

Although the screening vocabulary includes English terms, most keywords are model and product names or technical acronyms used internationally, such as ChatGPT, Copilot, GPT, and LLM. These terms should help detect AI-related requirements across posting languages. Moreover, the observed country rankings do

[11] Appendix A6 gives the full keyword list and the verification prompt.

not suggest a systematic advantage for English-speaking countries, as several non-English-speaking countries have higher AI-posting shares than the United States or the United Kingdom. One small data caveat is that the AI posting data for the United States only runs to October 2025, which means that we miss a small proportion of late adopters.

3.4 Descriptive Statistics

We first report statistics about AI adoption across the broader set of 395,638 affiliates.[12] These affiliates come from 41 countries, with monthly observations from January 2021 through March 2026. Countries range from advanced economies such as the United States, Germany, and Japan to middle-income and developing countries such as Brazil, Turkey, South Africa, India, Indonesia, and Egypt. The median sample affiliate has 112 positions at baseline, and together the sample affiliates hold 154 million positions in October 2022. The United States contributes the most affiliates (123,074), Luxembourg the fewest (489). Table A6 in Appendix A7 reports posting counts and AI shares by country.[13]

Figure 1 ranks countries by the AI share of their postings.[14] India tops the ranking ahead of Israel (0.53 percent), Slovakia (0.43 percent), and Luxembourg (0.36 percent), with Lithuania (0.32 percent) and Poland (0.31 percent) close behind; the United States ranks tenth of the 41 countries at 0.17 percent.[15] The AI share of

[12] This section reports statistics for the broader set of affiliates to show how the prevalence of AI job postings varies across industries and across countries. Throughout, an affiliate is treated if its company posted an AI job in any country, the treatment definition of Section 4. These statistics complement a small but growing set of reports on AI adoption around the world (Bick et al., 2024; Chatterji et al., 2025; Appel et al., 2025; Eurostat, 2025; Yotzov et al., 2026; Klein Teeselink and Carey, 2026). We define treatment at the company level because companies may adjust operations globally in response to adoption in some affiliates.

[13] Table A7 in Appendix A7 reports affiliate-level counts, baseline characteristics, and treatment shares by country.

[14] These counts cover the full universe of postings in the 41 sample countries, absent sample restrictions on the affiliates.

[15] Microsoft's telemetry-based estimate of the share of 15-to-64-year-olds who used a generative AI product in the first quarter of 2026 is 17.8 percent worldwide, 63.4 percent in Singapore, 42.2 percent in the United Kingdom, 38.1 percent in Israel, 31.3 percent in the United States, 22.5 percent in Japan, 19.1 percent in Brazil, and 17.6 percent in India (Microsoft AI Economy Institute, 2026). The Anthropic Economic Index places per-capita Claude use highest in Israel and Singapore, with the United States at 3.6 times and India at 0.3 times the level its working-age population would predict, although India is second only to the United States in total use (Appel et al., 2025). OpenAI's ranking of countries by ChatGPT messages per working-age person shows the same pattern, with Singapore first and India in the lower half (OpenAI, 2026). Israel and Singapore rank near the top in both our data and these sources, but India ranks first here and near the bottom in population use, while the United Kingdom, Japan, and Brazil rank far higher in population use than in AI hiring. The posting share used in this paper measures hiring for roles that involve AI rather than use by the population, and the sample reflects employers who advertise

postings has risen steadily and more than quadrupled between the end of 2023 and March 2026.

This ranking reflects which entities advertise on LinkedIn as much as how intensively they adopt AI. The AI share of postings differs by a factor of fifty across industries (Figure 2), so a country whose postings come disproportionately from software and IT services such as India will rank high even if its companies adopt at the same rate as similar companies elsewhere. In countries such as the United States, coverage reaches roughly two thirds of official employment, and the posting sample resembles the wider economy far more closely (Section 3.5). We therefore interpret Figure 1 as describing where AI hiring is happening on LinkedIn rather than as a measure of national adoption rates.[16]

Across affiliates, treatment lines up with the selection in our model. Larger and more exposed affiliates are more likely to be treated, that is, to belong to a company that adopts, and so are affiliates that pay higher wages (see column 4 of Table 2). A doubling of employment raises the probability that the affiliate is treated by 6 percentage points, against a sample mean of 15.4 percent. A one standard deviation (0.10) increase in the workforce's average exposure score adds 2.8 points, and a 10 percent higher average wage adds 3 points. The junior share of employment matters least: holding size, exposure, and pay fixed, a 10 percentage point higher junior share raises the probability by about 1 point.

AI jobs constitute only 0.60 percent of job postings in adopting companies. The median adopting company posted three AI jobs by March 2026, and one in four posted one (Table A8). AI jobs cluster in a few occupations, as four in five adopting companies posted at least one AI job in computer and mathematical occupations, followed at a distance by management (20 percent of adopting companies), business and financial operations (18 percent), and architecture and engineering (15 percent).[17] AI jobs are typically in more exposed occupations than other postings, and have an average exposure score of 0.62 against 0.49 for others. Approximately 58 percent of AI jobs fall in a highly exposed occupation group against 26 percent of other postings. For 94 percent of AI jobs, the affiliate was already posting in that occupation group before ChatGPT's release. AI postings are concentrated across

on LinkedIn.

[16] In the U.S. Census Bureau's Business Trends and Outlook Survey, 19.8 percent of firms reported using AI in May 2026, led by information (39.7 percent) and finance and insurance (33.9 percent) (U.S. Census Bureau, 2026; Bonney et al., 2026). In Eurostat's 2025 survey, 20 percent of EU enterprises with ten or more employees used AI, led by information and communication (63 percent) and professional, scientific, and technical activities (40 percent), with construction lowest at 11 percent (Eurostat, 2025). Our posting shares also put information first (0.63 percent) but place education and health care second (0.21 percent) and professional and business services fifth. Our adoption metric measures something different from these surveys. Surveys ask respondents whether their firm uses AI at all. Our measure records how this gets reflected in job postings.

[17] In Section 5.2.2, we use this variation and estimate treatment effects separately by the occupation of the company's first AI job.

Figure 1: AI share of job postings by country

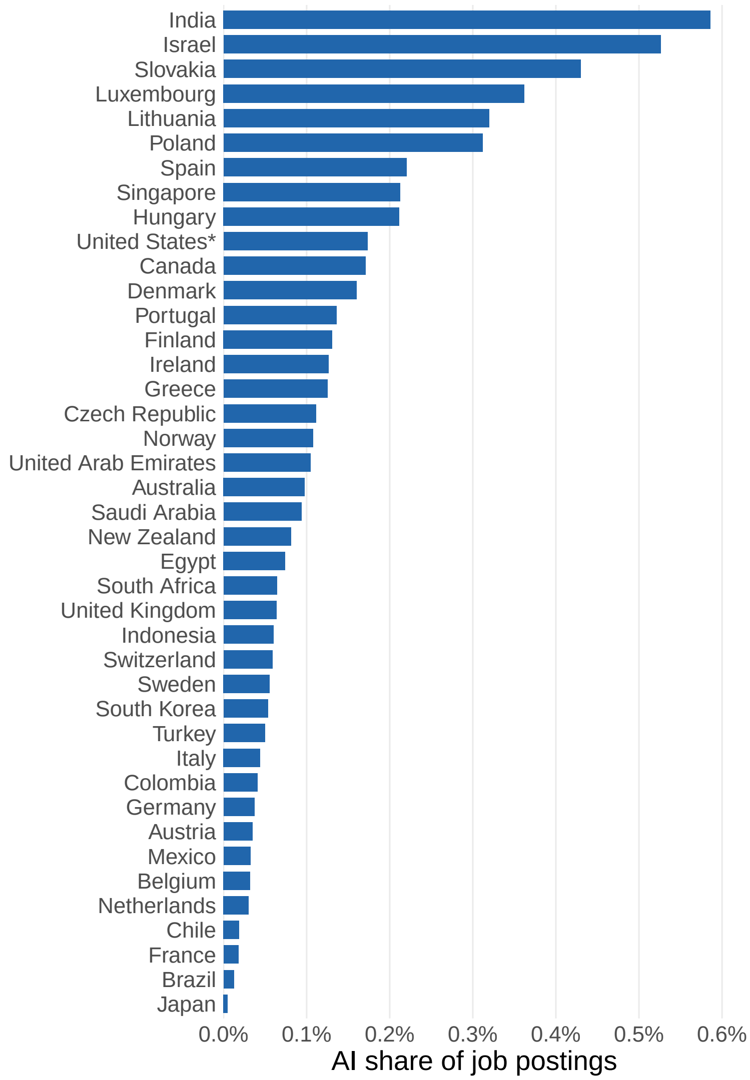

Notes: The figure ranks countries by the share of their job advertisements that involve generative AI. For each country, the bar reports LLM-verified AI job advertisements as a percentage of all job advertisements posted between November 2022 and March 2026. Postings come from the full Revelio job-listings universe in the country, with no affiliate-level sample restrictions, and each posting is counted once, on the date it first appears. The country sample is discussed in Section 3. It requires a minimum number of distinct AI-posting affiliates in a country and excludes Russia and Iran, where LinkedIn is officially blocked. *United States: posting data end in October 2025, so its AI and total postings are both counted through that month. Table A6 reports the underlying posting counts.

companies as well as occupations: the ten largest AI job posters, mainly online tutoring and data-annotation platforms, account for 40 percent of all AI postings.

Figure 2: AI share of job postings by industry

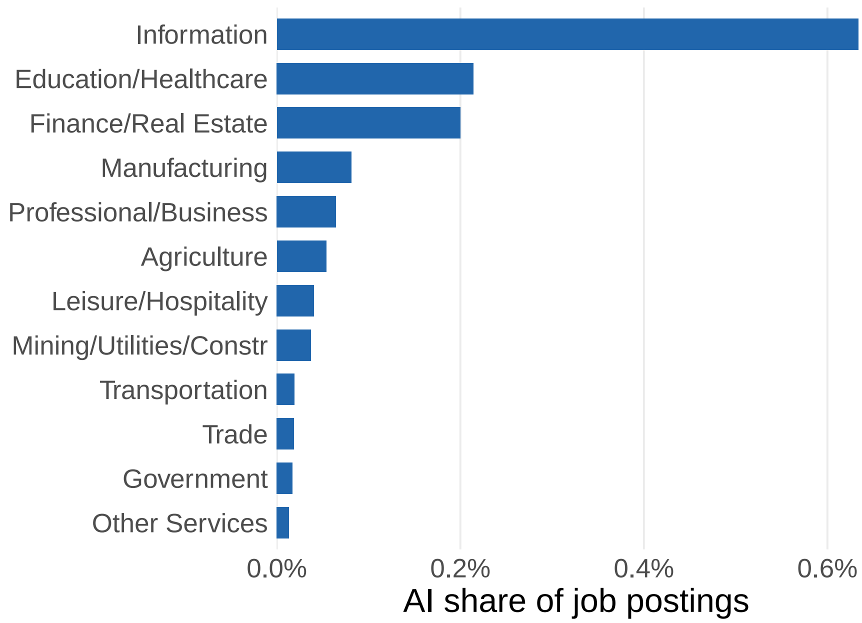

Notes: The figure ranks broad industry sectors by the share of job advertisements that involve generative AI, pooling all countries. For each sector, the bar reports LLM-verified AI job advertisements as a percentage of all job advertisements posted between November 2022 and March 2026. Sectors aggregate Revelio's industry classification into the broad groups used throughout the paper; postings that cannot be matched to a company with an industry code are excluded. United States postings are counted through October 2025. For the posting universe and the country sample, see Figure 1.

Because we treat adoption as binary, this concentration does not affect our estimates.[18]

Taken together, treated affiliates are larger, more exposed, and better paid than untreated affiliates in the same country and sector. This selection is why our empirical analysis compares treated affiliates to matched controls rather than to average affiliates and requires an instrument for treatment.

3.5 Representativeness

One may be concerned that LinkedIn profiles provide a non-representative sample of workers. To measure the degree of the problem, we benchmark Revelio against official labor force statistics at the one-digit occupation level (see Appendix A8 for the full analysis). Coverage varies across countries, with active Revelio positions roughly two-thirds as large as official employment in the United States and the

[18] Section 5.3 shows additional results for an intensity-of-treatment measure.

Table 2: Affiliate-level predictors of treatment

	Treated (0/1)			
	(1)	(2)	(3)	(4)
Log employment	0.091***	0.091***	0.093***	0.085***
	(0.006)	(0.006)	(0.006)	(0.004)
Average AI exposure		0.681***	0.505***	0.279***
		(0.072)	(0.064)	(0.024)
Junior share			−0.267***	0.105***
			(0.054)	(0.036)
Log average wage				0.336***
				(0.083)
Country FE	Yes	Yes	Yes	Yes
Sector FE	Yes	Yes	Yes	Yes
Outcome mean	0.154	0.154	0.154	0.154
Observations	395,554	395,554	395,554	395,554
Adjusted R^2	0.145	0.165	0.176	0.213

Notes: The table reports linear probability models of treatment on affiliate characteristics measured before ChatGPT's launch, adding one predictor at a time. The outcome is an indicator for the affiliate being treated, that is, for its company having posted at least one LLM-verified AI job advertisement in any country between November 2022 and March 2026. The sample is the affiliates satisfying the sample restrictions of Section 3 in every analysis country, including the United States, whose shorter adoption window ends in October 2025 and is absorbed by the country fixed effects. All predictors are measured in October 2022: log employment; the workforce's average AI exposure, defined as the occupation-level GPT-4 exposure score of Eloundou et al. (2024); the junior share of employment, defined as Revelio seniority levels 1–2; and the log of the affiliate's average imputed wage. *Outcome mean* is the mean of the dependent variable across sample affiliates. Standard errors, in parentheses, are clustered by country. */**/***: $p < 0.10/0.05/0.01$.

United Kingdom, a quarter in Brazil and Germany, and 8 percent in India. The occupational mix mirrors official shares well in high-coverage countries (a correlation of 0.73 in the United States and 0.87 in the United Kingdom), but relatively poorly in most low- and middle-income countries. In general, the data oversamples managers and professionals, and under-represents people working in agriculture, construction, and manual labor. Revelio employment also grows faster than official employment in most occupation groups, which is likely caused by platform growth rather than increased labor force participation.

These measurement facts are important inputs into our methodological design. Aggregate employment levels are confounded by LinkedIn coverage and entry, and this varies by occupation, seniority, country, and industry in ways that might confound aggregate analysis. Accordingly we do not study trends in aggregate employment outcomes in this paper. All estimation instead compares treated and control entities in the same country, sector, and month and matches on baseline size, AI exposure, and the junior share of employment. This ensures that, if the exclusion restriction of the instrument holds conditional on the matching, then coverage differences across countries and platform growth over time do not confound the internal validity of our results since representativeness is conditionally exogenous.[19]

4 Empirical Methodology

This section sets out our empirical methodology. Each observation in the analysis is an affiliate-month. We start by defining treatment and control, then describe approaches for addressing endogeneity, and finally set out the estimating equations. We conclude with a discussion of the assumptions under which the design identifies a causal effect of company adoption.

4.1 Treatment and Control

We measure adoption from job advertisements, as described in Section 3.3. We define an affiliate as treated if the associated *company* is an adopter, whether or not the AI job was advertised in the affiliate's country. A company that adopts AI may endogenously change its global operations, such as by shifting more customer service operations to AI chatbots. Consequently, we measure adoption at the company level rather than at the foreign affiliate level in our main specification.[20] An affiliate is a candidate for the *control* group if no company in its corporate group is an adopter. Affiliates whose own company did not adopt but whose corporate group contains an adopter are neither treated nor control and are excluded, since the group's adoption may affect them. In a robustness check we instead define treatment at the level of the parent's own company record, so that every affiliate of an adopting parent is treated.

Throughout the analyses, we date treatment at November 2022, when ChatGPT was launched, for every treated affiliate, rather than at the month of the company's

[19]Likewise, workers may revise their profiles in the years after their job ends (Bloom et al., 2026). Our estimation strategy requires that this behavior is exogenous given the matching and instrument.

[20]Only 30 percent of treated affiliates advertised an AI job in their own country. A variant treats an affiliate only when its company advertised in the affiliate's own country. Such affiliates tend to be part of large multinationals, especially in the technology sector, with poor matching candidates. Consequently the treated affiliates under that definition grew faster than their matched controls before ChatGPT's launch, so the variant fails the pre-trend test. Using the same sample of affiliates but with treatment defined at the company level similarly fails pre-trends, even as pre-trends are null with the same treatment definition in our main sample. This implies that restricting to foreign affiliates with local AI postings leads to a poor pool of counterfactual affiliates.

first AI job advertisement. In Section 5.3.2 we split treated affiliates by the date of the company's first advertisement and show heterogeneity in results.

4.2 Matching and Instrumental Variables

As our model and data show, AI adoption is endogenous, with larger, more profitable, and more AI exposed entities more likely to adopt the technology. As such, treated and control affiliates differ in many observed and unobserved characteristics that may correlate with both AI adoption and changes in workforce composition, and a simple comparison would combine the effects of treatment and selection. In addition, company shocks such as funding rounds, IPOs, R&D discoveries, and product launches might affect both sides of the equation in a way that will not be captured by any kind of fixed effects.

To separate treatment from selection, we proceed in two steps. Both steps exploit the structure of multinational corporate groups, so we restrict the sample to foreign affiliates. We first describe the matching procedure and the instrument. The next sections set out the estimating equations and discuss the assumptions under which the design identifies a causal effect of adoption.

Matching Equation (10) in the model predicts that adoption is a function of firm size and employment composition. In addition, Proposition 1 requires firms to operate in common labor and product markets. Accordingly, we match treated affiliates along these dimensions. Within each country, we match each treated foreign affiliate to up to five control foreign affiliates in the same 12-sector industry, selected by Mahalanobis distance on the affiliate's log employment, average AI exposure of its workforce, and junior share of employment, all measured in October 2022. We match with replacement and require each control to lie within half a standard deviation of the treated affiliate on all three variables; treated affiliates with no control inside these calipers are dropped. When a treated affiliate is matched to k controls, the treated affiliate receives weight one and each control weight $1/k$. Table A10 in Appendix A9 gives summary statistics for the matching sample, including the raw number of affiliates entering the matching procedure, the number of treated affiliates dropped by the calipers, and the number of controls per treated affiliate. The table also shows how the matching variables and the outcome variables compare between treated and control affiliates in the full sample and the matched sample.

We match only on baseline levels in October 2022, not growth rates prior to that. This matches the predictions of the model. Further, it helps us validate that our results are not driven by mean reversion or an Ashenfelter dip (Ashenfelter, 1978). Matching only on baseline levels makes satisfying pre-trends a more demanding and sharper test of the identifying assumptions.

Instrument. Even conditional on matching, treatment might be correlated with unobserved characteristics of affiliates in ways that affect employment. We thus instrument for treatment status. The instrument is based on two channels of technology diffusion. First, when firms around the parent's headquarters integrate generative AI, the parent gains information and working examples, and the technology becomes more salient, which make the parent more likely to adopt (Cai and Szeidl, 2018; Kalyani et al., 2025). Second, corporate practices and technology adoption decisions often flow from parent companies to their affiliates (Bloom et al., 2012; Guadalupe et al., 2012; Bilir and Morales, 2020). Peer adoption around the parent's headquarters can therefore drive treatment at the company and its affiliates.

We use this idea to construct an instrument for adoption. Let $p(f)$ denote affiliate f's parent, $a(f)$ the metropolitan area of that parent's headquarters, and $\mathcal{P}_a$ the set of multinational parents headquartered in a. The instrument is the leave-out adoption rate among the other parents in the same metropolitan area,

$$Z_f \;=\; \frac{1}{|\mathcal{P}_{a(f)}| - 1} \sum_{p \in \mathcal{P}_{a(f)},\, p \neq p(f)} A_p, \tag{14}$$

where A_p equals one if any company in parent p's corporate group is an adopter. Because the sum excludes $p(f)$, postings by the affiliate's own parent, or by any other company in its group, do not enter Z_f. We compute Z_f only in metropolitan areas with at least ten multinational parents, so that it never rests on fewer than nine peers. We standardize Z_f to mean zero and unit variance over the matched estimation sample, so that first-stage coefficients measure the effect of a one-standard-deviation increase in the peer adoption rate. The country-level analyses of Section 5.2.1 standardize within country.

4.3 Estimation Approach

We estimate our instrumented matched difference-in-differences model in long differences. Formally, let q measure months relative to the launch of ChatGPT in November 2022. The panel runs from $q = -22$ (January 2021) to $q = 40$ (March 2026), and October 2022 ($q = -1$) is the baseline from which we take differences. For employment outcomes, we use $\log(\max\{N, 1\})$, where N is the relevant employment count, retaining observations with zero junior or senior employment. For affiliate f in matched set m and outcome y, the long difference is $\Delta_q y_{fm} = y_{f,q} - y_{f,-1}$. For every month $q \neq -1$ we estimate, by two-stage least squares, the cross-sectional regression

$$\Delta_q y_{fm} \;=\; \beta_q \, D_f \;+\; \mu_{mq} \;+\; \varepsilon_{fmq}, \tag{15}$$

with first stage

$$D_f \;=\; \gamma_q \, Z_f \;+\; \kappa_{mq} \;+\; u_{fmq}, \tag{16}$$

where D_f indicates that affiliate f's company posted at least one AI job from November 2022 onward, in any country, and μ_{mq} and κ_{mq} are matched-set fixed effects, so that both stages compare affiliates only within matched sets. We cluster standard errors by the metropolitan area of the parent's headquarters, the level at which the instrument varies.

The coefficients β_q trace out an event study. For $q < -1$ they are placebo long differences that test whether treated and control affiliates were on parallel paths before the introduction of LLMs; for $q \geq 0$ they give the treatment effect in each month after October 2022. We summarize post-period effects with the coefficient at March 2026 ($q = 40$).

Our headline estimates pool all 41 countries in a single regression. Because matched sets are formed within country and sector, the matched-set fixed effects absorb any shock common to a set, including country- and sector-specific shocks over the window. The pooled coefficient is a weighted average of within-country effects. The effect of adoption need not be the same in the United States as in India, so we also estimate (15) country by country and combine the endpoint estimates ($q = 40$) in a random-effects meta-analysis (DerSimonian and Laird, 1986), which lets us examine the distribution of effects around the world.

The long-difference specification matters for interpretation. Standard 2SLS estimators do not in general estimate a positively weighted average of treatment effects unless the specification is fully saturated in the covariates (Blandhol et al., 2026). When covariates enter additively, as the affiliate and month fixed effects of a panel event study do, some treatment effects can receive negative weights, so that one could estimate a negative effect even if adoption raised the outcome at every affiliate. Long differencing avoids this problem. Differencing from October 2022 removes the affiliate effect by construction, and the only covariates left in each monthly cross-section are matched-set indicators, which is a saturated specification. Under standard LATE assumptions, each monthly coefficient is therefore a positively weighted average of matched-set effects for compliers, the affiliates whose treatment status responds to the peer environment around their parent's headquarters. Under this interpretation, the estimates capture effects in affiliates on the margin of treatment, who are pushed into the treatment group because their parent's peers adopted AI.

4.4 Identification Assumptions

For the instrument to be valid, peer adoption around the parent's headquarters must be unrelated to the other determinants of the affiliate's workforce. Our design rules out several sources of confounding. First, because the instrument is built from other parents' adoption decisions, it does not respond to shocks to the affiliate's company, such as funding rounds or product launches. Second, because the peers sit in a different country from the affiliate, the instrument is also insulated from shocks to the affiliate's local market. Third, matching within country and sector on baseline

AI exposure controls for shocks that hit a whole sector in a country symmetrically, and for factors that track AI exposure itself, such as interest-rate changes or the scope for remote work (Dingel and Neiman, 2020; Lambert and Schindler, 2026).

One remaining concern is that a shock to the parent's home market could affect both peer adoption and the affiliate's outcomes. A boom in the metropolitan area of the parent's headquarters could raise peer adoption, because firms have more cash to spare for AI investment, and at the same time raise the parent's cash flow, which funds expansion abroad. This pattern predicts that the instrument should raise overall affiliate employment, whereas we find only a modest and imprecise increase. It may nonetheless bias our estimates. We address the concern in two ways. First, in a robustness exercise we add the employment growth of the other multinational parents headquartered in the same metropolitan area as a 'bad' control in both stages, which absorbs any generic expansion around the headquarters. Because peer growth partly reflects the peers' own adoption, this control can absorb part of the true effect, so the check is conservative. Second, we consider robustness to using an alternative instrument that uses a peer group of multinational parents in the same headquarters *country and sector* as the focal parent rather than in the same metropolitan area. This is a different source of variation that may be less subject to shocks local to the parent's metropolitan area. Section 5.3.1 reports these checks.

The exclusion restriction requires, in addition, that the parent's peer adoption affects the affiliate's workforce only through the company's adoption. Most direct channels are less plausible, since the peers operate in a different country and in different product and labor markets from the affiliate, but three are harder to dismiss. The first is the parent's own adoption. A parent that integrates generative AI may centralize exposed tasks at headquarters and reduce its affiliates' junior ranks whether or not the company adopts. Our estimates would then capture the effect of the corporate group's adoption on the affiliate rather than of its company's adoption. To gauge it, we compare our main estimates with estimates in which adoption is measured by the parent's own postings rather than the company's. Two further channels arise because adoption by nearby firms changes the parent's environment in ways that can reach the affiliate without any effect of adoption in its company. One is organizational imitation: local peers reorganize, and a parent may copy their practices, fewer junior hires or flatter teams, across its group irrespective of AI. As a robustness we rebuild the instrument from peers outside the parent's sector, under the idea that firms imitate these organizational practices mainly within their industry; if organizational imitation drove the results, they should weaken once same-sector peers are removed, lending more credence to AI adoption as the relevant channel. The other is the local labor market: wide AI adoption near headquarters may bid up or down workers' wages, which may affect the parent's employment decisions abroad. In that case the instrument moves the affiliate's workforce through relocation rather than AI. As discussed below, OLS and synthetic difference-in-differences estimates that do not rely on the cross-border instrument are less subject to this concern.

Last, our approach requires monotonicity: peer adoption around the parent may make an affiliate's company more likely to adopt or leave it unaffected, but it should not make it less likely (Imbens and Angrist, 1994). To examine this assumption, we estimate the first stage separately by parent sector and by affiliate size and check that no subgroup responds negatively to the instrument.

Despite these checks, we acknowledge our instrument is not perfect, and the exclusion restriction remains untestable. We therefore complement the IV with two alternative designs. The first is a matched difference-in-differences event study, which applies our matching procedure and then compares adopters with their controls directly, without instrumenting adoption. The second is synthetic difference-in-differences (Arkhangelsky et al., 2021), which replaces the nearest-neighbor matches with estimated weights on control affiliates and pre-treatment months, chosen so that the synthetic control group tracks the treated group before November 2022. Both designs use the same treatment definition, timing, eligibility rules, and outcomes as our main analysis but a broader sample: because neither requires the cross-border instrument, we estimate both on every eligible affiliate in the 41 analysis countries, domestic and foreign affiliates alike. The identifying assumptions differ as well: the event study relies on parallel trends within matched sets, and synthetic difference-in-differences on parallel trends after reweighting. Appendix A14 provides full methodological details for both designs. The OLS design in particular is more vulnerable to confounding by unobserved affiliate shocks, so we interpret it as a complementary robustness exercise rather than as a replacement for the IV.

5 Results

This section shows the main results for the effect of company AI adoption on affiliate-level workforce composition. The goal of our empirical exercise is to examine for which groups of workers and which occupations AI has been labor saving, and for which it has been labor expanding, to date. We furthermore examine heterogeneity in these effects across countries, and show how the estimates change under alternative identification strategies. We end with a model-based decomposition of the estimated employment changes into scale and bias terms.

Before turning to the treatment effects, we first establish that the instrument is relevant: affiliates whose parents are surrounded by adopters are indeed more likely to be treated, that is, to belong to a company that adopts. Table 3 shows that a one-standard-deviation increase in the peer adoption rate raises the probability that the affiliate's company adopts by 16 percentage points, which gives an effective first-stage F statistic of 104.2, well above conventional weak-instrument thresholds. In other words, the pattern is consistent with generative AI diffusing along corporate networks, from the firms around a parent's headquarters to that parent's affiliates abroad. Section 5.3.1 provides supporting evidence, showing that the first stage is positive in every parent sector and affiliate size tercile, in line with the monotonicity assumption, and it stays strong when we exclude any of the ten metropolitan areas that contribute the most treated affiliates.

5.1 Workforce Composition

Starting with the effect of AI adoption on workforce seniority, Panel A in Figure 3 plots the effect of adoption on junior employment as a fraction of an affiliate's total employment. Before the introduction of ChatGPT, the junior share followed similar trajectories in treated and control affiliates. After the introduction of LLMs, however, the junior share begins to decline almost immediately, and continues to gradually decline for three years. By March 2026, the junior share of employment in treated affiliates decreased by 1.9 percentage points compared to control affiliates, as shown by the end-period estimates in Table 3. Juniors hold 57.1 percent of positions at matched treated affiliates at baseline (an employment-weighted share), so the 1.9-point drop is equivalent to 3.3 percent of that baseline share.

A falling junior share can reflect fewer juniors, more seniors, or both. Panels B and C estimate the effect of adoption on each group separately. Junior employment at treated affiliates falls quickly after the launch of ChatGPT, and decreases by 2.9 percent relative to matched controls within a year, although the estimates are somewhat imprecise. The gap then stays constant around that margin. Senior employment, by contrast, barely moves for the first year and a half, and then climbs steadily, ending 6.7 percent higher at treated affiliates. The decrease in the junior share is therefore primarily driven by the increase in senior employment; junior employment is lower at treated affiliates, but not significantly so. To put these magnitudes in terms of people: the average matched treated affiliate has roughly 255 positions in October 2022, 145 of them junior. The point estimates imply that by March 2026 it employs about 8 more seniors than it would have without adoption, and about 4 fewer juniors.

To better understand the relative decline in junior employment, we next examine the extent to which the change in junior positions tend to be concentrated in high AI-exposure occupations. Our model shows that such a concentration would be consistent with the interpretation that AI is labor saving for high-exposure junior roles in particular. We therefore estimate the effect of AI adoption on the difference between two shares of junior employment: the share in the most exposed occupations (top quartile of exposure) and the share in the least exposed occupations (bottom quartile). By taking the difference of the two shares, we investigate shifts in exposure conditional on total junior employment, which allows us to abstract from all factors that move junior employment as a whole, including affiliate-level scale effects and overall shifts between juniors and seniors. In other words, the estimate tells us how AI adoption reallocates juniors across occupations, holding fixed how many juniors

Figure 3: Event study plots for the effect of AI adoption on junior versus senior employment

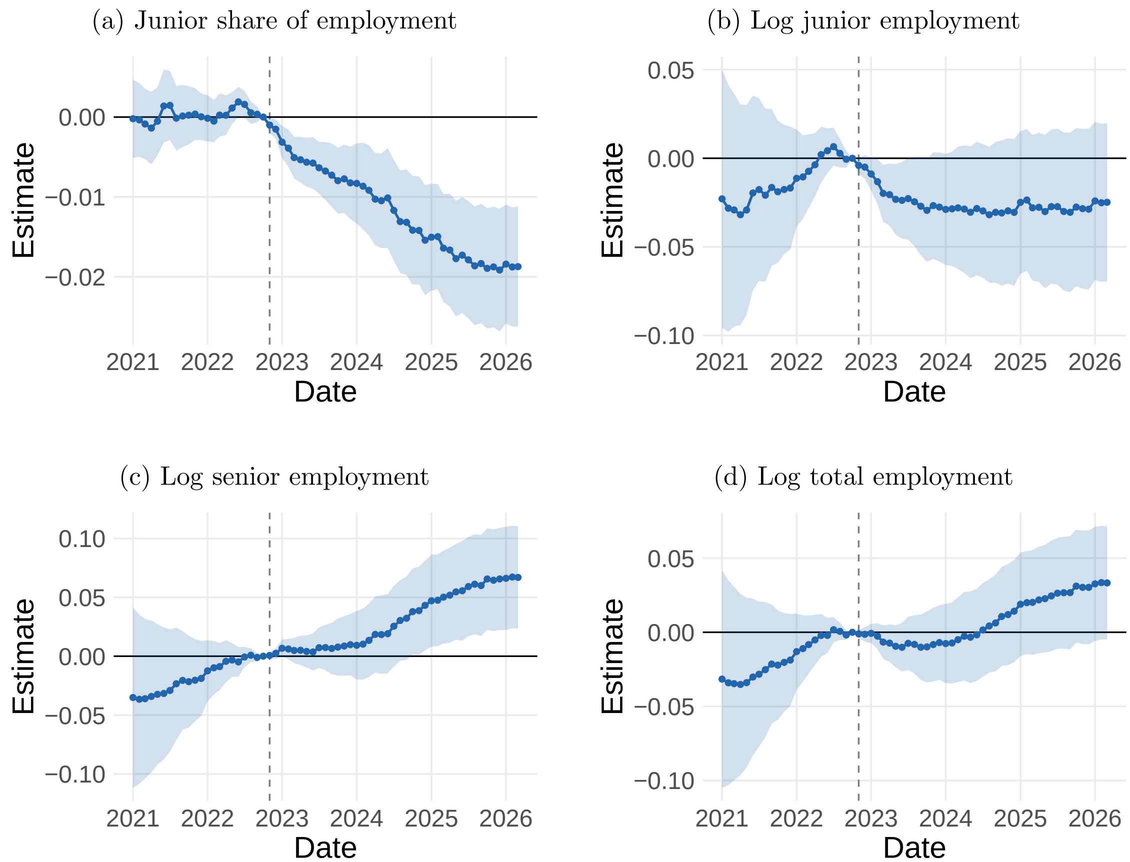

Notes: The figure plots the estimated effect of treatment, adoption of generative AI by the affiliate's company, on the seniority mix of the affiliate's workforce, month by month. Each point is a two-stage least squares coefficient β_q from Equation (15), in which the outcome enters as the long difference between the month on the horizontal axis and October 2022, the month before ChatGPT's launch. The sample is the matched sample of Section 4: foreign affiliates of multinational groups whose company posted at least one AI job advertisement from November 2022 onward, each paired with control foreign affiliates matched within country and sector on baseline employment, average AI exposure, and junior employment share. Treatment is instrumented by the leave-out AI adoption rate among the other multinational parents headquartered in the same metropolitan area as the affiliate's own parent, and match fixed effects restrict every comparison to affiliates within the same matched set. Panel A plots the junior share of employment, Panel B log junior employment, Panel C log senior employment, and Panel D log total employment, where junior covers Revelio seniority levels 1–2 and senior levels 3–7. Shaded bands are 95% confidence intervals, clustered by the metropolitan area of the parent's headquarters. The dashed vertical line marks ChatGPT's launch in November 2022.

Table 3: Effect of AI adoption on employment and workforce composition

	Log employment			Junior share	Q4−Q1 exposure share		
	Total	Junior	Senior		Total	Junior	Senior
AI adoption	0.033*	−0.025	0.067***	−0.019***	0.001	−0.006	0.009**
	(0.020)	(0.023)	(0.022)	(0.004)	(0.003)	(0.004)	(0.004)
Match fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Countries	41	41	41	41	41	41	41
Treated affiliates	26,102	26,102	26,102	26,102	26,102	26,102	26,102
Control affiliates	27,697	27,697	27,697	27,697	27,697	27,697	27,697
Effective first-stage F	104.2	104.2	104.2	104.2	104.2	104.2	104.2

Notes: The table reports the estimated effect of AI adoption on each column outcome at the March 2026 endpoint of the panel, that is, the coefficient β_q from Equation (15) in the final month, estimated by two-stage least squares and pooled across countries. *AI adoption* indicates that the affiliate's company posted at least one AI job advertisement after November 2022, instrumented by the leave-out AI adoption rate among the other multinational parents headquartered in the same metropolitan area as the affiliate's own parent. The sample is the matched sample of Section 4: adopting foreign affiliates, each paired with non-adopting foreign affiliates matched within country and sector on baseline employment, average AI exposure, and junior employment share, with match fixed effects restricting every comparison to affiliates within the same matched set. Sample counts describe the matched design. Affiliates are distinct company–country pairs; control affiliates reused across matched sets are counted once. Junior covers Revelio seniority levels 1–2 and senior levels 3–7; the Q4−Q1 exposure share is the difference between the share of employment in occupations in the top quartile of AI exposure and the share in occupations in the bottom quartile. The effective first-stage F is the cluster-robust F statistic of the design first stage, which is common to all outcomes and, with a single instrument, equals the Montiel Olea–Pflueger effective F. Standard errors, in parentheses, are clustered by the metropolitan area of the parent's headquarters. The corresponding estimates from the matched and synthetic difference-in-differences designs are in Table A11. */**/***: $p < 0.10/0.05/0.01$.

the affiliate employs. We then repeat this analysis for seniors.

Figure 4 shows the results. Among juniors (Panel A), the share working in the most exposed occupations falls relative to the share in the least exposed occupations, with the difference falling by 0.6 percentage points by March 2026, although the estimate is not statistically significant. Among seniors (Panel B), the shift goes in the opposite direction, as the difference between the top- and bottom-quartile shares of senior employment rises by 0.9 percentage points. Taken together, these estimates are consistent with the interpretation that AI is labor saving for junior roles and labor expanding for senior roles.

The results so far describe which workers treated affiliates employ, but they say less about whether treated affiliates grow relative to controls. Panel D of Figure 3 plots the effect of adoption on log total employment. Total employment at treated affiliates tracks that of their matched controls for the first year and a half after

Figure 4: Event study plots for the effect of AI adoption on employment shares in high- versus low-exposure occupations

Notes: The figure plots the estimated effect of AI adoption on where an affiliate's workers sit in the AI-exposure distribution. The outcome in each panel is the difference between the share of employment in occupations in the top quartile of AI exposure and the share in occupations in the bottom quartile, so that the estimates describe how adoption reallocates workers across occupations while holding fixed how many the affiliate employs. Panel A measures this contrast within junior employment, Panel B within senior employment, and Panel C across the whole workforce. We assign detailed O*NET-SOC occupations to exposure quartiles using the GPT-4 beta score of Eloundou et al. (2024). Cutoffs are the employment-weighted quartiles of this score among positions active on 1 October 2022 in the pooled profile data across all available countries, before restricting to the estimation sample. The cutoffs are common to all affiliates and countries. For the design, sample, and all other definitions, see Figure 3.

ChatGPT's launch: the junior decline and the senior increase offset each other, so the workforce is recomposed rather than resized. From mid-2024 the senior increase begins to dominate, and by March 2026 total employment is 3.3 percent higher at treated affiliates (Table 3). The estimate is marginally significant at the 10% level.

This estimate is consistent with a small increase in the workforce of treated affiliates relative to controls, but does not by itself isolate the scale effect from changes in worker-specific labor demand. We discuss estimation of the scale term and of the bias of AI across worker types in Section 5.4.

Two caveats apply. First, the matched comparison nets out any effect of AI that reaches adopters and non-adopters alike, so the estimate says that treated affiliates grow *relative* to control affiliates, not that employment rises in the aggregate; Appendix A3 shows that the two can differ in sign. Second, the design compares affiliates, not corporate groups, so growth at a treated affiliate may come at the expense of other parts of the same group. Section 5.3.1 returns to this point when adoption is measured at the level of the parent.

The second dimension of affiliate-level workforce recomposition we consider is occupation. We estimate how AI adoption affects which types of occupations people hold within those affiliates. We focus on employment shares across 22 two-digit occupation groups in our data. Using shares rather than levels allows us to isolate changes in workforce composition from changes in overall affiliate size. For each occupation in the analysis, we re-match treated and control affiliates, keeping only the affiliates that already employed at least one person in that occupation in October 2022, and adding the occupation's baseline employment share as an additional matching variable. We thus use 22 different matched samples for this analysis. We consider two outcomes, namely employment in occupation o as a fraction of total affiliate employment, and the share of juniors in occupation o.

The left panel of Figure 5 shows the results. In contrast to the junior vs. senior changes, most occupations move comparatively little. In 21 of the 22 groups, the share effect is smaller than half a percentage point in absolute value. The main exception is computer and mathematical occupations, which is also the most AI-exposed of the 22 groups. This group *grows* by 0.8 percentage points, suggesting that AI is labor expanding rather than labor saving for these occupations. Labor saving cases (occupations with shrinking shares) are slightly less apparent, and we only find a somewhat sizable decline (-0.3 p.p.) for business and financial operations, although even this decline is non-significant. Figure A10 repeats the analysis with a Bonferroni correction across the 22 groups. The gain in computer and mathematical occupations remains statistically significant.

That occupation shares do not move all that much suggests the shift from juniors to seniors happens because of shifts within occupations rather than between them. We analyze within-occupation shifts by examining how the share of juniors in each occupation changes in treated affiliates compared to control affiliates. The right panel of Figure 5 shows the results. Adoption lowers the junior share in 16 of the 22 groups. Five declines are statistically significant: legal (-4.9 percentage points), computer and mathematical (-3.3), business and financial operations (-3.2), sales (-2.6), and management (-1.0), of which computer and mathematical and business and financial operations survive a Bonferroni correction (Figure A10). The five

Figure 5: Effect of AI adoption on affiliate-level occupation shares

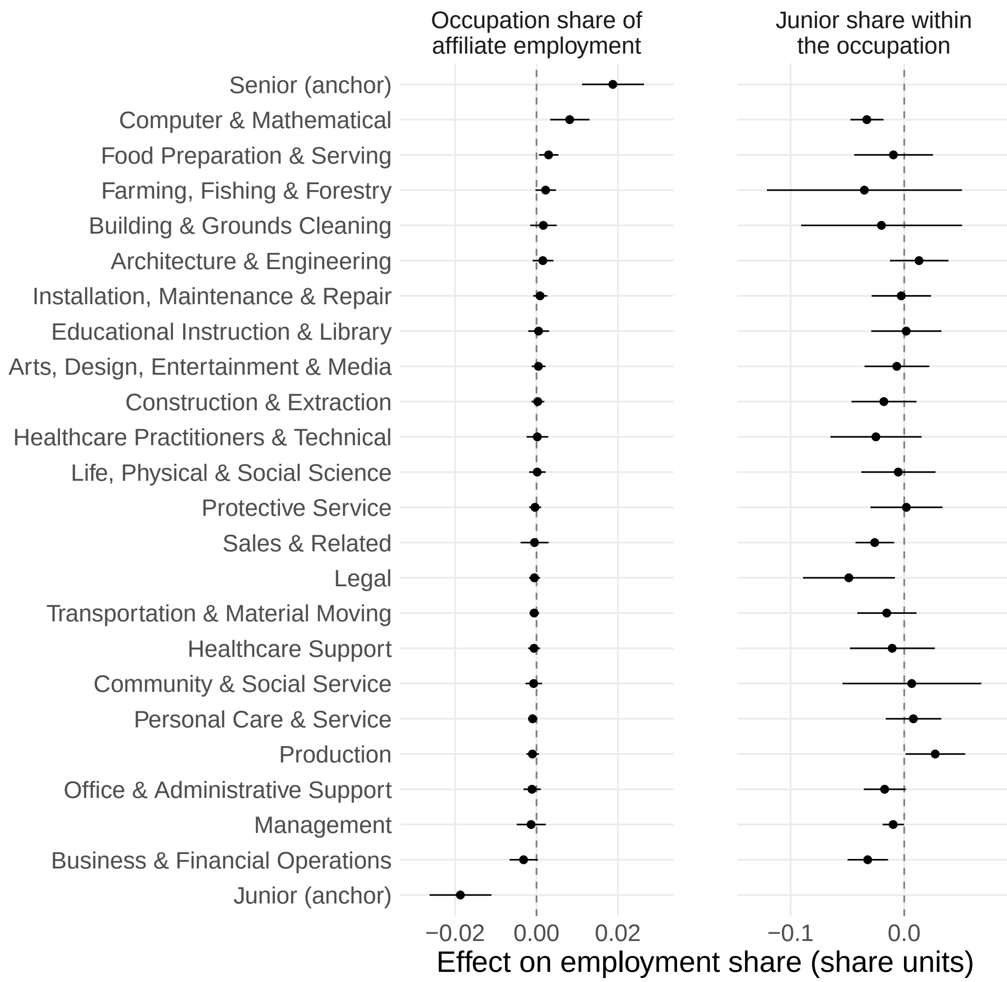

Notes: The figure ranks occupation groups by the estimated effect of AI adoption on their place in the affiliate. Each estimate is the two-stage least squares coefficient from Equation (15) at the March 2026 endpoint, in share units. The left panel plots the occupation's share of affiliate employment and the right panel the junior share of employment within the occupation, both ordered by the left panel's estimates; the junior and senior anchors have no within-occupation junior share and so appear in the left panel only. Occupations are the two-digit SOC major groups, each estimated on its own matched sample: the affiliates that employed the occupation in October 2022, re-matched with the occupation's baseline share added to the main design's matching variables. Whiskers are 95% confidence intervals, clustered by the metropolitan area of the parent's headquarters. For the treatment definition, instrument, and matched sample, see Figure 3.

occupations with significant declines are all in the upper half of our exposure ranking with three of them in the top five.

Taken together, our results show that AI is labor expanding for seniors and for workers in computer and mathematical occupations, whose employment shares rise at treated affiliates, and labor saving for juniors, lowering the relative demand for entry-level roles across professions.

5.2 Heterogeneity

The pooled estimates assign a single effect to every treated affiliate, from a Japanese manufacturer's plant in Brazil to an American software company's London office. This section explores how the effects differ. We first estimate them country by country (Section 5.2.1). We then split the pooled sample by the affiliate's baseline AI exposure, size, and sector, and by the occupation of the company's first AI job (Section 5.2.2).

5.2.1 Country-Level Heterogeneity

Our evidence so far pools a wide set of countries with different levels of development, labor market institutions, and digital infrastructure. These factors could lead to vastly different treatment effects across countries, for example because firms in high-wage economies may have more to gain from replacing junior labor, or firms with dense digital infrastructure may find the technology easier to deploy. The current section explores how the observed effects differ across these countries. This exercise also informs us about the generalizability of our findings.

We estimate our main IV routine separately for each country. We focus on the set of 31 countries with a sufficiently strong first-stage F statistic (10 or higher), since the instrument is too weak elsewhere to support a country-specific estimate. We combine the resulting endpoint estimates (March 2026) in a random-effects meta-analysis, which gives both a mean effect across countries and a measure of how widely the effects vary around it. We then correlate the country estimates with five characteristics (measured in 2022, except digital readiness in 2024 and the ICT employment share in the latest available year up to 2022): GDP per capita at purchasing power parity, digital readiness, the ICT share of employment, labor market rigidity, and ICT service exports as a share of GDP. Because the country estimates differ widely in precision, the meta-analysis weights countries by the inverse of the sum of their sampling variance and the estimated between-country variance, while the correlations use inverse sampling-variance weights. The between-country variance is estimated by restricted maximum likelihood (REML).

Figure 6 shows the results. The decline in the junior share is widespread, and we observe negative point estimates for 23 of the 31 countries. Seven of these are statistically significant, and occur in countries with vastly different institutions and levels of development (United States, the United Kingdom, Spain, Poland, Brazil, Mexico, and Saudi Arabia). A random-effects meta-analysis puts the mean effect at -1.1 percentage points, with moderate heterogeneity across countries ($I^2 = 42$

Figure 6: Effect of AI adoption on employment and workforce composition, by country

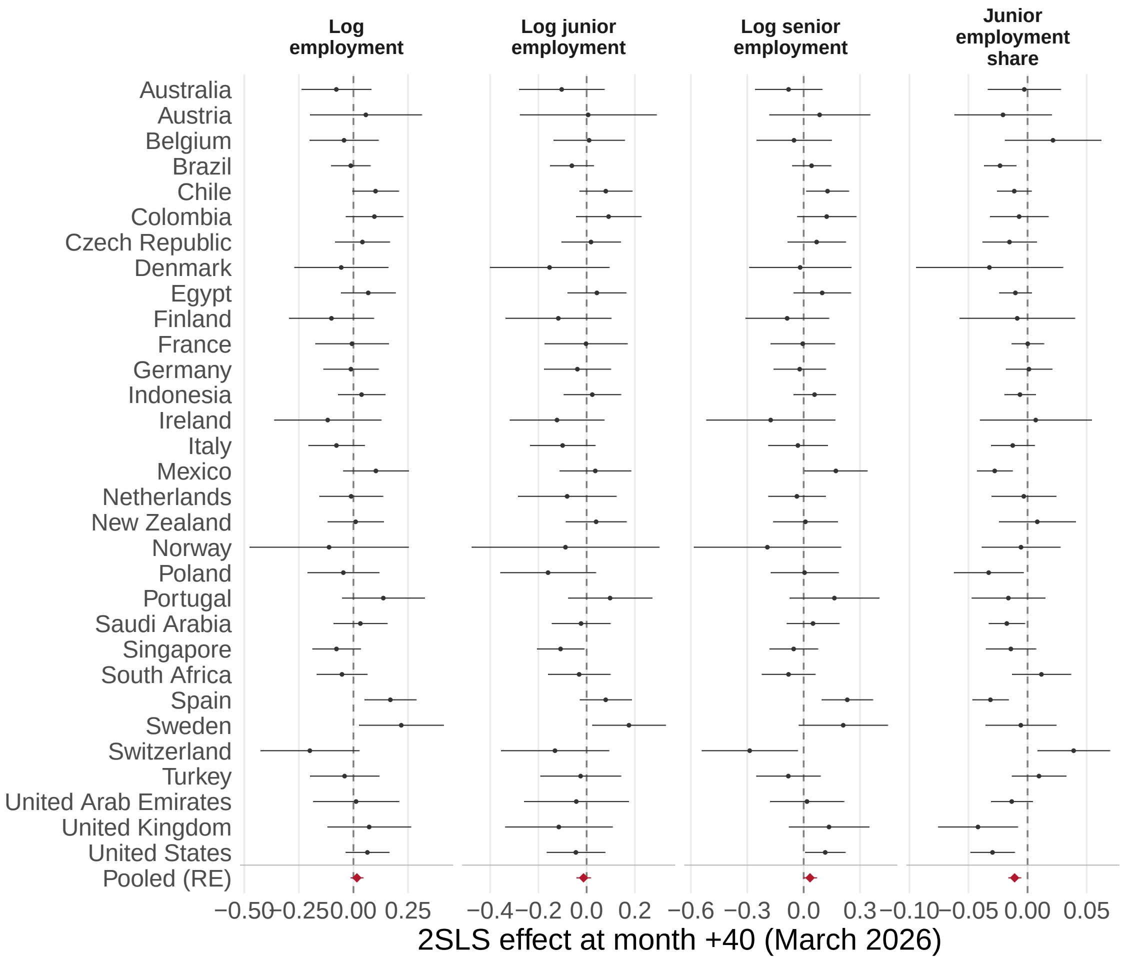

Notes: The figure shows how the estimated effect of AI adoption varies across countries. Each row reports a country-specific two-stage least squares estimate, obtained by estimating Equation (15) on that country alone and reading the coefficient at the March 2026 endpoint of the panel; the horizontal line is its 95% confidence interval. The panels report the four headline outcomes: log total employment, log junior employment, log senior employment, and the junior share of employment. Countries are listed alphabetically, and countries in which the instrument is too weak to support a country-specific estimate are omitted. The bottom row reports the random-effects meta-analysis of the country estimates, which weights each country by the inverse of the sum of its sampling variance and the between-country variance estimated by restricted maximum likelihood (REML). For the treatment definition, the instrument, and the matched sample, see Figure 3.

percent). This aligns with our pooled result, which showed a 1.9 percentage point decline in the junior share.

The other employment margins also mirror the pooled results: total and junior employment show no significant change, and senior employment rises by 3.4 percent on average, which is marginally significant. Two countries show significant total-employment gains (Sweden and Spain), and none shows a significant decline. For all three employment outcomes, the cross-country dispersion is small enough that we cannot reject a common effect (I^2 between 18 and 29 percent).

Next, we test what factors explain the differences in effects across countries, and consider whether the effects of AI adoption depend on how rich, digitized, or regulated a country is. Appendix Figures A11–A15 plot the country-level estimates against five country characteristics. Junior employment losses run somewhat deeper in richer economies ($r = -0.37$) and in more digitized ones ($r = -0.29$ for digital readiness, $r = -0.20$ for ICT service exports), as we would expect if high-wage firms have more to gain from replacing junior labor and digitized economies find the technology easier to deploy. The losses are smaller where labor regulation is more rigid ($r = +0.20$), consistent with the higher cost or difficulty of reducing junior employment.[21] Among these correlations, only one association is statistically significant, which is partly explained by the fact that the sample consists of 31 relatively imprecise country estimates. For the other outcomes (total employment, senior employment, and the junior share), we do not find meaningful heterogeneity.

Taken together, these results indicate that the recomposition we document is not just a rich-country phenomenon. The junior share falls in Brazil and Mexico by about as much as it does in the United States, the United Kingdom, and Spain, and it is negative, though smaller and imprecisely estimated, in Egypt, Colombia, and Indonesia; country characteristics such as income, digital readiness, ICT employment, and labor regulation explain little of the remaining variation in treatment effects. Given the small sample size, however, we take all evidence in this section as suggestive.

5.2.2 Additional Heterogeneity

Next, we go from country-level heterogeneity to affiliate-level heterogeneity. We consider heterogeneous treatment effects by the affiliate's baseline AI exposure, size, and sector (technology versus non-technology), and examine whether the effects differ depending on the occupation in which the AI job is posted. Each analysis re-estimates our main empirical model on a subset of the matched sample while keeping treatment, instrument, matching, eligibility, and clustering the same. In every subsample, treated affiliates keep the same matched controls as in the main analysis. We caution that for certain cuts of the data we find evidence of pre-trends in employment outcomes, so the estimates should be viewed as somewhat more

[21]Mirroring our findings, Klein Teeselink and Carey (2026) show that countries with more rigid labor markets mostly adjust by reducing new hiring.

suggestive relative to the main results in the paper.[22]

Starting with exposure, we compare treatment effects among high-exposure and low-exposure affiliates. Baseline exposure measures how many tasks within the affiliate AI can perform. Hence, higher exposure should arguably lead to larger treatment effects. To test this, we split the treated affiliates at their median baseline exposure. Figure A16 shows the results for both groups.[23] The results are almost entirely driven by the above-median group. Among these affiliates, adoption lowers the junior share by 3.4 percentage points, raises senior employment by 10.4 percent, and lowers junior employment by 6.5 percent. Below the median, the pooled coefficients are close to zero and statistically insignificant, although the computer and mathematical share still gains 0.4 percentage points (Figure A18). The computer and mathematical result also strengthens among above-median affiliates, where the occupation's employment share gains 1.3 percentage points and the junior share within it falls by 4.7 points, both larger than our main estimates (Figures A17 and A18). In other words, the affiliates whose existing work can be performed by AI account for virtually all of the recomposition we observe. This result ties the effect of technological feasibility (exposure) to that of adoption.

We next consider size, and split matched treated affiliates at their pooled median baseline employment of 122 positions (Figure A19). The junior share effect is nearly identical on the two sides, falling by 1.8–1.9 percentage points in both. Smaller treated affiliates grow after adoption, with total employment up 5.7 percent and senior employment up 8.2 percent, whereas we find a senior-employment increase of 5.0 percent for larger treated affiliates. In the occupation-level analysis, the computer and mathematical employment share gains 1.1 percentage points among smaller treated affiliates and 0.5 points among larger ones, and the junior share within the occupation falls on both sides of the split (Figures A20 and A21). Adoption thus shifts an affiliate's workforce composition toward seniors across the size distribution, while the growth effect of adoption is concentrated among the smaller affiliates.

The third split separates technology affiliates, defined as those in the information sector, from all other affiliates. The information sector has the highest AI share of job postings, three times that of the next sector, and the most exposed workforce in our data (Section 3), so the exposure logic of the first split predicts stronger effects among technology affiliates. The split also bears on a common objection to firm-level AI studies, namely that they describe the technology sector rather than the economy, or reflect the unwinding of pandemic-era overhiring, although it should be noted that all our analyses match within country and sector, so any correction that hit the whole technology sector applies to treated and control affiliates alike.

[22]Note also that scales of the dependent variables differ across figures in this section given the heterogeneous effect sizes.

[23]Both groups have a sufficiently strong first stage ($F = 116$ for below median, and 70 for above).

The results show that adoption lowers the junior share by 3.9 percentage points, raises senior employment by 14.6 percent, and raises total employment by 7.8 percent among technology affiliates. Outside of tech, the same pattern appears, but in smaller magnitudes. The junior share falls by 0.9 points, senior employment rises by 2.9 percent, and total employment remains stable (Figure A22). The same holds for the occupation-level analysis, as the computer and mathematical employment share increases by 1.6 percentage points among technology affiliates and by 0.4 points in other industries (Figures A23 and A24). The recomposition thus appears in both groups, but is larger among technology affiliates.

Our fourth heterogeneity check considers how effects differ depending on the occupation in which the company first advertises an AI job.[24] Unlike exposure, size, and sector, this analysis describes heterogeneity in the type of treatment rather than in a pre-determined affiliate trait. We separately estimate our treatment effects for each occupation as long as the occupation has a first-stage F statistic above 10. The companies of 70.8 percent of matched treated affiliates first advertise AI jobs in computer and mathematical occupations; 10 out of 22 occupations have a sufficiently large F statistic. Figure A25 shows that the junior-share effect is negative in all ten categories and significant in six. The results for total junior and senior employment (Figure A26 and Figure A27) are more mixed, with some positive and some negative estimates, most of which are insignificant. We conclude that the decline in the junior share is not a by-product of hiring AI engineers, and that affiliates of companies whose first AI job advertisement is for a lawyer, a manager, or a designer shift away from junior workers as well.

Across the four splits, the decline in the junior share appears across affiliate size, both inside and outside the technology sector, and across a number of occupation groups. As such, this result appears to be robust and widespread.

5.3 Robustness Checks

5.3.1 Instrument Validity

The validity of our IV estimates is based on the assumption that peer adoption around the parent's headquarters affects the affiliate's workforce only through whether the affiliate's company adopts. We work through a range of concerns in this section, summarizing every check by the March 2026 endpoint estimates of the main effects unless stated otherwise.

The first concern is that a local boom could raise peer adoption and the parent's cash flow at the same time, and that cash flow funds expansion abroad. To test this, we add local employment growth as a control. We measure this variable by the average employment growth between the pre- and post-ChatGPT periods of other multinational parents headquartered in the same metropolitan area as the

[24] This is based on the occupation code of the company's first classifiable AI job listing.

affiliate's parent. This is a 'bad' control, because post-launch growth is affected by local adoption rates, so the control absorbs part of the treatment effect along with any generic local boom. The check is therefore conservative, biasing our estimates toward zero. Our estimates barely move when we add this control (Figure A28 and Figure A44). The first stage stays strong at $F = 116.6$, the junior share still falls by 1.9 percentage points, senior employment still rises by 6.9 percent, junior employment falls by an insignificant 2.1 percent, and seniors still reallocate toward more exposed occupations.

A second way to address potential shocks to the parent's locality is to change the peer group itself. To do so, we rebuild the instrument from multinational parents in the same headquarters country and sector as the focal parent, rather than the same metropolitan area. The instrument is weaker, with a first-stage F of 20.1, and the confidence intervals widen accordingly, but the point estimates move if anything away from zero. The junior share falls by 2.3 percentage points and senior employment rises by 9.8 percent, both significant, while junior employment falls by an insignificant 2.2 percent, close to our main estimate (Figure A29). The reallocation of seniors toward highly exposed occupations nearly doubles to 1.7 percentage points, and the computer and mathematical employment share goes up by 1.2 points (Figure A45). The estimates are qualitatively similar using this alternative source of variation for the instrument, which offers evidence that local conditions around the headquarters may not be sufficient to explain the results.

Another possibility is that a parent that adopts generative AI may centralize exposed work at headquarters and reduce its affiliate's junior ranks whether or not the affiliate's own company adopts. To assess this channel, we re-estimate the model with adoption measured at the parent company rather than at the affiliate's company, so an affiliate counts as treated only when its parent's own company posts an AI job. The control pool stays as it is, namely corporate groups that never posted an AI job anywhere. Affiliates whose own company adopted under a non-adopting headquarters are neither treated under the new definition nor clean controls, so we exclude them from the sample. If centralization at the parent drives our results, parent adoption should reproduce our results in full, whereas if adoption by the affiliate's own company matters, the headquarters-based estimates should be smaller.[25] Figure A30 shows that the estimates are indeed smaller. The decline in the junior share shrinks from 1.9 to 1.1 percentage points and now comes from junior employment, which falls by an imprecisely estimated 3.6 percent, as the rise in senior employment goes toward zero. The computer and mathematical share increases by a reduced 0.6 points. Parent adoption thus seems to have more muted effects on its own, with larger impacts on workforce composition from company adoption.

[25]This robustness check does not isolate the direct effect of headquarters adoption, because the companies of many affiliates of adopting parents adopt themselves, so the estimates combine the direct channel with company adoption that follows from headquarters adoption.

The next concern is imitation. A parent surrounded by adopters may copy their organizational hierarchies and pass those practices on to its affiliates, irrespective of AI. Because imitation of this kind is most likely to occur within industries, we rebuild the instrument from metro peer parents outside the focal parent's own sector, so as to exclude same-industry practices transmission from the instrument. The instrument remains strong, with a first-stage F of 100.9. The junior share falls by 1.5 percentage points and senior employment rises by 4.6 percent, both smaller than the main estimates, while junior employment falls by 2.7 percent, so the pattern is qualitatively unchanged (Figure A28 and Figure A46).

Our interpretation of the estimates as a local average treatment effect further requires monotonicity, which holds that peer adoption around the parent never makes an affiliate less likely to be treated. To test this, we estimate the first stage separately within parent sectors and affiliate size terciles (Figure A31). All eleven sector estimates are positive, ranging from 0.07 in other services to 0.32 in government, and ten are significant, as are all three size terciles. These results support the monotonicity assumption.

Another concern is that a handful of large technology hubs are driving most of the identifying variation in the instrument. We re-estimate the pooled endpoint after dropping, one at a time, each of the ten metropolitan areas that contribute the most treated affiliates (Figure A32; Figure A47 repeats the exercise for the occupation shares). The effect on the junior share ranges from -1.5 percentage points, when we drop San Francisco, to -2.0 points, when we drop Paris, and the effect on senior employment ranges from 5.7 to 7.5 percent. The first-stage F never falls below 90. Removing New York, London, San Jose, or Tokyo changes the junior-share estimate by at most 0.1 percentage points. As such, the results are robust to excluding any particular metropolitan area.

5.3.2 Additional Robustness Checks

Next, we consider a set of additional robustness checks that vary our identification strategy and the sample construction.

We start by dropping the instrument entirely and re-estimating the treatment effect in two ways: a matched difference-in-differences event study that compares treated affiliates with their nearest-neighbor controls directly, and a synthetic difference-in-differences (SDID) that replaces the matches with weights on control affiliates and pre-treatment months, chosen so that the synthetic comparison group tracks the treated affiliates before November 2022 (Arkhangelsky et al., 2021). Appendix A14 gives additional details on the methodologies. Because neither design relies on the HQ peer instrument, these analyses need not be restricted to affiliates of multinational companies. To nonetheless keep our analyses comparable, we estimate the analyses both for the affiliate sample used in the main analyses, as well as for the full sample of affiliates. This allows us to separate differences across estimators from differences in samples. Throughout, we report both designs at the March 2026 endpoint.

For the foreign-affiliate sample, both designs reproduce the main compositional shifts (Figures A35 and A37). The junior share falls by 1.5 percentage points in the matched design and by 1.4 points in the synthetic design, in line with the 1.9-point 2SLS decline. Similarly, the share of computer and mathematical employment increases by 0.4 percentage points in the matched design and 0.3 points in the synthetic design, against 0.8 points in the 2SLS estimate (Figures A49 and A51). The synthetic estimates also mirror the main analysis on the remaining outcomes: by March 2026, total employment is 5.4 percent higher and senior employment 7.2 percent higher at treated affiliates than at their synthetic control, junior employment is 1.5 percent higher, and the top-minus-bottom exposure-quartile share gap rises among seniors (0.5 percentage points) while falling among juniors (0.2 points). The synthetic design fits the pre-treatment months within each country by construction, and the pooled pre-treatment gaps never exceed 0.3 percent of employment or 0.1 points of the junior share. The matched design, by contrast, shows clear pre-treatment divergences between treated and control affiliates for every employment level. These divergences corroborate the need for our IV analysis and force us to disregard the matched design's estimates for those outcomes. In particular, the matched design finds junior employment 7.0 percent and total employment 10.8 percent higher at treated affiliates by March 2026, the opposite sign to the instrumented estimate for juniors, which is what one expects if adopters were already on steeper growth paths. The synthetic design gives results very similar to the 2SLS estimates under different identifying assumptions, which offers further support for the main results.

When we extend both designs to the full sample of affiliates, the results change little. The junior share falls by 1.3 percentage points in both the matched and the synthetic design, while the share of computer and mathematical employment increases (Table A11 and figs. A36, A48 and A50). The synthetic employment effects are slightly smaller in magnitude. By March 2026, senior employment is 4.3 percent higher and junior employment 1.1 percent lower at treated affiliates than at their synthetic control, and total employment is 2.7 percent higher. The matched design again shows the same pre-treatment divergences as on the foreign-affiliate sample. Taken together, the compositional shifts are consistent across estimators and samples.

We next replace the binary adoption indicator with a continuous measure of how heavily the company hires for AI work. We measure treatment intensity by the ratio of its worldwide AI openings from November 2022 onward to its job postings in the 22 months before the launch.[26] The results in Figure A38 are qualitatively identical

[26]We use logs to address outliers, and add 1 to each ratio to deal with the fact that all control affiliates have a ratio of zero.

to our main results although the first stage weakens from 104.2 to 40.1. By March 2026, a 0.1 log-point higher adoption intensity (roughly the adoption intensity of the ninetieth-percentile adopter) lowers the junior share by 4.1 percentage points, raises senior employment by 16 percent, and reallocates seniors toward exposed occupations, while junior and total employment do not move significantly. Similarly, the share of employees working in Computer and Mathematical occupations increases with adoption. In other words, the workforce changes we observe hold whether we measure adoption at the extensive margin (whether companies adopt) or incorporate the intensive margin (how intensely they adopt).

Another concern is that identifying AI adoption through job advertisements may select companies recruiting in particular occupations, especially computer and mathematical occupations. Employment growth in those occupations among treated affiliates could therefore partly reflect occupation-specific recruitment associated with our adoption measure. To address this concern, we redo the matching with an additional restriction on the controls. We identify the two-digit occupation category of each treated company's first classifiable AI posting anywhere in our data, and impose that each control affiliate must have at least one job advertised by its company in that category in the affiliate's country in the post-treatment period. The resulting sample contains 43,781 unique affiliates (22,909 treated and 20,872 controls), forming 22,909 matched sets, and the first-stage F increases from 104.2 to 112.0. The junior share still falls by 1.9 percentage points, although the decline is now driven more by junior employment, which falls by approximately 6.0 percent, than by senior employment, which increases by an insignificant 3.5 percent (Figure A39). The computer and mathematical employment share still gains 0.8 percentage points and the junior share within that occupation still falls by 3.6 points (Figure A52). Thus, the compositional shifts we observe do not appear to be a mechanical feature of our adoption definition.

Last, the main specification requires each control affiliate's company to have posted at least one job anywhere in our data from November 2022 onward. Because this conditions the control pool on post-period hiring activity, which is itself a function of treatment status, we re-match after dropping the requirement, which raises the number of distinct control affiliates from 27,697 to 27,923. The results in Figures A43 and A54 are essentially unchanged.

Overall, across a range of robustness exercises there appears to be a decline in the junior share of employment in treated affiliates. This tends to coincide with a modest increase in overall employment driven by seniors. Senior employment tends to become more AI-exposed, and junior employment less so. The share of workers in computer and mathematical occupations increases. Diagnostics for assessing the robustness of the results include pre-trends, which are largely insignificant in the main results but show some prior trends in certain heterogeneity analyses; alternative sources of identifying variation, such as the country-by-sector instrument and synthetic difference-in-differences; and the other exercises considered above.

Timing A remaining concern involves timing. Some of our results emerge within a year of the launch of ChatGPT, yet the companies of many treated affiliates post their first AI job only in 2024 or 2025. Although companies plausibly start using AI in their workflows before it appears in their advertisements, these effects nonetheless seem early. We consider additional analyses to understand the timing of the effects.

Figure A40 estimates the effects separately for affiliates whose company's first AI advertisement appeared in 2023, in 2024, and in 2025 or later, each cohort against its own matched controls. The first stage is strong for the 2023 and 2024 cohorts ($F = 227$ and 75) and weak for the 2025 cohort ($F = 24$).

For the 2023 and 2024 cohorts, no pre-launch coefficient on the junior share, senior employment, or total employment is significant, and the 2024 cohort tends to respond later than the 2023 cohort. By December 2023, the junior share has fallen by 0.8 percentage points in the 2023 cohort and 0.6 points in the 2024 cohort; senior employment has risen by 2.4 percent in the 2023 cohort and is unchanged in the 2024 cohort, and total employment is 0.8 percent higher against 1.8 percent lower. Once the 2024 cohort starts advertising, its estimates catch up: by March 2026 the junior share is down 1.6 points in the 2023 cohort and 2.0 points in the 2024 cohort, and senior employment is up 8.5 and 6.1 percent. The 2024 cohort's junior share does begin to fall during 2023, before its first advertisement, at roughly two thirds of the 2023 cohort's pace.

The 2025 cohort behaves differently. Its estimates are noisier, and its employment levels drift upward over the pre-period, from about 10 log points below the control path in January 2021 to zero at the launch. Its junior share falls by 1.5 points by December 2023 and 2.7 points by December 2024, before any company in the cohort has advertised an AI job, and by March 2026 its junior employment is 11 percent lower with total employment 3 percent lower, both imprecise. Companies in this cohort also differ from earlier adopters. They are less often headquartered in the United States (37 against 48 percent), their affiliates are smaller (median baseline employment of 112 against 127 positions) and start with a higher junior share (57 against 53 percent), and they post a median of three AI advertisements in total, against 25 for earlier adopters. The notable pre-trends and small first-stage F statistic suggest that the matched controls provide a poor counterfactual for later adopters and that the instrument is fairly weak.[27]

Where the design is cleanest, with strong first stages and matched controls that track the treated affiliates before the launch, the response therefore follows the timing of adoption: the 2024 cohort moves later than the 2023 cohort. Where the first stage is weak and the pre-launch fit is poorer, the estimated response precedes the company's first advertisement. We read the 2025 cohort's estimates as less

[27]In the 2025 cohort the instrument must predict adoption among companies that did not adopt until 2025 despite heavy adoption by their parent's peers. This selection in timing could work against the peer effect and result in a weak first stage.

reliable rather than as strong evidence that the response precedes adoption, but we cannot rule out that part of the response among the latest adopters reflects factors other than their own adoption.

As an additional test, we compare early adopters to similar late adopters. Rather than matching to never-treated affiliates, we match affiliates whose company first advertised in 2024 or earlier to affiliates whose company first advertised in 2025 or later (Figure A41). The junior share of the earlier group is already 1.6 points lower than that of the later group by December 2024, before any late adopter has advertised an AI job, and 1.8 points lower by March 2026 (first-stage $F = 81$), with senior and total employment changes of similar size to the main estimates but imprecise. Because late adopters begin adopting in January 2025, the coefficients from that month onward compare earlier with later adoption rather than adoption with none. This suggests that among comparable firms, early adopters see earlier impacts than late adopters.

Finally, we re-run the design with a placebo launch date of January 2022, keeping the original sample, matching, and instrument, and examine the coefficients over the ten months between the placebo treatment date and the launch of ChatGPT, which is the end of the placebo sample (Figure A42 and Figure A53). Over that window the junior share moves by at most 0.2 percentage points and ends it exactly at the baseline, while total, junior, and senior employment, along with the reallocation of juniors and seniors across exposure levels, do not differ significantly from zero in the post period. No occupation share moves by more than 0.3 percentage points over the window, and only two of the 44 occupation-level confidence intervals exclude zero.

5.4 Scale and Bias

Finally, we use these empirical results to estimate the scale effect and labor demand shifters in the model of Section 2. Remark 1 shows that the scale term and the labor-demand shifters can be recovered from log employment in the least-exposed cells and from differences across cells.

We do not estimate those cell-level logs directly because employment in seniority-by-AI-exposure cells is often zero (Chen and Roth, 2024). Instead, we use a first-order approximation that can be computed directly from the log-employment and share coefficients reported in our results. Appendix A5 gives the exact computation. We approximate the scale term for a seniority level as the log employment change for the least AI-exposed quartile of jobs. We approximate the bias of AI for a seniority as the difference between the log employment change for the most-exposed quartile and that for the least-exposed quartile. We allow the scale term to vary by seniority, consistent with the nested production function of Appendix A4. The scale terms may also differ by seniority if the least-exposed quartile is itself affected by AI ($\alpha_{Q1,s} \neq 0$). The estimated bias and scale terms should accordingly be interpreted

Figure 7: Event study plots for the bias and scale terms of AI adoption

Notes: The figure plots the bias and scale terms of Proposition 1, in log units, month by month. Panels A and B plot the bias of AI toward the most-exposed occupations within juniors and within seniors, $\alpha_{Q4,s} - \alpha_{Q1,s}$, the implied effect of adoption on log employment in the top quartile of AI exposure minus that in the bottom quartile. Panels C and D plot the scale term for each group, the implied effect on log employment in the bottom exposure quartile. Both are computed from the two-stage least squares long-difference coefficients of Equation (15) on log group employment and on the top- and bottom-quartile exposure shares within the group, as described in Appendix A5. Sample, instrument, matching, and weights as in Figure 3. Junior covers Revelio seniority levels 1–2 and senior levels 3–7. Shaded bands are 95% confidence intervals from delta-method standard errors clustered by the metropolitan area of the parent's headquarters. The dashed vertical line marks ChatGPT's launch in November 2022.

as an approximation.

Figure 7 traces the bias and scale terms month by month, and Table 4 reports them at the March 2026 endpoint. AI adoption is biased toward the most-exposed seniors: their employment rises by 5.0 percent relative to the least-exposed seniors, an estimate significant at the 5 percent level. Among juniors the bias runs the other

Table 4: Bias and scale terms of AI adoption in log units

	Bias: $\alpha_{Q4} - \alpha_{Q1}$		Scale term			Senior − junior
	Juniors	Seniors	Juniors	Seniors	All	scale
2SLS	−0.020	0.050**	−0.018	0.043	0.025	0.061**
	(0.014)	(0.023)	(0.021)	(0.028)	(0.019)	(0.025)
OLS	−0.001	0.042***	0.069***	0.117***	0.097***	0.048***
	(0.004)	(0.006)	(0.007)	(0.010)	(0.008)	(0.007)
Match fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Effective first-stage F	104.2	104.2	104.2	104.2	104.2	104.2

Notes: The table maps the estimates of Table 3 into the objects of Proposition 1 at the March 2026 endpoint. For seniority group s and exposure cell c, cell employment is group employment times the cell's share of the group, so to first order the effect of adoption on log cell employment equals the effect on log group employment plus the effect on the cell share divided by its baseline share. The bias is the implied effect on log employment in the top exposure quartile minus that in the bottom quartile within the group; the scale term is the implied effect on log employment in the bottom quartile, which equals the firm-level scale term when AI leaves the least-exposed occupations' task content unchanged; the last column is the difference between the senior and junior scale terms, zero under a single firm-level scale term. Baseline shares are October 2022 employment-weighted shares among matched adopting foreign affiliates (bottom quartile: 21.3 percent of juniors, 6.8 percent of seniors; top quartile: 35.2 and 31.6 percent). Standard errors, in parentheses, are delta-method standard errors from a regression that stacks the underlying share and log-employment long differences, with match fixed effects and clustering by the metropolitan area of the parent's headquarters. Sample, instrument, and weights as in Table 3. */**/***: $p < 0.10/0.05/0.01$.

way, at −2.0 percent relative to the least-exposed juniors, though the estimate is not statistically distinguishable from zero at the endpoint. The implied scale terms are +4.3 percent for seniors, −1.8 percent for juniors, and +2.5 percent for all workers, none individually significant, and the last consistent with the modest total-employment effect in Figure 3. The difference between the senior and junior scale terms, 6.1 percentage points, is significant.

Overall, these results give suggestive evidence that AI adoption leads to a modest scale effect. AI is biased against exposed juniors and biased towards exposed seniors.

6 Conclusion

It is notoriously difficult to know which workers a new technology will replace, and generative AI is unlikely to prove an exception. Exposure scores cannot settle the question, because they measure the potential for labor-saving and labor-expanding uses alike. We pair a task-based model of firm-level adoption with large-scale employment data to measure which jobs gain and lose from AI adoption. Our model shows how endogenous AI adoption changes a firm's employment through labor-saving and labor-expanding shifts in labor demand and through a scale effect. It shows that comparisons between adopters and otherwise similar non-adopters reveal, for each group of workers, whether AI is on net labor saving or labor expanding.

We test the framework on 1.25 billion job postings and 154 million positions in 41 countries, coding companies as AI adopters once they post a job advertisement that involves generative AI. We instrument company adoption in the foreign-affiliate sample with the adoption rate among the other multinational parents headquartered in the same metropolitan area as the affiliate's parent. By March 2026, adoption has lowered the junior share of employment by 1.9 percentage points relative to matched controls. Junior employment falls by 2.5 percent, although this estimate is imprecise, and senior employment rises by 6.7 percent. The share of employees in computer and mathematical occupations increases by 0.8 percentage points, and within the occupation employment shifts from juniors to seniors. The decline in the junior share is consistent across the world, appearing in 23 of the 31 countries we can estimate separately. Taken together, our evidence is consistent with the conclusion that generative AI has so far been labor saving for junior workers and labor expanding for senior workers and for workers in computer and mathematical occupations. Through the lens of the model, adoption shifts labor demand toward the most AI-exposed seniors by 5 percent relative to less-exposed seniors, and away from the most AI-exposed juniors by 2 percent relative to less-exposed juniors. Scale effects are modest, with a point estimate of 2.5 percent across the affiliate's workforce.

We acknowledge some limitations. First, LinkedIn profiles over-represent managers and professionals and cover some countries better than others, so our estimates describe parts of the economy visible on the platform. Second, our instrument is imperfect, and although our robustness checks help address the exclusion-restriction channels we can test, no set of checks rules out every channel. Relatedly, our instrumented estimates come from affiliates of foreign multinationals, which are larger and more internationally integrated than the average affiliate. The matched event study and the synthetic difference-in-differences, however, which require no instrument and extend to domestic affiliates, show similar recomposition. Third, we do not observe the exact moment when a company starts using generative AI. Job advertisements likely lag AI adoption, which is why we date treatment at the launch of ChatGPT rather than at the company's first AI posting. Moreover, companies that use the technology without ever advertising AI jobs count as non-adopters in our data. This likely biases our estimates toward zero, however, suggesting that the true effects are likely to be, if anything, larger than the ones we report. Fourth, our panel covers just over three years after the launch of ChatGPT. The estimates therefore capture the initial reorganization of treated affiliates, which need not be the same as the long-run equilibrium in which wages, training, and career choices have had time to adjust.

AI adoption has not (yet) led to large workforce reductions at treated affiliates, but our results indicate that it has led to a shift in composition away from junior workers. These shifts may have implications for the pipeline of experienced workers that firms will draw on in the future (Ide, 2025). Our model highlights that more work is needed to assess if these findings apply in equilibrium, and our empirical results do not make predictions about how they will persist. Policy reports such as World Economic Forum (2026) urge employers to keep hiring at the entry level and governments to build other routes into skilled work. Whether these interventions are needed depends on these questions about generalization and persistence.

Our results suggest several directions for future research. First, treated affiliates and their matched controls hire in the same labor market and face the same wages, leaving open the question of whether senior wages rise and junior pay declines with the observed employment changes. Second, we measure recomposition across existing occupations. Like earlier technologies, AI will also create new kinds of work (Autor et al., 2024), which our classification cannot yet record. Finally, as the technology continues to improve, AI may be labor saving and labor expanding for a different set of workers than it is today, and future research can use our framework to find out.

References

Acemoglu, D., Anderson, G. W., Beede, D. N., Buffington, C., Childress, E. E., Dinlersoz, E., Foster, L. S., Goldschlag, N., Haltiwanger, J., Kroff, Z., Restrepo, P. and Zolas, N. (2025), Automation and the workforce: A firm-level view from the 2019 annual business survey, *in* S. Basu, L. Eldridge, J. Haltiwanger and E. Strassner, eds, 'Technology, Productivity, and Economic Growth', Vol. 83 of *Studies in Income and Wealth*, University of Chicago Press, chapter 1, pp. 13–55. **URL:** `https://www.nber.org/books-and-chapters/technology-productivity-and-economic-growth/automation-and-workforce-firm-level-view-2019-annual-business-survey`

Acemoglu, D. and Autor, D. (2011), Skills, Tasks and Technologies: Implications for Employment and Earnings, *in* O. Ashenfelter and D. Card, eds, 'Handbook of Labor Economics', Vol. 4B, Elsevier, pp. 1043–1171. **URL:** `https://www.sciencedirect.com/science/article/pii/S0169721811024105`

Acemoglu, D., Autor, D., Hazell, J. and Restrepo, P. (2022), 'Artificial Intelligence and Jobs: Evidence from Online Vacancies', *Journal of Labor Economics* **40**(S1), S293–S340. **URL:** `https://www.journals.uchicago.edu/doi/abs/10.1086/718327`

Acemoglu, D., Kong, F. and Restrepo, P. (2025), Tasks at work: Comparative advantage, technology and labor demand, *in* C. Dustmann and T. Lemieux, eds, 'Handbook of Labor Economics', Vol. 6, Elsevier, pp. 1–114.

Acemoglu, D., Lelarge, C. and Restrepo, P. (2020), 'Competing with Robots: Firm-Level Evidence from France', *AEA Papers and Proceedings* **110**, 383–388. **URL:** `https://www.aeaweb.org/articles?id=10.1257/pandp.20201003`

Acemoglu, D. and Restrepo, P. (2018), 'The race between man and machine: Implications of technology for growth, factor shares, and employment', *American Economic Review* **108**(6), 1488–1542.

Acemoglu, D. and Restrepo, P. (2019), 'Automation and New Tasks: How Technology Displaces and Reinstates Labor', *Journal of Economic Perspectives* **33**(2), 3–30.

Acemoglu, D. and Restrepo, P. (2022), 'Tasks, Automation, and the Rise in U.S. Wage Inequality', *Econometrica* **90**(5), 1973–2016.

Adams, Z., Billa, B., Bishop, L., Deakin, S. and Shroff, T. (2023), 'CBR labour regulation index (dataset of 117 countries, 1970–2022)'. 2023 update, published

December 2023; data year 2022. Accessed: 2026-08-03.
URL: https://www.repository.cam.ac.uk/items/938d5a0d-3799-4c5a-8103-8a7355628ef3

Aghion, P., Antonin, C., Bunel, S. and Jaravel, X. (2026), 'Modern manufacturing capital, labor demand, and product market dynamics: Evidence from France', Working paper, conditionally accepted at *American Economic Journal: Macroeconomics*.
URL: https://sites.google.com/view/simon-bunel/research

Aguirre, T. and Meloni, L. (2026), Early Labor-Market Responses to Generative AI in a Developing Country: Evidence from Brazil. Working paper, GovAI and University of São Paulo.
URL: https://luismeloni.com/labour_ai_brazil_2026.pdf

Aldasoro, I., Gambacorta, L., Pal, R., Revoltella, D., Weiss, C. and Wolski, M. (2026), AI adoption, productivity and employment: Evidence from European firms, EIB Working Paper 2026/02, European Investment Bank (EIB).
URL: https://www.econstor.eu/handle/10419/335876

Althoff, L. and Reichardt, H. (2026), Task-specific technical change and comparative advantage, Working Paper 35353, National Bureau of Economic Research.
URL: https://www.nber.org/papers/w35353

Appel, R., McCrory, P., Tamkin, A., Stern, M., McCain, M. and Neylon, T. (2025), Anthropic Economic Index report: Uneven geographic and enterprise AI adoption, Report, Anthropic.
URL: https://www.anthropic.com/research/anthropic-economic-index-september-2025-report

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W. and Wager, S. (2021), 'Synthetic difference-in-differences', *American Economic Review* **111**(12), 4088–4118.

Ashenfelter, O. (1978), 'Estimating the effect of training programs on earnings', *The Review of Economics and Statistics* **60**(1), 47–57.

Autor, D., Chin, C., Salomons, A. and Seegmiller, B. (2024), 'New Frontiers: The Origins and Content of New Work, 1940–2018', *The Quarterly Journal of Economics* **139**(3), 1399–1465.
URL: https://doi.org/10.1093/qje/qjae008

Autor, D. H., Levy, F. and Murnane, R. J. (2003), 'The Skill Content of Recent Technological Change: An Empirical Exploration', *The Quarterly Journal of Economics* **118**(4), 1279–1333.
URL: https://doi.org/10.1162/003355303322552801

Autor, D. and Thompson, N. (2025), Expertise, Working Paper 33941, National Bureau of Economic Research.
URL: `https://www.nber.org/papers/w33941`

Azar, J., Gine, M. and Sanz-Espín, J. (2025), 'AI Is Already Eroding Wages: Quasi-Experimental Evidence From Occupational Exposure'.
URL: `https://papers.ssrn.com/abstract=5842084`

Babina, T. (2026), Understanding firms' AI efforts and their economic impact, Working Paper 35123, National Bureau of Economic Research.
URL: `https://www.nber.org/papers/w35123`

Babina, T., Fedyk, A., He, A. X. and Hodson, J. (2024), 'Artificial intelligence, firm growth, and product innovation', *Journal of Financial Economics* **151**, 103745.
URL: `https://www.sciencedirect.com/science/article/pii/S0304405X2300185X`

Bessen, J., Goos, M., Salomons, A. and van den Berge, W. (2025), 'What happens to workers at firms that automate?', *The Review of Economics and Statistics* **107**(1), 125–141.

Bick, A., Blandin, A. and Deming, D. J. (2024), The Rapid Adoption of Generative AI, Working Paper 32966, National Bureau of Economic Research.
URL: `https://www.nber.org/papers/w32966`

Bilir, L. K. and Morales, E. (2020), 'Innovation in the global firm', *Journal of Political Economy* **128**(4), 1566–1625.

Blandhol, C., Bonney, J., Mogstad, M. and Torgovitsky, A. (2026), 'When is TSLS actually LATE?', *Review of Economic Studies*, advance article.
URL: `https://doi.org/10.1093/restud/rdag029`

Bloom, N., Moore, G., Simon, L. K. and Wilkie-Rogers, C. (2026), Time Travel on Professional Profiles, Working Paper 35546, National Bureau of Economic Research.
URL: `https://www.nber.org/papers/w35546`

Bloom, N., Sadun, R. and Van Reenen, J. (2012), 'Americans do IT better: US multinationals and the productivity miracle', *American Economic Review* **102**(1), 167–201.

Bonfiglioli, A., Crinò, R., Fadinger, H. and Gancia, G. (2024), 'Robot imports and firm-level outcomes', *The Economic Journal* **134**(664), 3428–3444.

Bonin, S. M., Darougheh, S. and Kuchler, A. (2026), AI Adoption and Firm Size: Employment Effects in Monthly Administrative Data, Working Paper 223, Danmarks Nationalbank. Supersedes the slides "Prompting Change: Firm-Level AI Adoption and Worker-Level Outcomes in Linked Administrative Data" presented at the CCBS Monetary Policy Conference, Bank of England, 24 June 2026.
URL: https://www.nationalbanken.dk/media/m2rn4owp/ai-adoption-and-firm-size-employment-effects-in-monthly-administrative-data.pdf

Bonney, K., Breaux, C. L., Dinlersoz, E., Foster, L. S., Haltiwanger, J. C. and Pande, A. A. (2026), The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks, Working Paper 35141, National Bureau of Economic Research.
URL: https://www.nber.org/papers/w35141

Brynjolfsson, E., Chandar, B. and Chen, R. (2026), Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence, Working paper, Stanford Digital Economy Lab. August 2026 version.
URL: https://digitaleconomy.stanford.edu/app/uploads/2026/08/Canaries_August2026.pdf

Cai, J. and Szeidl, A. (2018), 'Interfirm relationships and business performance', *The Quarterly Journal of Economics* **133**(3), 1229–1282.

Chandar, B. (2025), Tracking Employment changes in AI-Exposed Jobs, SSRN Working Paper 5384519, Social Science Research Network.
URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5384519

Chatterji, A., Cunningham, T., Deming, D. J., Hitzig, Z., Ong, C., Shan, C. Y. and Wadman, K. (2025), How People Use ChatGPT, Working Paper 34255, National Bureau of Economic Research.
URL: https://www.nber.org/papers/w34255

Chen, F. and Stratton, J. (2026), Artificial intelligence in the firm. Preliminary draft, January 7, 2026.
URL: https://fion.ac/jellyfish.pdf

Chen, J. and Roth, J. (2024), 'Logs with Zeros? Some Problems and Solutions', *The Quarterly Journal of Economics* **139**(2), 891–936.
URL: https://doi.org/10.1093/qje/qjad054

DerSimonian, R. and Laird, N. (1986), 'Meta-analysis in clinical trials', *Controlled Clinical Trials* **7**(3), 177–188.

Dingel, J. I. and Neiman, B. (2020), 'How many jobs can be done at home?', *Journal of Public Economics* **189**, 104235.

URL: https://www.sciencedirect.com/science/article/pii/S0047272720300992

Dixon, J., Hong, B. and Wu, L. (2021), 'The robot revolution: Managerial and employment consequences for firms', *Management Science* **67**(9), 5586–5605.

Eckhardt, S. and Goldschlag, N. (2025), AI and Jobs: The Final Word (Until the Next One), Report, Economic Innovation Group.
URL: https://eig.org/ai-and-jobs-the-final-word/

Eloundou, T., Manning, S., Mishkin, P. and Rock, D. (2024), 'GPTs are GPTs: Labor market impact potential of LLMs', *Science* **384**(6702), 1306–1308.
URL: https://www.science.org/doi/10.1126/science.adj0998

Engberg, E., Hellsten, M., Javed, F., Lodefalk, M., Sabolová, R., Schroeder, S. and Tang, A. (2025), 'Artificial intelligence, hiring and employment: Job postings evidence from Sweden', *Applied Economics Letters*, advance online publication.
URL: https://doi.org/10.1080/13504851.2025.2497431

Eurostat (2025), 'Use of artificial intelligence in enterprises', Statistics Explained.
URL: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Use_of_artificial_intelligence_in_enterprises

Feigenbaum, J. and Gross, D. P. (2024), 'Organizational and Economic Obstacles to Automation: A Cautionary Tale from AT&T in the Twentieth Century', *Management Science* **70**(12), 8520–8540.
URL: https://doi.org/10.1287/mnsc.2022.01760

Felten, E., Raj, M. and Seamans, R. (2023), How will Language Modelers like ChatGPT Affect Occupations and Industries?, Technical report, arXiv. arXiv:2303.01157 [econ].
URL: http://arxiv.org/abs/2303.01157

Frank, M. R., Sabet, A. J., Simon, L., Bana, S. H. and Yu, R. (2026), 'AI-exposed jobs deteriorated before ChatGPT'. arXiv:2601.02554 [econ].
URL: http://arxiv.org/abs/2601.02554

Freund, L. B. and Mann, L. F. (2025), Job transformation, specialization, and the labor market effects of AI, RFBerlin Discussion Paper 117/25, ROCKWOOL Foundation Berlin.
URL: https://www.rfberlin.com/wp-content/uploads/2025/11/25117.pdf

Gimbel, M., Kinder, M., Kendall, J. and Lee, M. (2025), 'Evaluating the Impact of AI on the Labor Market: Current State of Affairs', Report, The Budget Lab at Yale.

URL: https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs

Gmyrek, P., Berg, J. and Bescond, D. (2023), Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality, SSRN Scholarly Paper, Social Science Research Network, Rochester, NY.
URL: https://papers.ssrn.com/abstract=4584219

Gmyrek, P., Berg, J., Kamiński, K., Konopczyński, F., Ładna, A., Nafradi, B., Rosłaniec, K. and Troszyński, M. (2025), 'Generative AI and jobs: A 2025 update', ILO Research Brief, International Labour Organization, Geneva.
URL: https://researchrepository.ilo.org/esploro/outputs/encyclopediaEntry/Generative-AI-and-jobs-a-2025/995653517302676

Grattafiori, A., Dubey, A., Jauhri, A., Pandey, A., Kadian, A. et al. (2024), 'The Llama 3 Herd of Models'. arXiv:2407.21783.
URL: https://arxiv.org/abs/2407.21783

Guadalupe, M., Kuzmina, O. and Thomas, C. (2012), 'Innovation and foreign ownership', *American Economic Review* **102**(7), 3594–3627.

Hampole, M., Papanikolaou, D., Schmidt, L. D. W. and Seegmiller, B. (2025), Artificial Intelligence and the Labor Market, Working Paper 33509, National Bureau of Economic Research.
URL: https://www.nber.org/papers/w33509

Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., Troy, K. K., Amodei, D., Kaplan, J., Clark, J. and Ganguli, D. (2025), Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations, Technical report, arXiv. arXiv:2503.04761 [cs].
URL: http://arxiv.org/abs/2503.04761

Hanney, O. and Kim, O. (2026), 'Tomorrow's AI Won't Meet Today's World', *In Development*.
URL: https://indevelopmentmag.com/tomorrows-ai-wont-meet-todays-world/

Hernæs, Ø. and Kostøl, A. R. (2026), Has AI Widened Employment Gaps? Tracking Early-Career Employment by Occupational Exposure in Norway, IZA Discussion Paper 18767, IZA@LISER Network.
URL: https://docs.iza.org/dp18767.pdf

Hirvonen, J., Stenhammar, A. and Tuhkuri, J. (2022), New evidence on the effect of technology on employment and skill demand, Working Paper 93, ETLA.

URL: https://www.etla.fi/wp-content/uploads/ETLA-Working-Papers-93.pdf

Hosseini Maasoum, S. M. and Lichtinger, G. (2025), 'Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data'.
URL: https://papers.ssrn.com/abstract=5425555

Hsieh, C.-T., Hurst, E., Jones, C. I. and Klenow, P. J. (2019), 'The Allocation of Talent and U.S. Economic Growth', *Econometrica* **87**(5), 1439–1474.

Huang, J., Liu, Y., Wang, H. and Yu, S. (2026), Who Takes the Hit? The Uneven Impacts of Generative AI on Labor Demand Across Countries, Working paper, World Bank. Background paper for the World Development Report 2026: Artificial Intelligence for Development. July 27, 2026 version.

Hubmer, J. and Restrepo, P. (2026), 'Not a typical firm: Capital–labor substitution and firms' labor shares', *American Economic Journal: Macroeconomics* **18**(2), 34–71.

Humlum, A. (2019), Robot adoption and labor market dynamics, Working paper, Princeton University.
URL: https://economics.yale.edu/sites/default/files/humlumjmp_111419.pdf

Humlum, A. and Vestergaard, E. (2025), Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI, Working Paper 33777, National Bureau of Economic Research. Revised March 2026; previously circulated as "Large Language Models, Small Labor Market Effects".
URL: https://www.nber.org/papers/w33777

Ide, E. (2025), 'Automation, AI, and the intergenerational transmission of knowledge'. arXiv:2507.16078.
URL: https://arxiv.org/abs/2507.16078

Imbens, G. W. and Angrist, J. D. (1994), 'Identification and estimation of local average treatment effects', *Econometrica* **62**(2), 467–475.

International Institute for Management Development (2024), 'IMD world digital competitiveness ranking 2024'. Overall score, retrieved from the IMD World Competitiveness Statbase. Accessed: 2026-01-23.
URL: https://www.imd.org/centers/wcc/world-competitiveness-center/rankings/world-digital-competitiveness-ranking/

International Labour Organization (2026), 'Employment by sex and economic activity, ILOSTAT database'. Employment in ISIC Rev. 4 section J and total employment, 2019–2022, retrieved from the ILO SDMX gateway. Accessed: 2026-07-07.
URL: `https://ilostat.ilo.org/data/`

Kalyani, A., Bloom, N., Carvalho, M., Hassan, T. A., Lerner, J. and Tahoun, A. (2025), 'The diffusion of new technologies', *The Quarterly Journal of Economics* **140**(2), 1299–1365.

Kauhanen, A. and Rouvinen, P. (2025), 'Assessing early labour market effects of generative AI: Evidence from population data', *Applied Economics Letters*, advance online publication.
URL: `https://doi.org/10.1080/13504851.2025.2513973`

Kharazian, A., Simon, L. and Stevens, R. (2026), A New Look at AI's Impact on Jobs: Firm-Level AI Spending and Workforce Adjustment, Ramp Economics Lab Working Paper ramp-ai-jobs-2026, Ramp Economics Lab.
URL: `https://ramp.com/data/ai-jobs-impact`

Klein Teeselink, B. (2025), Generative AI and Labor Market Outcomes: Evidence from the United Kingdom, SSRN Scholarly Paper 5516798, Social Science Research Network, Rochester, NY.
URL: `https://papers.ssrn.com/abstract=5516798`

Klein Teeselink, B. and Carey, D. (2026), 'AI, automation, and expertise'.
URL: `https://papers.ssrn.com/abstract=6134506`

Koch, M., Manuylov, I. and Smolka, M. (2021), 'Robots and firms', *The Economic Journal* **131**(638), 2553–2584.

Kulpa, W. (1997), 'The Poincaré–Miranda theorem', *The American Mathematical Monthly* **104**(6), 545–550.

Lambert, P. J. and Schindler, Y. (2026), The Broken Ladder: AI, Remote Work, and Early-Career Hiring, SSRN Scholarly Paper 6787638, Social Science Research Network, Rochester, NY.
URL: `https://papers.ssrn.com/abstract=6787638`

Lindner, A., Muraközy, B., Reizer, B. and Schreiner, R. (2022), Firm-level technological change and skill demand, Discussion Paper 17421, CEPR.
URL: `https://cepr.org/publications/dp17421`

Massenkoff, M. and McCrory, P. (2026), Labor market impacts of AI: A new measure and early evidence, Report, Anthropic.
URL: `https://cdn.sanity.io/files/4zrzovbb/website/2b5bbaf2c1eb81dbf6e6fb813c1a24e35a64d376.pdf`

McFadden, D. (1963), 'Constant elasticity of substitution production functions', *The Review of Economic Studies* **30**(2), 73–83.

Microsoft AI Economy Institute (2026), Global AI Diffusion Q1 2026: Trends and Insights, AI Diffusion Report, Microsoft AI Economy Institute.
URL: https://www.microsoft.com/en-us/research/wp-content/uploads/2026/05/Microsoft-AI-Diffusion-Report-2026-Q1.pdf

Miranda, C. (1940), 'Un'osservazione su un teorema di Brouwer', *Bollettino dell'Unione Matematica Italiana* **3**, 5–7.

OpenAI (2026), 'From asking to doing: How the world is putting ChatGPT to work', OpenAI Signals; country ranking of consumer ChatGPT messages per working-age person, April–June 2026, data at openai.com/signals/data.
URL: https://openai.com/index/how-the-world-is-putting-chatgpt-to-work/

Orr, C., Tucker, L. C. and Warren, L. (2026), Graduating into Disruption: Labor Market Outcomes for AI-Exposed College Majors, Working Paper CES-26-56, U.S. Census Bureau, Center for Economic Studies.
URL: https://www.census.gov/library/working-papers/2026/adrm/CES-WP-26-56.html

Ravanilla, N. (2026), 'Generative AI and labor displacement in the Global South: Evidence from Philippine business process outsourcing', Working paper, University of California San Diego.
URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7313383

Tomlinson, K., Jaffe, S., Wang, W., Counts, S. and Suri, S. (2025), 'Working with AI: Measuring the Applicability of Generative AI to Occupations'. arXiv:2507.07935 [cs].
URL: http://arxiv.org/abs/2507.07935

Tucker, L. C. (2026), You're (not) Hired: Artificial Intelligence and Early Career Hiring in the Quarterly Workforce Indicators, Working Paper CES-26-27, U.S. Census Bureau, Center for Economic Studies.
URL: https://www.census.gov/library/working-papers/2026/adrm/CES-WP-26-27.html

U.S. Census Bureau (2026), 'Large firms with at least 20 employees biggest AI users', America Counts, Business Trends and Outlook Survey.
URL: https://www.census.gov/library/stories/2026/05/ai-use-businesses.html

Uzawa, H. (1962), 'Production functions with constant elasticities of substitution', *The Review of Economic Studies* **29**(4), 291–299.

World Bank (2026), 'World development indicators'. Accessed: 2026-03-04 and 2026-07-07.
URL: https://databank.worldbank.org/source/world-development-indicators

World Economic Forum (2026), 'Artificial intelligence and the future of entry-level work: A framework for safeguarding and reinventing early career pathways'. Insight Report, June 2026, in collaboration with PwC. Accessed: 2026-08-11.
URL: https://www.weforum.org/publications/artificial-intelligence-and-the-future-of-entry-level-work-a-framework-for-safeguarding-and-reinventing-early-career-pathways/

Yotzov, I., Barrero, J. M., Bloom, N., Bunn, P., Davis, S. J., Foster, K. M., Jalca, A., Meyer, B. H., Mizen, P., Navarrete, M. A., Smietanka, P., Thwaites, G. and Wang, B. Z. (2026), Firm Data on AI, Working Paper 34836, National Bureau of Economic Research.
URL: https://www.nber.org/papers/w34836

A1 Proofs for the Firm-Level Model

A1.1 The reduced-form production function (3)

Task aggregator. Output of firm f is a CES aggregator over a continuum of tasks $\mathcal{Z}_f \subset \mathbb{R}_+$ with elasticity of substitution η:

$$y_f = \left[\int_{\mathcal{Z}_f} q_f(z)^{(\eta-1)/\eta} dz\right]^{\eta/(\eta-1)}, \tag{17}$$

where $q_f(z)$ is the quantity of task-z output. The set of tasks $\mathcal{Z}_f$ may differ across firms.

Task assignment. Each task z is performed by a single factor g (a labor type or, at adopters, AI capital) with productivity $\psi_f^g(z)$, a measurable function of z bounded between positive constants. If factor g performs task z, then $q_f(z) = \psi_f^g(z) \cdot x_f^g(z)$, where $x_f^g(z)$ is the input quantity of factor g allocated to task z. Firms minimize cost by assigning each task to the factor with the lowest unit cost $w_g/\psi_f^g(z)$. Let $\mathcal{Z}_f^g = \{z \in \mathcal{Z}_f : g \text{ does task } z\}$ denote the set of tasks performed by factor g at firm f, a measurable set of finite positive measure, so that Ψ_f^g in (18) below is finite and positive for either range of η, with $\bigcup_g \mathcal{Z}_f^g = \mathcal{Z}_f$ and $\mathcal{Z}_f^g \cap \mathcal{Z}_f^{g'} = \emptyset$ for $g \neq g'$.

Effective factor productivity. Define

$$\Psi_f^g \equiv \left(\int_{\mathcal{Z}_f^g} \psi_f^g(z)^{\eta-1} dz\right)^{1/(\eta-1)}, \tag{18}$$

the firm- and factor-specific integrated effective productivity over the factor's task set, so that $(\Psi_f^g)^{\eta-1} = \int_{\mathcal{Z}_f^g} \psi_f^g(z)^{\eta-1} dz$ is the factor's effective task content of Section 2.1.

Derivation. This subsection derives the reduced-form CES production function in factor inputs (3) from the task aggregator (17). The key step is the optimization of $x_f^g(z)$ across tasks $z \in \mathcal{Z}_f^g$ for given total factor input $x_f^g \equiv \int_{\mathcal{Z}_f^g} x_f^g(z)\, dz$.

Setup. Substituting the assignment $q_f(z) = \psi_f^g(z)\, x_f^g(z)$ on $\mathcal{Z}_f^g$ into (17) and using the disjoint partition $\mathcal{Z}_f = \bigsqcup_g \mathcal{Z}_f^g$ (in the first integral, g is the factor performing task z):

$$y_f^{(\eta-1)/\eta} = \int_{\mathcal{Z}_f} \left(\psi_f^g(z)\, x_f^g(z)\right)^{(\eta-1)/\eta} dz = \sum_g \int_{\mathcal{Z}_f^g} \left(\psi_f^g(z)\, x_f^g(z)\right)^{(\eta-1)/\eta} dz. \tag{19}$$

This is not yet (3) because the right-hand side still depends on the full task-level profile $x_f^g(\cdot)$, not just on the totals $\{x_f^g\}$.

Task-level allocation. The firm chooses the task-level profiles $\{x_f^g(\cdot)\}$ to minimize cost $\sum_g w_g \, x_f^g$ subject to producing y_f through the task aggregator (17). We solve this in two stages. Since the cost objective depends on the profiles only through their totals $\{x_f^g\}$, the firm can first choose the profiles to make output as large as possible holding the totals fixed, and then minimize cost over the totals. The first stage separates across factors: each profile $x_f^g(\cdot)$ enters the split aggregator (19) only through its own term and contributes additively to $y_f^{(\eta-1)/\eta}$.

One direction-of-optimization subtlety deserves note. For $\eta > 1$ the exponent $(\eta-1)/\eta$ is positive, y is increasing in each term of (19), and the first stage *maximizes* $\int_{\mathcal{Z}_f^g}(\psi_f^g x_f^g)^{(\eta-1)/\eta} dz$ subject to $\int_{\mathcal{Z}_f^g} x_f^g(z)\, dz = x_f^g$. For $\eta < 1$ the exponent is negative: $y_f = \left[\sum_g \int (\psi x)^{(\eta-1)/\eta}\right]^{\eta/(\eta-1)}$ has a negative outer exponent, so y is *decreasing* in each inner term, and making output as large as possible means *minimizing* the same integral. In both cases the optimizer is interior and characterized by the same stationarity condition:

$$\frac{\eta-1}{\eta}\, \psi_f^g(z)^{(\eta-1)/\eta}\, x_f^g(z)^{-1/\eta} = \lambda_g \qquad \Longrightarrow \qquad x_f^g(z) = \left(\frac{\eta-1}{\eta\,\lambda_g}\right)^{\eta} \psi_f^g(z)^{\eta-1},$$

for a multiplier λ_g pinned down by the budget. (For $\eta > 1$ the inner problem is concave and the condition is sufficient for the max; for $\eta < 1$ the map $x \mapsto x^{(\eta-1)/\eta}$ is convex and the same condition is sufficient for the min.)

Write $x_f^g(z) = K_g\, \psi_f^g(z)^{\eta-1}$ where $K_g \equiv [(\eta-1)/(\eta\,\lambda_g)]^{\eta}$ is constant in z. Integrating over $\mathcal{Z}_f^g$ and using the budget and the definition (18):

$$x_f^g \;=\; K_g \int_{\mathcal{Z}_f^g} \psi_f^g(z)^{\eta-1}\, dz \;=\; K_g\, (\Psi_f^g)^{\eta-1} \quad \Longrightarrow \quad K_g \;=\; \frac{x_f^g}{(\Psi_f^g)^{\eta-1}},$$

so the optimal within-factor allocation is

$$x_f^g(z) \;=\; \frac{\psi_f^g(z)^{\eta-1}}{(\Psi_f^g)^{\eta-1}}\, x_f^g. \tag{20}$$

For $\eta > 1$ the allocation loads on higher-productivity tasks; for $\eta < 1$ (complements) it loads on lower-productivity tasks, as expected.

Substitute and aggregate. Plugging (20) into the factor-g term of (19), the integrand becomes

$$\left(\psi_f^g(z)\, x_f^g(z)\right)^{(\eta-1)/\eta} \;=\; (x_f^g)^{(\eta-1)/\eta}\, (\Psi_f^g)^{-(\eta-1)^2/\eta}\, \psi_f^g(z)^{\eta-1},$$

and integrating over $\mathcal{Z}_f^g$ using $\int_{\mathcal{Z}_f^g} \psi_f^g(z)^{\eta-1}\,dz = (\Psi_f^g)^{\eta-1}$:

$$\int_{\mathcal{Z}_f^g} \big(\psi_f^g(z)\, x_f^g(z)\big)^{(\eta-1)/\eta}\, dz \;=\; (x_f^g)^{(\eta-1)/\eta}\, (\Psi_f^g)^{(\eta-1)\,-\,(\eta-1)^2/\eta} \;=\; \big(\Psi_f^g\, x_f^g\big)^{(\eta-1)/\eta},$$

since the exponent on Ψ_f^g collapses to $(\eta-1)\big[1-(\eta-1)/\eta\big] = (\eta-1)/\eta$. Substituting into (19) and raising both sides to $\eta/(\eta-1)$ yields (3). □

A1.2 Additive decomposition of the labor-demand shifter (5)

Fix a type $t = (o,s)$ and firm f. Write $\mathcal{Z}_f^{t,\text{pre}}$ and $\mathcal{Z}_f^{t,\text{post}}$ for the type's task sets before and after adoption, ψ_f^t and $\psi_f^{t,\text{post}}$ for its per-task productivities before and after, $\Psi_f^{t,\text{pre}} \equiv \Psi_f^{t,0}$ and $\Psi_f^{t,\text{post}} \equiv \Psi_f^{t,A}$ for the corresponding effective productivities (18), and $\mathcal{Z}_f^{t,r} \equiv \mathcal{Z}_f^{t,\text{pre}} \cap \mathcal{Z}_f^{t,\text{post}}$ for the retained-task domain. The three components of (5) are log ratios of integrated effective task content $\int \psi^{\eta-1}\,dz$ on different domains:

- $\alpha_{\text{disp}}^{(o,s,f)} \;=\; \ln\left(\dfrac{\int_{\mathcal{Z}_f^{t,\text{pre}}} \psi_f^t(z)^{\eta-1}\,dz}{\int_{\mathcal{Z}_f^{t,r}} \psi_f^t(z)^{\eta-1}\,dz}\right) \;\geq\; 0$: log shrinkage of effective task content from labor's task set — tasks the type no longer performs, reassigned to AI capital under channel (i).
- $\alpha_{\text{reinst}}^{(o,s,f)} \;=\; \ln\left(\dfrac{\int_{\mathcal{Z}_f^{t,\text{post}}} \psi_f^{t,\text{post}}(z)^{\eta-1}\,dz}{\int_{\mathcal{Z}_f^{t,r}} \psi_f^{t,\text{post}}(z)^{\eta-1}\,dz}\right) \;\geq\; 0$: log expansion of effective task content from reinstatement of new tasks under channel (ii).
- $\alpha_{\text{prod}}^{(o,s,f)} \;=\; \ln\left(\dfrac{\int_{\mathcal{Z}_f^{t,r}} \psi_f^{t,\text{post}}(z)^{\eta-1}\,dz}{\int_{\mathcal{Z}_f^{t,r}} \psi_f^t(z)^{\eta-1}\,dz}\right)$: log change in integrated effective task content on the retained domain from per-task productivity gains under channel (iii). When $\psi_f^{t,\text{post}}(z) \geq \psi_f^t(z)$ pointwise on $\mathcal{Z}_f^{t,r}$, this term is ≥ 0 for $\eta > 1$ (since $\psi^{\eta-1}$ is increasing in ψ) and ≤ 0 for $\eta < 1$ (since $\psi^{\eta-1}$ is decreasing).

For this fixed (t,f), abbreviate the four integrated effective task-content terms entering the three components as

$$P \equiv \int_{\mathcal{Z}_f^{t,\text{pre}}} \psi_f^{t\,\eta-1}, \qquad R_{\text{pre}} \equiv \int_{\mathcal{Z}_f^{t,r}} \psi_f^{t\,\eta-1},$$
$$R_{\text{post}} \equiv \int_{\mathcal{Z}_f^{t,r}} \psi_f^{t,\text{post}\,\eta-1}, \qquad Q \equiv \int_{\mathcal{Z}_f^{t,\text{post}}} \psi_f^{t,\text{post}\,\eta-1},$$

so that $\alpha_{\text{disp}} = \ln(P/R_{\text{pre}})$, $\alpha_{\text{reinst}} = \ln(Q/R_{\text{post}})$, $\alpha_{\text{prod}} = \ln(R_{\text{post}}/R_{\text{pre}})$, and $\alpha^{(o,s,f)} = \ln(Q/P)$ by (4).

Lemma 1 (Additive decomposition of the shifter)**.** *Let $\eta > 0$, $\eta \neq 1$, and suppose P, R_{pre}, R_{post}, Q are finite and strictly positive—equivalently, given the bounds on ψ, that the retained domain has positive measure, $|\mathcal{Z}_f^{t,r}| > 0$. Then:*

$$\alpha_{reinst}^{(o,s,f)} \;-\; \alpha_{disp}^{(o,s,f)} \;+\; \alpha_{prod}^{(o,s,f)} \;=\; \alpha^{(o,s,f)}.$$

Proof. All four quantities lie in $(0, \infty)$, so each log of a ratio expands as a difference of logs:

$$\begin{aligned}\alpha_{\text{reinst}} - \alpha_{\text{disp}} + \alpha_{\text{prod}} &= \big(\ln Q - \ln R_{\text{post}}\big) - \big(\ln P - \ln R_{\text{pre}}\big) + \big(\ln R_{\text{post}} - \ln R_{\text{pre}}\big) \\ &= \ln Q - \ln P,\end{aligned}$$

since $\ln R_{\text{post}}$ and $\ln R_{\text{pre}}$ each appear twice with opposite signs. By (18), $Q = (\Psi_f^{t,\text{post}})^{\eta-1}$ and $P = (\Psi_f^{t,\text{pre}})^{\eta-1}$, so $\ln(Q/P) = \ln\big(\Psi_f^{t,\text{post}}/\Psi_f^{t,\text{pre}}\big)^{\eta-1} = \alpha^{(o,s,f)}$ by (4). □

A1.3 Derivation of the factor demands and proof of Proposition 1

Derivation of (6)–(7). The reduced-form production function (3) is a standard CES with effective input prices w_g/Ψ_f^g. Substitute $\tilde{x}^g \equiv \Psi_f^g x^g$ and $\tilde{w}^g \equiv w_g/\Psi_f^g$; the production function becomes the symmetric CES $y_f = [\sum_g (\tilde{x}^g)^{(\eta-1)/\eta}]^{\eta/(\eta-1)}$, and the cost objective becomes $\sum_g \tilde{w}^g \tilde{x}^g$. The standard CES cost-minimization formulas give $\tilde{x}^g = (\tilde{w}^g)^{-\eta} c_f^\eta y_f$ and $c_f = [\sum_g (\tilde{w}^g)^{1-\eta}]^{1/(1-\eta)}$. Substituting back yields (6) and (7).

For non-adopters, AI capital is not in the firm's factor set: the unit-cost sum in (7) is over the labor types $\mathcal{T}$ only, and $x_f^{k,NA} = 0$ trivially. □

Proof of Proposition 1. By (6) and (4),

$$\ell_f^{t,A} = e^{\alpha^{(o,s,f)}} (\Psi_f^{t,0})^{\eta-1} w_t^{-\eta} c_f^{A,\eta} y_f^A, \qquad \ell_{f'}^{t,NA} = (\Psi_{f'}^{t,0})^{\eta-1} w_t^{-\eta} c_{f'}^{NA,\eta} y_{f'}^{NA}.$$

The twin assumption gives $\Psi_f^{t,0} = \Psi_{f'}^{t,0}$ and the common-factor-price assumption gives the same w_t (and the same r, embedded in c_f^A) in both expressions, so

$$\frac{\ell_f^{t,A}}{\ell_{f'}^{t,NA}} = e^{\alpha^{(o,s,f)}} \left(\frac{c_f^A}{c_{f'}^{NA}}\right)^\eta \frac{y_f^A}{y_{f'}^{NA}}.$$

By (8), $y_f^d = \xi_f \mu^{-\sigma} (c_f^d)^{-\sigma} P^{\sigma-1} E$. Common P and E together with $\xi_f = \xi_{f'}$ give $y_f^A/y_{f'}^{NA} = (c_f^A/c_{f'}^{NA})^{-\sigma}$, so the ratio is $e^{\alpha^{(o,s,f)}}(c_f^A/c_{f'}^{NA})^{\eta-\sigma} = e^{\alpha^{(o,s,f)}}(c_{f'}^{NA}/c_f^A)^{\sigma-\eta}$, which is (11). □

A2 Labor Supply and Equilibrium

This appendix closes the model: we specify workers and capitalists, derive labor supply, define equilibrium, and prove existence. Throughout the appendices, $d_f \in \{A, NA\}$ denotes firm f's adoption status; $\Psi_f^{t,A}$ and $\Psi_f^{t,NA} \equiv \Psi_f^{t,0}$ are its post- and pre-adoption effective productivities, so that $(\Psi_f^{t,A})^{\eta-1} = e^{\alpha(t,f)}(\Psi_f^{t,0})^{\eta-1}$ by (4); and c_f^d, y_f^d, π_f^d, and $x_f^{g,d}$ are the unit cost, output, variable profit, and factor demands of Section 2.3 under status d. Appendix A3 characterizes the general-equilibrium response to adoption to first order; Appendix A3.6 solves the nonlinear system numerically.

A2.1 Workers and capitalists

There is a population N^s of workers of each seniority $s \in \mathcal{S}$. A worker i of seniority s chooses an option j from the market occupations $\mathcal{O}$ and the outside option $j = 0$. Choosing a market occupation o, the worker supplies one unit of type-(o, s) labor, is paid the wage $w_{o,s}$ of Section 2.3, and consumes $w_{o,s}/P$ units of the bundle (1). Choosing the outside option, the worker supplies no market labor and engages in home production of value $w_{0,s}$, an exogenous quantity denominated in the same nominal units as wages, worth $w_{0,s}/P$ units of the bundle.[28] Utility from option j is consumption times an option-specific taste shock $z_i^{(j,s)}$:

$$u_i(j, s) \;=\; z_i^{(j,s)} \cdot w_{j,s}/P. \tag{21}$$

The taste shocks $(z_i^{(j,s)})_{j\in\mathcal{O}\cup\{0\}}$ are independent across options and workers, with Fréchet marginals

$$\Pr\big(z_i^{(j,s)} \le z\big) \;=\; \exp\big(-T_{j,s} \cdot z^{-\zeta}\big), \tag{22}$$

with shape $\zeta > 0$ and scales $T_{j,s} > 0$. Throughout, j indexes options in $\mathcal{O} \cup \{0\}$, while a cell (o, s) always has $o \in \mathcal{O}$: the set $\mathcal{T} = \mathcal{O} \times \mathcal{S}$ contains market cells only, and the outside option of seniority s is written $(0, s)$.

Firms are owned by capitalists, who supply no labor. AI capital is supplied perfectly elastically at the exogenous rental rate r, and capitalists receive the rental flow rK on the units rented, $K = \int \mathbf{1}_{d_f=A}\, k_f\, df$, and aggregate dividends $\Pi - \bar{\mathcal{F}}$, where $\Pi = \int \pi_f\, df$ is total variable profit (9) and $\bar{\mathcal{F}} = \int \mathbf{1}_{d_f=A}\, \mathcal{F}_f\, df$; adoption fixed costs are spent on goods, an adopting firm purchasing $\mathcal{F}_f/P$ units of the consumption bundle. Capitalists spend their income on the bundle.

[28] Denominating the outside option in the same units as wages, and hence as sectoral expenditure, means that its real value $w_{0,s}/P$ rises one-for-one with real output E/P: a technology that lowers the price of market output raises the value of non-employment along with real wages. This is the treatment in Hsieh et al. (2019), where home production is one more sector whose return is determined jointly with market wages and whose output counts toward total output, so that participation depends only on market returns relative to the home sector.

All spending on the bundle—worker consumption, capitalist consumption, and adopter fixed costs—is allocated across varieties to maximize (1), so aggregate variety demand is (2) with E equal to total spending. Workers employed in the sector spend $\sum_{(o,s)} w_{o,s} L_{o,s}$, where $L_{o,s}$ is the mass of workers in cell (o, s) (Appendix A2.2); capitalists spend $rK + \Pi - \bar{\mathcal{F}}$; adopters spend $\bar{\mathcal{F}}$. Total spending on the bundle therefore equals total market income:

$$E \;=\; \sum_{(o,s)\in\mathcal{T}} w_{o,s}\, L_{o,s} \;+\; rK \;+\; \Pi. \tag{23}$$

We normalize $E = 1$; the remaining nominal primitives—$w_{0,s}$, r, and the adoption costs $\mathcal{F}_f$—are thereby denominated in units of sectoral expenditure.

A2.2 Occupation choice and labor supply

Worker i chooses the option maximizing utility (21), that is, maximizing $z_i^{(j,s)} w_{j,s}$.

Lemma 2 (Fréchet labor supply)**.** *The mass of seniority-s workers choosing market occupation o is*

$$L_{o,s} \;=\; \frac{T_{o,s}\, w_{o,s}^{\zeta}}{D^s} \cdot N^s, \qquad D^s \;\equiv\; \sum_{j\in\mathcal{O}\cup\{0\}} T_{j,s}\, w_{j,s}^{\zeta}. \tag{24}$$

Aggregate seniority-s market employment is $L^s = \sum_{o\in\mathcal{O}} L_{o,s} = (1 - \pi_0^s) N^s$, where $\pi_0^s = T_{0,s}\, w_{0,s}^{\zeta} / D^s$ is the outside-option share.

Proof. The value of option j is $V_j \equiv z^{(j,s)} w_{j,s}$, with $\Pr(V_j \le v) = \exp(-S_j v^{-\zeta})$ where $S_j \equiv T_{j,s}\, w_{j,s}^{\zeta}$: Fréchet with scale S_j and shape ζ. Write $S \equiv \sum_j S_j = D^s$. For any j and $x \ge 0$, using $dF_j(v) = S_j\, \zeta v^{-\zeta-1} e^{-S_j v^{-\zeta}} dv$ and independence across options,

$$\begin{aligned}
\Pr\bigl(j \text{ wins},\ V_j \le x\bigr) &= \int_0^x \prod_{j'\neq j} e^{-S_{j'} v^{-\zeta}}\; dF_j(v) \\
&= \int_0^x S_j\, \zeta v^{-\zeta-1}\, e^{-S v^{-\zeta}}\, dv \;=\; \frac{S_j}{S}\, e^{-S x^{-\zeta}}.
\end{aligned}$$

Letting $x \to \infty$ gives $\Pr(j \text{ wins}) = S_j/S$; multiplying by N^s gives (24), and summing over $o \in \mathcal{O}$ gives L^s. $\square$

Each worker supplies one unit of labor, so $L_{o,s}$ is simultaneously the cell's employment count, its labor input in production, and its wage bill divided by $w_{o,s}$. The supply primitives $\{T_{j,s}, w_{0,s}, N^s\}$ are fixed across the AI and no-AI regimes; adoption affects labor supply only through the market wages.

A2.3 Equilibrium definition

Aggregate firm demand for type (o, s), integrating the cost-minimizing factor demands (6) across firms, is

$$L^D_{o,s} \;=\; \int_{\mathcal{F}} \ell_f^{(o,s),d_f}\, df \;=\; w_{o,s}^{-\eta} \int_{\mathcal{F}} \big(\Psi_f^{(o,s),d_f}\big)^{\eta-1} \big(c_f^{d_f}\big)^{\eta} y_f^{d_f}\, df. \tag{25}$$

Definition 1. *An* equilibrium *consists of wages $\{w_{o,s}\}_{(o,s)\in\mathcal{T}}$, a sectoral price index P, worker masses $\{L_{o,s}\}_{(o,s)\in\mathcal{T}}$, and a measurable profile of adoption decisions, prices, outputs, and factor demands $\{d_f, p_f, y_f, (\ell^t_f)_t, k_f\}_{f\in\mathcal{F}}$, with $E = 1$ (Appendix A2.1), satisfying the optimality conditions (i)–(v) and the market-clearing conditions (vi)–(vii):*

(i) Output equals product demand (2)*: $y_f = \xi_f\, p_f^{-\sigma}\, P^{\sigma-1} E$.*

(ii) Factor demands minimize the cost of producing y_f given d_f: every factor input satisfies (6)*, with unit cost $c_f^{d_f}$ given by* (7)*.*

(iii) Prices carry the CES markup: $p_f = \mu\, c_f^{d_f}$ (Section 2.3).

(iv) Adoption is optimal for almost every firm: $d_f = A$ iff $\pi^A_f - \pi^{NA}_f \geq \mathcal{F}_f$, the cutoff (10)*.*

(v) Workers choose the option maximizing utility (21)*, so the worker masses are* (24) *(Lemma 2): $L_{o,s} = T_{o,s}\, w^{\zeta}_{o,s}\, N^s/D^s$.*

(vi) The goods market clears: spending on varieties exhausts expenditure,

$$\int_{\mathcal{F}} p_f\, y_f\, df \;=\; E.$$

Substituting (i), $\int p_f y_f\, df = E\, P^{\sigma-1} \int \xi_f\, p_f^{1-\sigma} df$, so the condition holds iff the price index is consistent with posted prices:

$$P^{1-\sigma} \;=\; \int_{\mathcal{F}} \xi_f\, p_f^{1-\sigma}\, df. \tag{26}$$

(vii) Every type-(o, s) labor market clears: $L^D_{o,s} = L_{o,s}$, with demand $L^D_{o,s}$ from (25) *and the worker masses from (v).*

The conditions match the unknowns one for one. Given $(\{w_{o,s}\}, P, \{d_f\})$, conditions (i)–(iii) and (v) determine every remaining object uniquely: (7) gives $c_f^{d_f}$, (iii) gives p_f, (i) gives y_f, (ii) gives the factor demands, and (v) gives the worker masses. The remaining conditions then restrict the remaining unknowns: (iv) restricts each

d_f, (vi) restricts P, and (vii) is $|\mathcal{T}|$ restrictions on the $|\mathcal{T}|$ wages. Proposition 2 shows a solution exists.

Two candidate conditions are absent because the listed conditions imply them. Firm-level goods clearing (firm f produces the quantity demanded at its posted price) is condition (i) itself: a price-setting firm chooses a point on its demand curve and has no separate supply schedule. The income identity (23) is Walras' law rather than an independent condition. Under constant returns each firm's total cost equals its factor payments, so aggregate revenue is $\int_{\mathcal{F}} p_f y_f \, df = \sum_{(o,s)\in\mathcal{T}} w_{o,s} L^D_{o,s} + rK + \Pi$. Subtracting (23) with $E = 1$ gives $\int_{\mathcal{F}} p_f y_f \, df - 1 = \sum_{(o,s)\in\mathcal{T}} w_{o,s}\big(L^D_{o,s} - L_{o,s}\big)$. The identity therefore holds once the goods market clears (vi) and every labor market clears (vii).

A2.4 Existence of equilibrium

Assumption 1.

(a) Elasticities and supply primitives: $\eta > 0$, $\eta \neq 1$; $\sigma > \max\{1, \eta\}$; $\zeta > 0$; $r > 0$; and, for every seniority s, $N^s > 0$, $w_{0,s} > 0$, and $T_{j,s} > 0$ for all $j \in \mathcal{O} \cup \{0\}$.

(b) Firm primitives are measurable functions of f with uniform bounds: $\mathcal{F}_f \geq 0$; $\xi_f \in [\underline{\xi}, \bar{\xi}] \subset (0, \infty)$ with $\int \xi_f \, df = 1$; $\Psi^{t,d}_f \in [\underline{\Psi}, \bar{\Psi}] \subset (0, \infty)$ for all t and both adoption statuses d (baseline productivities and labor-demand shifters are bounded); and $\Psi^{k,A}_f \in [\underline{\Psi}, \bar{\Psi}]$.

(c) No atom of indifferent adopters: for every wage vector $\mathbf{w} \in \mathbb{R}^{|\mathcal{T}|}_{++}$ and price index $P > 0$, the set $\{f : \pi^A_f(\mathbf{w}, P) - \pi^{NA}_f(\mathbf{w}, P) = \mathcal{F}_f\}$ has Lebesgue measure zero. This holds if, conditional on the production primitives, $\mathcal{F}_f$ is continuously distributed across firms.

Fix a compact box $\mathcal{W}_{\text{box}} = \prod_{(o,s)\in\mathcal{T}} [w^{\min}, w^{\max}] \subset \mathbb{R}^{|\mathcal{T}|}_{++}$ of wage vectors, with arbitrary $0 < w^{\min} < w^{\max}$. Two consequences of Assumption 1(b) are used repeatedly. Let κ_1 and κ_2 be the smaller and the larger of $\underline{\Psi}^{\eta-1}$ and $\bar{\Psi}^{\eta-1}$. Since $(w_g/\Psi^g_f)^{1-\eta} = w_g^{1-\eta}(\Psi^g_f)^{\eta-1}$, every labor term in the unit cost (7) satisfies $\kappa_1 w_g^{1-\eta} \leq (w_g/\Psi^g_f)^{1-\eta} \leq \kappa_2 w_g^{1-\eta}$. And on $\mathcal{W}_{\text{box}}$, every firm's unit cost at its realized adoption status d_f satisfies

$$c_{\text{lo}} \;\leq\; c^d_f \;\leq\; c_{\text{hi}} \;\equiv\; |\mathcal{T}|^{1/(1-\eta)} \, \frac{w^{\max}}{\underline{\Psi}}, \tag{27}$$

for some constant $c_{\text{lo}} > 0$ depending only on the box and the primitive bounds. For the upper bound: c^{NA}_f is increasing in every wage and decreasing in every productivity, so it is at most its value at $w_g = w^{\max}$ and $\Psi^g_f = \underline{\Psi}$ for all g, which is c_{hi}; and an adopter has $c^A_f \leq c^{NA}_f$, because the cutoff (10) with $\mathcal{F}_f \geq 0$ requires $(c^A_f)^{1-\sigma} \geq (c^{NA}_f)^{1-\sigma}$. For the lower bound: c^{NA}_f is at least its value at $w_g = w^{\min}$ and

$\Psi_f^g = \bar{\Psi}$; for c_f^A, when $\eta < 1$ the AI term is a positive addition to the labor sum in (7), and when $\eta > 1$ the AI term is at most $(r/\bar{\Psi})^{1-\eta}$, so the sum is bounded above and c_f^A, with the negative exponent $1/(1-\eta)$, is bounded below.

An equilibrium is a wage vector at which excess demand vanishes, but the price index and the adoption decisions are simultaneously determined at any candidate wage vector. We first resolve that inner fixed point in the price index and adoption decisions. At given $(P, \mathbf{w})$ the cutoff (10) determines adoption firm by firm, so the inner problem reduces to a fixed point in the scalar index alone. Adoption incentives increase in P (the right-hand side of the cutoff (10) is decreasing in P), while more adoption lowers P (adopters price lower). Lemma 3 shows that the inner fixed point exists and is unique.

Lemma 3 (Price-index/adoption fixed point)**.** *For* $\mathbf{w} \in \mathcal{W}_{\text{box}}$ *and* $P > 0$*, let* $d_f(P, \mathbf{w}) = A$ *iff the cutoff condition* (10) *holds at* $(\mathbf{w}, P)$*, and let*

$$G(P, \mathbf{w}) \equiv \left[\int_{\mathcal{F}} \xi_f \big(\mu\, c_f^{d_f(P,\mathbf{w})}(\mathbf{w})\big)^{1-\sigma} df \right]^{1/(1-\sigma)}$$

be the price index generated by the prices firms post when they take the index to be P*. Then:*

(i) Boundedness: *there exist* $0 < P_{\text{lo}} \leq P_{\text{hi}} < \infty$*, independent of* $\mathbf{w} \in \mathcal{W}_{\text{box}}$*, such that* $G(P, \mathbf{w}) \in [P_{\text{lo}}, P_{\text{hi}}]$ *for all* $P > 0$*.*

(ii) Continuity: *G is jointly continuous in* $(P, \mathbf{w})$*.*

(iii) Monotonicity: *G is weakly decreasing in P.*

Consequently, for each $\mathbf{w} \in \mathcal{W}_{\text{box}}$ *there is exactly one* $P^*(\mathbf{w})$ *with* $G(P^*(\mathbf{w}), \mathbf{w}) = P^*(\mathbf{w})$*: existence follows from (i)–(ii), uniqueness from (iii).*

Proof. Boundedness (i). Set $P_{\text{lo}} = \mu c_{\text{lo}}$ and $P_{\text{hi}} = \mu c_{\text{hi}}$ from the cost bounds (27). Since $\int \xi_f df = 1$, G is a generalized $(1-\sigma)$-mean of the values $\mu c_f^{d_f} \in [\mu c_{\text{lo}}, \mu c_{\text{hi}}]$, which gives (i).

Continuity (ii). Both $c_f^A(\mathbf{w})$ and $c_f^{NA}(\mathbf{w})$ are continuous in $\mathbf{w}$ for every f. Take $(P_n, \mathbf{w}_n) \to (P, \mathbf{w})$. The integrand converges for every firm whose adoption decision is locally constant, i.e., every firm not exactly indifferent at $(P, \mathbf{w})$—a set of full measure by Assumption 1(c)—because costs are continuous and the cutoff comparison is strict for such firms. The integrand is uniformly bounded, so dominated convergence gives $G(P_n, \mathbf{w}_n) \to G(P, \mathbf{w})$.

Monotonicity (iii). The cutoff (10) has right-hand side proportional to $P^{-(\sigma-1)}$, which is decreasing in P (as $\sigma > 1$), so the set of adopters $\{f : d_f = A\}$ weakly expands as P rises. Every firm that switches into adoption weakly lowers its price: the cutoff with $\mathcal{F}_f \geq 0$ requires $(c_f^A)^{1-\sigma} \geq (c_f^{NA})^{1-\sigma}$, hence $c_f^A \leq c_f^{NA}$ since $\sigma > 1$.

A lower price weakly raises the integrand $\xi_f(\mu c)^{1-\sigma}$ (as $1-\sigma<0$), hence weakly raises $G^{1-\sigma}$ and weakly lowers G.

Fixed point. Fix $\mathbf{w}$ and let $h(P) \equiv G(P,\mathbf{w}) - P$ on $[P_{\mathrm{lo}}, P_{\mathrm{hi}}]$. By (i), $h(P_{\mathrm{lo}}) = G(P_{\mathrm{lo}},\mathbf{w}) - P_{\mathrm{lo}} \geq 0$ and $h(P_{\mathrm{hi}}) \leq 0$; by (ii), h is continuous; the intermediate value theorem gives a root—existence uses only (i) and (ii). By (iii), h is *strictly* decreasing (weakly decreasing G minus strictly increasing identity), so the root is unique. □

Given $\mathbf{w}$, write $d_f(\mathbf{w}) \equiv d_f(P^*(\mathbf{w}),\mathbf{w})$ and define excess demand

$$Z_{o,s}(\mathbf{w}) \;\equiv\; L^D_{o,s}(\mathbf{w}) - L^S_{o,s}(\mathbf{w}), \tag{28}$$

with L^D from (25) evaluated at $(P^*(\mathbf{w}), d_f(\mathbf{w}))$ and $E=1$, and L^S the supply (24). The outer step requires these solved objects to vary continuously with wages; this is a separate claim from the inner loop, and we record it next.

Corollary 1 (Continuity of the inner solution)**.** *$P^*(\cdot)$ and $Z(\cdot)$ are continuous on $\mathcal{W}_{\mathrm{box}}$.*

Proof. P^*. Let $\mathbf{w}_n \to \mathbf{w}$ and set $P_n = P^*(\mathbf{w}_n) \in [P_{\mathrm{lo}}, P_{\mathrm{hi}}]$. Consider any convergent subsequence $P_{n_j} \to \bar{P}$. By definition of P^*, each term satisfies $G(P_{n_j}, \mathbf{w}_{n_j}) = P_{n_j}$. As $j \to \infty$, $(P_{n_j}, \mathbf{w}_{n_j}) \to (\bar{P}, \mathbf{w})$, so the left side converges to $G(\bar{P},\mathbf{w})$ by the joint continuity of G (Lemma 3(ii)), while the right side converges to $\bar{P}$; hence $G(\bar{P},\mathbf{w}) = \bar{P}$. By Lemma 3 the fixed point at $\mathbf{w}$ is unique, so $\bar{P} = P^*(\mathbf{w})$.

Next, every subsequence of the bounded sequence $\{P_n\}$ has a further subsequence converging to $P^*(\mathbf{w})$, so $P_n \to P^*(\mathbf{w})$ by the subsubsequence convergence criterion. To see this, suppose, for contradiction, that $P_n \not\to P^*(\mathbf{w})$. Then there is an $\varepsilon > 0$ and a subsequence $\{P_{n_k}\}$ whose every term stays at distance $\geq \varepsilon$ from $P^*(\mathbf{w})$. The subsequence $\{P_{n_k}\}$ lies in the compact interval $[P_{\mathrm{lo}}, P_{\mathrm{hi}}]$, so by the Bolzano–Weierstrass theorem it has a further convergent subsequence, which converges to $P^*(\mathbf{w})$ by the previous paragraph. But every term of this sub-subsequence is at least ε away from $P^*(\mathbf{w})$, so its limit is too—a contradiction. Hence $P_n \to P^*(\mathbf{w})$.

Z. As $\mathbf{w}_n \to \mathbf{w}$, $(P^*(\mathbf{w}_n), \mathbf{w}_n) \to (P^*(\mathbf{w}), \mathbf{w})$ by the previous step. For every firm not indifferent at $(P^*(\mathbf{w}),\mathbf{w})$—a set of full measure by Assumption 1(c)—the adoption decision is locally constant, so the integrand of L^D in (25) converges firm by firm; it is uniformly bounded on $\mathcal{W}_{\mathrm{box}} \times [P_{\mathrm{lo}}, P_{\mathrm{hi}}]$ (the cost bounds (27) and $w_{o,s} \geq w^{\min}$), so dominated convergence gives $L^D_{o,s}(\mathbf{w}_n) \to L^D_{o,s}(\mathbf{w})$. The supply (24) is continuous in $\mathbf{w}$. Hence $Z(\mathbf{w}_n) \to Z(\mathbf{w})$. □

Existence is now a problem in wages alone. Conditions (i)–(iii) and (v) of Definition 1 hold at every candidate $(\mathbf{w}, P, \{d_f\})$, because they define the prices, outputs, factor demands, and worker masses. For every $\mathbf{w} \in \mathcal{W}_{\mathrm{box}}$, Lemma 3 selects $P^*(\mathbf{w})$ and $\{d_f(\mathbf{w})\}$ at which conditions (iv) and (vi) hold as well: $d_f(\mathbf{w})$ is optimal for almost every firm by construction, and $P^*(\mathbf{w}) = G(P^*(\mathbf{w}),\mathbf{w})$ is the index consistency (26). A wage vector $\mathbf{w}$, together with these selections and the allocation they define, is therefore an equilibrium if and only if condition (vii) holds, that is, $Z(\mathbf{w}) = 0$. Proposition 2 shows that such a wage vector exists. Its proof bounds excess demand through the *wage-bill identity*: from (6), factor (o,s)'s wage bill at firm f is its cost share times the firm's total cost, and total cost is revenue over the markup, while condition (vi) at $P^*(\mathbf{w})$ gives $\int_{\mathcal{F}} p_f y_f \, df = E = 1$:

$$w_{o,s}\, L^D_{o,s} \;=\; \int_{\mathcal{F}} \underbrace{\frac{(w_{o,s}/\Psi_f^{(o,s),d_f})^{1-\eta}}{\sum_g (w_g/\Psi_f^{g,d_f})^{1-\eta}}}_{\text{cost share } \bar{s}_f^{(o,s)}(\mathbf{w}) \in [0,1]} \cdot \frac{p_f y_f}{\mu}\, df. \tag{29}$$

Proposition 2 (Equilibrium existence)**.** *Under Assumption 1, an equilibrium in the sense of Definition 1 exists.*

Proof. For each cell $(o,s) \in \mathcal{T}$, let

$$\mathcal{W}^{\max}_{o,s} \;\equiv\; \{\mathbf{w} \in \mathcal{W}_{\text{box}} : w_{o,s} = w^{\max}\}, \qquad \mathcal{W}^{\min}_{o,s} \;\equiv\; \{\mathbf{w} \in \mathcal{W}_{\text{box}} : w_{o,s} = w^{\min}\}$$

be the two faces of the box on which cell (o,s)'s wage is at its ceiling or at its floor, the other wages being anywhere in $[w^{\min}, w^{\max}]$; the boundary of $\mathcal{W}_{\text{box}}$ is the union of these $2|\mathcal{T}|$ faces. We show that there are a floor $w^{\min}$ and a ceiling $w^{\max}$ such that, for every (o,s), supply exceeds demand in cell (o,s) on $\mathcal{W}^{\max}_{o,s}$ and demand exceeds supply on $\mathcal{W}^{\min}_{o,s}$. Since Z is continuous on the box (Corollary 1), the Poincaré–Miranda theorem then yields a zero of Z in the box—an equilibrium by the reduction above.

Step 1: the upper faces. Fix any $(o,s) \in \mathcal{T}$ and any $\mathbf{w} \in \mathcal{W}^{\max}_{o,s}$. From (29), since cost shares are at most one and $\int py = 1$,

$$L^D_{o,s} \;\le\; \frac{1}{\mu\, w^{\max}},$$

so demand is arbitrarily small for $w^{\max}$ large enough. For supply, let $w^{\max}\mathbf{1}$ denote the wage vector with every market wage equal to $w^{\max}$. In the labor supply equation (24), the denominator D^s is increasing in every market wage and every market wage in $\mathbf{w}$ is at most $w^{\max}$, so $D^s(\mathbf{w}) \le D^s(w^{\max}\mathbf{1})$. The numerator $T_{o,s}(w^{\max})^{\zeta} N^s$ is the same at $\mathbf{w}$ and at $w^{\max}\mathbf{1}$. Hence

$$\begin{aligned} L^S_{o,s}(\mathbf{w}) \;\ge\; L^S_{o,s}(w^{\max}\mathbf{1}) &= \frac{T_{o,s}\,(w^{\max})^{\zeta}}{T_{0,s}\, w_{0,s}^{\zeta} + (w^{\max})^{\zeta} \sum_{o' \in \mathcal{O}} T_{o',s}}\, N^s \\ &\xrightarrow[w^{\max}\to\infty]{} \; \frac{T_{o,s}}{\sum_{o'\in\mathcal{O}} T_{o',s}}\, N^s \;>\; 0, \end{aligned}$$

so supply is bounded away from zero, uniformly on $\mathcal{W}^{\max}_{o,s}$, for $w^{\max}$ large enough. Choose $w^{\max}$ large enough that $L^D_{o,s} < L^S_{o,s}$ on $\mathcal{W}^{\max}_{o,s}$ for every (o,s); the choice does not involve $w^{\min}$. Then $Z_{o,s} < 0$ on $\mathcal{W}^{\max}_{o,s}$ for every (o,s).

Step 2: the lower faces. Fix any $(o,s) \in \mathcal{T}$ and any $\mathbf{w} \in \mathcal{W}^{\min}_{o,s}$. Supply satisfies $L^S_{o,s} \le N^s$ always (by (24), in fact $L^S_{o,s} \to 0$ as $w^{\min} \to 0$, the home-production value $w_{0,s}$ being a primitive), so it suffices to show that $L^D_{o,s}$ exceeds $\max_s N^s$ on $\mathcal{W}^{\min}_{o,s}$ once $w^{\min}$ is small. By (29) with $\int py = 1$,

$$L^D_{o,s} \;\ge\; \frac{\inf_f \bar{s}^{(o,s)}_f}{\mu\, w^{\min}},$$

so it suffices to bound the cost share of the cheap factor from below, uniformly over firms and over the face. Each labor term in the cost share is the product of a wage part and a productivity part, $(w_g/\Psi^g_f)^{1-\eta} = w_g^{1-\eta}\,(\Psi^g_f)^{\eta-1}$, and the productivity part lies in $[\kappa_1, \kappa_2]$ by (27).

Case $\eta < 1$. The denominator of the cost share in (29) is $(c^{d_f}_f)^{1-\eta}$, and $c^{d_f}_f \le c_{\text{hi}}$ by (27), where c_{hi} depends on $w^{\max}$ but not on $w^{\min}$, so

$$\bar{s}^{(o,s)}_f \;\ge\; \frac{\kappa_1\,(w^{\min})^{1-\eta}}{c_{\text{hi}}^{1-\eta}}.$$

Case $\eta > 1$. The share is the own term divided by the sum of all terms, so it is increasing in the own term and decreasing in every other term; a lower bound on the own term and upper bounds on the other terms therefore give a lower bound on the share. For the own term the wage part is exactly $(w^{\min})^{1-\eta}$, because $w_{o,s} = w^{\min}$ on the face, so the own term is at least $\kappa_1 (w^{\min})^{1-\eta}$. For a competing labor term, w_g can be anywhere in $[w^{\min}, w^{\max}]$; since $1-\eta < 0$, the wage part $w_g^{1-\eta}$ is decreasing in w_g and is largest at $w_g = w^{\min}$—every competitor as cheap as the own factor—so each of the $|\mathcal{T}| - 1$ competing terms is at most $\kappa_2 (w^{\min})^{1-\eta}$. The AI term, present only at adopters, is at most $\kappa_3 \equiv (r/\bar{\Psi})^{1-\eta}$, since $\Psi^{k,A}_f \le \bar{\Psi}$. Hence

$$\begin{aligned}
\bar{s}^{(o,s)}_f \;&\ge\; \frac{\kappa_1 (w^{\min})^{1-\eta}}{\kappa_1 (w^{\min})^{1-\eta} + (|\mathcal{T}|-1)\,\kappa_2 (w^{\min})^{1-\eta} + \kappa_3} \\
&=\; \frac{\kappa_1}{\kappa_1 + (|\mathcal{T}|-1)\,\kappa_2 + \kappa_3\,(w^{\min})^{\eta-1}},
\end{aligned}$$

dividing numerator and denominator by $(w^{\min})^{1-\eta}$. We restrict the floor to $w^{\min} \le 1$; the floor is ours to choose, and it is sent to zero below. Then $(w^{\min})^{\eta-1} \le 1$, so the denominator is at most $\kappa_1 + (|\mathcal{T}|-1)\,\kappa_2 + \kappa_3$, and

$$\bar{s}^{(o,s)}_f \;\ge\; \frac{\kappa_1}{\kappa_1 + (|\mathcal{T}|-1)\,\kappa_2 + \kappa_3} \;\equiv\; s_{\min} \;>\; 0.$$

Substituting each case's share bound into $L^D_{o,s} \ge \inf_f \bar{s}^{(o,s)}_f/(\mu w^{\min})$ gives

$$L^D_{o,s} \;\ge\; \frac{\kappa_1}{\mu\, c_{\text{hi}}^{1-\eta}}\,(w^{\min})^{-\eta} \quad (\eta < 1), \qquad L^D_{o,s} \;\ge\; \frac{s_{\min}}{\mu}\,(w^{\min})^{-1} \quad (\eta > 1),$$

and both bounds diverge as $w^{\min} \to 0$. Choose $w^{\min} \leq 1$ small enough that the relevant bound exceeds $\max_s N^s$; then $Z_{o,s} > 0$ on $\mathcal{W}^{\min}_{o,s}$ for every (o, s).

Step 3: a zero of excess demand. On the box $\mathcal{W}_{\text{box}}$ so constructed, $Z_{o,s} > 0$ on $\mathcal{W}^{\min}_{o,s}$ and $Z_{o,s} < 0$ on $\mathcal{W}^{\max}_{o,s}$ for every $(o, s) \in \mathcal{T}$ (Steps 1–2), and Z is continuous on $\mathcal{W}_{\text{box}}$ (Corollary 1). By the Poincaré–Miranda theorem,[29] Z has a zero $\mathbf{w}^* \in \mathcal{W}_{\text{box}}$.

Step 4: the zero is an equilibrium. $Z(\mathbf{w}^*) = 0$ says that all labor markets clear, which is condition (vii) of Definition 1. Conditions (i)–(vi) hold at $\mathbf{w}^*$ by the reduction preceding the proposition, so $\mathbf{w}^*$, together with $P^*(\mathbf{w}^*)$, $\{d_f(\mathbf{w}^*)\}$, and the allocation they define, is an equilibrium. □

A3 Matched Differences and General-Equilibrium Changes

Proposition 1 compares an adopter with a non-adopting twin at common wages. The general-equilibrium change in cell employment when a set of firms adopts is a different object: wages, the price index, and the outside-option margin adjust, and the adjustment affects adopters and non-adopters alike. This appendix relates the two to first order; Section 2.5 summarizes the result. For simplicity, we restrict attention to a single seniority, $|\mathcal{S}| = 1$, and drop the index s: cells are occupations $o \in \mathcal{O}$, the Fréchet denominator of (24) is D, and the outside-option share of Lemma 2 is $\pi_0 = T_0 w_0^{\zeta}/D$, the fraction of workers who choose the outside option.

A3.1 The perturbation

The *no-AI equilibrium* is the equilibrium of Appendix A2 in which AI is unavailable, so that no firm adopts; a bar denotes a quantity's value there, so $\bar{\mathbf{w}}$ is the no-AI wage vector. We linearize around it, along the following family of economies.

Definition 2 (The ϵ-economies)**.** *For $\epsilon \in [0, \bar{\epsilon})$, the* ϵ-economy *differs from the no-AI economy only in adoption costs and in the technology available to adopters:*

(i) Who can adopt. *A set $\mathcal{A} \subset \mathcal{F}$ of positive measure has adoption cost $\mathcal{F}_f = 0$; every other firm has an adoption cost exceeding any attainable gain.*

(ii) Labor-demand shifters. *Adopters' shifters are $\alpha^{(o,f)}(\epsilon) = \epsilon\, a^{(o,f)}$, with a measurable and bounded.*

[29] The Poincaré–Miranda theorem: if $f : \prod_{i=1}^{n}[a_i, b_i] \to \mathbb{R}^n$ is continuous and, for every i, f_i takes opposite signs (or vanishes) on the two opposite faces $x_i = a_i$ and $x_i = b_i$, then f has a zero. It is the n-dimensional intermediate value theorem, stated and proved by Poincaré in 1883–1886; Miranda (1940) showed it equivalent to Brouwer's fixed-point theorem, and Kulpa (1997) gives a direct proof.

(iii) AI cost term. *Adopters' AI cost term in the unit cost* (7) *is* $\epsilon\, k_f$*, with* $k_f \geq 0$ *measurable and bounded (*$k_f = 0$ *meaning no AI capital).*

(iv) Adoption lowers cost. *With* $\bar{s}^o_f \equiv (\bar{w}_o/\Psi^{o,0}_f)^{1-\eta}/(c^{NA}_f(\bar{\mathbf{w}}))^{1-\eta}$ *firm* f*'s baseline cost share of occupation* o*, for all* $f \in \mathcal{A}$ *and some* $\underline{b} > 0$*,*

$$b_f \;\equiv\; \frac{1}{\eta - 1}\Big[\sum_{o \in \mathcal{O}} \bar{s}^o_f\, a^{(o,f)} + \frac{k_f}{(c^{NA}_f(\bar{\mathbf{w}}))^{1-\eta}}\Big] \;\geq\; \underline{b} \;>\; 0. \tag{30}$$

Condition (iv) ensures that adoption lowers cost to first order, so that the firms in $\mathcal{A}$ adopt.

A3.2 The decomposition

The decomposition that follows in Proposition 3 compares two objects. The first object compared is the matched-DiD estimand. By Proposition 1, an adopter and its twin differ in cell-o employment, to first order, by

$$\alpha^{(o,f)} + (\sigma - \eta)\ln\big(c^{NA}_f / c^A_f\big) \;=\; \alpha^{(o,f)} + (\sigma - \eta)\,\delta_f,$$

where $\delta_f \equiv \ln c^{NA}_f(\bar{\mathbf{w}}) - \ln c^A_f(\bar{\mathbf{w}};\epsilon)$ is the adopter's log unit-cost reduction at the no-AI wages. Averaged over adopters and weighted by employment, the matched-DiD estimand[30] for cell o is

$$\tau_o \;\equiv\; \frac{1}{\rho_o}\int_{\mathcal{A}} \bar{\lambda}^o_f \big[\alpha^{(o,f)} + (\sigma - \eta)\,\delta_f\big]\, df, \tag{31}$$

where $\bar{\lambda}^o_f \equiv \bar{\ell}^o_f/\bar{L}_o$ is firm f's share density of occupation-o employment and $\rho_o \equiv \int_{\mathcal{A}} \bar{\lambda}^o_f\, df$ is adopters' total share of baseline occupation employment.

The second object is the general-equilibrium change in cell employment. For this, define $\omega_f \;\equiv\; \sum_{o' \in \mathcal{O}} \bar{s}^{o'}_f\, d\ln w_{o'}$, the cost-share-weighted average of the wage changes, which is how wages move firm f's unit cost to first order; its average over the firms employing occupation o, $\omega_o \equiv \int \bar{\lambda}^o_f\, \omega_f\, df$; and $d\ln D$ and $d\ln P$, the first order changes in the Fréchet denominator of (24) and in the price index.

Proposition 3 (Matched-DiD estimand and general-equilibrium change)**.**

(a) Fix a no-AI equilibrium with wages $\bar{\mathbf{w}}$*. For all sufficiently small* $\epsilon > 0$*, the* ϵ*-economy has an equilibrium with adopter set* $\mathcal{A}$ *and wages* $\mathbf{w}(\epsilon)$*, where* $\mathbf{w}(\cdot)$ *is continuously differentiable and* $\mathbf{w}(0) = \bar{\mathbf{w}}$*. Its price index and allocation are those that conditions (i)–(iii), (v), and (vi) of Definition 1 assign to* $\mathbf{w}(\epsilon)$*. Along this path, every equilibrium quantity is continuously differentiable in* ϵ*.*

[30] The matched estimator of the empirical section weights each treated firm equally (controls receive weight $1/k$ per matched pair), whereas τ_o weights adopters by their baseline share of cell employment, the weighting under which adopters' contributions add up to the change in cell employment.

(b) Along this path,

$$\begin{aligned} d\ln L_o \;&=\; \frac{\zeta}{\eta+\zeta}\Big[\rho_o\,\tau_o \;-\; (\sigma-\eta)\,\omega_o\Big] \;+\; \Gamma, \\ \Gamma \;&\equiv\; \frac{\zeta\,(\sigma-1)}{\eta+\zeta}\,d\ln P \;-\; \frac{\eta}{\eta+\zeta}\,d\ln D. \end{aligned} \tag{32}$$

The proof is in Appendix A3.5. The first term is the matched-DiD estimand scaled by $\rho_o\,\zeta/(\eta+\zeta) < 1$: adopters are a fraction ρ_o of the cell, and their additional demand raises the cell's wage, which moves employment only through the supply elasticity ζ against the demand elasticity η. The second term is the only other occupation-specific one: a rise in the wages that occupation-o employers pay, $\omega_o > 0$, raises their unit costs and lowers their sales, and since $\sigma > \eta$ the lost scale reduces occupation-o employment. The remaining term Γ is common to every occupation: the price-index term and the outside-option term shift every cell's employment alike. The second and third terms contain the equilibrium wage changes, so (32) does not by itself sign the difference between $d\ln L_o$ and τ_o.

A3.3 A closed-form case

Example 1 (Uniform productivity gain)**.** *In Definition 2, let adopters share the same technology and the same presence in every occupation:*

- *(a)* $a^{(o,f)} = (\eta-1)\,\tilde{a}_0$ *for every* $o \in \mathcal{O}$ *and* $f \in \mathcal{A}$*, with* $\tilde{a}_0 > 0$*;*
- *(b)* $k_f = \tilde{k}\,(c_f^{NA}(\bar{\mathbf{w}}))^{1-\eta}$ *for every* $f \in \mathcal{A}$*, with* $\tilde{k} \geq 0$ *and* $\tilde{a}_0 + \tilde{k}/(\eta-1) > 0$*, so that (iv) holds;*
- *(c)* $\rho_o = \rho$ *for every* $o \in \mathcal{O}$*.*

Write $\tilde{\alpha}_0 \equiv \epsilon\,\tilde{a}_0$ *for adopters' productivity gain,* $\tilde{\kappa} \equiv \epsilon\,\tilde{k}$ *for AI's first-order share of an adopter's cost, and* $\kappa \equiv \rho\,\tilde{\kappa}$ *for its share of sector cost. Then for every* $o \in \mathcal{O}$*,*

$$\begin{aligned} \tau_o \;=\; (\sigma-1)\,\tilde{\alpha}_0 + \frac{\sigma-\eta}{\eta-1}\,\tilde{\kappa}, \qquad d\ln w_o \;&=\; -\,\frac{\kappa}{1+\zeta\pi_0}, \\ d\ln L_o \;&=\; -\,\frac{\zeta\pi_0}{1+\zeta\pi_0}\,\kappa. \end{aligned} \tag{33}$$

Moreover, writing $\delta \equiv \tilde{\alpha}_0 + \tilde{\kappa}/(\eta-1) > 0$ *for adopters' log cost reduction, the price index falls,* $d\ln P = -\rho\delta - \kappa/(1+\zeta\pi_0) < 0$*, and labor's share of income falls from* $1/\mu$ *to* $(1-\kappa)/\mu$*.*

The proof is in Appendix A3.5.6. The DiD estimand τ_o is positive when the productivity gain $\tilde{\alpha}_0$ is large enough. The AI cost term $(\sigma-\eta)/(\eta-1)\,\tilde{\kappa}$ is positive when tasks are substitutes, $\eta > 1$, and negative when tasks are complements, $\eta < 1$. The general-equilibrium change is never positive: it is zero without AI capital and negative with it, in proportion to AI's share of sector cost, and it does not depend on $\tilde{\alpha}_0$.

The reason is as follows. Adoption lowers prices, so real output rises. But adopters spend part of their costs on AI capital, so the labor share falls, from $1/\mu$ to $(1-\kappa)/\mu$ of income. The value of the outside option rises with real output, so its value relative to employment rises. This reduces the share of workers who are employed, in proportion to the participation margin $\zeta\pi_0$, where ζ is the Fréchet parameter and π_0 is the outside-option share.

A3.4 Simulated equilibria

The results of Appendices A3.1–A3.3 are first order. This subsection solves the full nonlinear equilibrium of Definition 1, with endogenous adoption, and asks which combinations of signs of the matched-DiD estimand and the general-equilibrium change occur. The economy has two occupations and one seniority; occupation 1 is the cheaper one at the no-AI equilibrium. Firms share the same production technology and differ in their demand shifters ξ_f and adoption costs $\mathcal{F}_f$. The elasticities are $\eta = 0.5$, $\sigma = 4$, and $\zeta = 2$, and the outside-option share is $\pi_0 = 0.37$. Adopters' technology is described by $\tilde{\alpha}_o$, the log change in occupation-o labor productivity, so that $\alpha^{(o,f)} = (\eta - 1)\tilde{\alpha}_o$, and by $\tilde{\kappa}$, AI's share of an adopter's cost at the no-AI wages. Appendix A3.6 gives every parameter value and the solution method. For each scenario, τ_o is the matched-DiD estimand of Proposition 1: each adopter's occupation-o employment relative to its own no-adoption counterfactual at the new equilibrium wages and price index, averaged over adopters with baseline employment weights. $d \ln L_o$ is the log change in occupation-o employment between the no-AI equilibrium and the equilibrium with AI.

All four combinations of signs occur for occupation 1. In the first scenario, adopters' productivity gain is in the expensive occupation; with complementary tasks this raises their demand for occupation 1, and the sector employs more occupation-1 workers at a higher wage. In the second, a uniform gain with AI capital raises adopters' employment relative to their twins, while AI spending displaces the sector's wage bill, wages fall, and workers leave for the outside option. In the third, the gain is concentrated in the expensive occupation and AI capital absorbs most of the cost reduction. Adopters employ fewer occupation-1 workers than their twins, because their scale gain is small and AI displaces labor; occupation-1 employment nonetheless rises, because adopters' demand tilts toward occupation 1, its wage falls by less than the other occupation's, and workers move from occupation 2 into occupation 1. In the fourth, a uniform gain with AI capital absorbing most of it lowers both. The third and fourth scenarios are so-so technologies in the sense of Acemoglu and Restrepo (2019): AI performs a large share of adopters' tasks at

Scenario	Adopters' technology $\tilde{\alpha}_1$	$\tilde{\alpha}_2$	$\tilde{\kappa}$	Share adopting	Occupation 1 (cheap) τ_1	$d \ln L_1$	Occupation 2 (expensive) τ_2	$d \ln L_2$
Gain in the expensive occupation only	0.00	0.40	0.00	0.99	+0.831	+0.088	+0.631	−0.071
Uniform gain, with AI capital	0.30	0.30	0.10	0.96	+0.141	−0.046	+0.141	−0.046
Gain mainly in the expensive occupation, AI capital absorbs most of it	0.13	0.68	0.20	0.18	−0.046	+0.004	−0.321	−0.032
Uniform gain, AI capital absorbs most of it	0.50	0.50	0.20	0.79	−0.135	−0.081	−0.135	−0.081

Table A1: Matched-DiD estimand and general-equilibrium change by occupation in simulated equilibria. $\tilde{\alpha}_o$: log labor-productivity gain at adopters in occupation o; $\tilde{\kappa}$: AI's share of an adopter's cost at the no-AI wages; share adopting: share of firms that adopt in the equilibrium with AI. τ_o: adopters' log occupation-o employment relative to their no-adoption counterfactual at the new equilibrium (Proposition 1), baseline-employment weighted. $d \ln L_o$: log change in occupation-o employment between the two equilibria. $\eta = 0.5$, $\sigma = 4$, $\zeta = 2$, $\pi_0 = 0.37$; two occupations, one seniority, 500 firms with a common technology (Appendix A3.6).

a cost close to labor's, so displacement is large and the productivity gain small; in the third, the displacement falls on the expensive occupation. The matched-DiD estimand and the general-equilibrium change are distinct quantities, and the sign of one does not determine the sign of the other.

A3.5 Proofs of Proposition 3 and Example 1

Appendices A3.5.1–A3.5.5 prove Proposition 3; Appendix A3.5.6 proves Example 1. The proof of Proposition 3 has five steps. The maps that define the ϵ-economy are smooth (A3.5.1); their derivatives with respect to wages and the technology are computed (A3.5.2)—derivatives of demand and supply schedules at arbitrary wage changes, not yet changes in equilibrium quantities; the Jacobian of the market-clearing conditions is nonsingular (A3.5.3); the implicit function theorem turns these into the equilibrium path of part (a) (A3.5.4); and part (b) follows by evaluating the derivatives of A3.5.2 in the direction the path moves (A3.5.5).

A3.5.1 The maps are smooth

Fix the adopter set $\mathcal{A}$. This subsection establishes the regularity that the rest of the appendix uses: every map that defines the ϵ-economy at given wages—unit costs, prices, outputs, labor demands, the price index, and aggregate labor demand and supply—is C^∞ in $(\mathbf{w}, \epsilon)$ near $(\bar{\mathbf{w}}, 0)$, with derivatives bounded uniformly across firms, and the adopter's cost reduction δ_f has derivative b_f at $\epsilon = 0$. Lemma 5 differentiates these maps, and the implicit function theorem in A3.5.4 uses their smoothness. The fixed adopter set is what makes smoothness available: were adoption re-optimized along the family, the maps would contain the indicator of each firm's adoption decision.

The maps are as follows. A non-adopter's unit cost is $c_f^{NA}(\mathbf{w}) = [\sum_o (w_o/\Psi_f^{o,0})^{1-\eta}]^{1/(1-\eta)}$. For an adopter, (4) gives $(\Psi_f^{o,A})^{\eta-1} = e^{\epsilon a^{(o,f)}}(\Psi_f^{o,0})^{\eta-1}$, so by Definition 2(ii)–(iii)

$$c_f^A(\mathbf{w};\epsilon) \;=\; B_f(\mathbf{w};\epsilon)^{1/(1-\eta)}, \qquad B_f(\mathbf{w};\epsilon) \;\equiv\; \sum_{o\in\mathcal{O}} e^{\epsilon a^{(o,f)}} \left(w_o/\Psi_f^{o,0}\right)^{1-\eta} + \epsilon\, k_f. \quad (34)$$

At $\epsilon = 0$, $c_f^A(\mathbf{w}; 0) = c_f^{NA}(\mathbf{w})$: the 0-economy is the no-AI economy. Write $c_f = c_f^A$ for $f \in \mathcal{A}$ and $c_f = c_f^{NA}$ otherwise. The remaining objects are $p_f = \mu c_f$; the price index $P(\mathbf{w};\epsilon) = [\int \xi_f(\mu c_f)^{1-\sigma} df]^{1/(1-\sigma)}$; $y_f = \xi_f \mu^{-\sigma} c_f^{-\sigma} P^{\sigma-1}$; $\ell_f^o = (\Psi_f^o)^{\eta-1} w_o^{-\eta} c_f^\eta y_f$; $L_o^D = \int \ell_f^o\, df$; and $L_o^S = T_o w_o^\zeta N/D(\mathbf{w})$. The adopter's log unit-cost reduction at the no-AI wages is

$$\delta_f(\epsilon) \;\equiv\; \ln c_f^{NA}(\bar{\mathbf{w}}) - \ln c_f^A(\bar{\mathbf{w}};\epsilon). \quad (35)$$

Lemma 4 (Smoothness)**.** *Let $\mathcal{W} = [w^{\min}, w^{\max}]^n$ be a compact box containing $\bar{\mathbf{w}}$ in its interior. There is $\epsilon_0 > 0$ such that, on the interior of $\mathcal{W} \times [-\epsilon_0, \epsilon_0]$:*

(i) the maps c_f, p_f, y_f, ℓ_f^o are C^∞ in $(\mathbf{w}, \epsilon)$ with derivatives of every order bounded uniformly in f, and P, L_o^D, L_o^S are C^∞;

(ii) $\delta_f(0) = 0$ and $\delta_f'(0) = b_f$, with b_f as in (30); hence $\delta_f(\epsilon) = \epsilon\, b_f + O(\epsilon^2)$.

Proof. The firm-level maps are compositions of exponentials, powers, and sums, so they are C^∞ at every point where $B_f > 0$; the only failure point is $B_f = 0$, where the power $1/(1-\eta)$ in (34) is not smooth. Choose ϵ_0 so that B_f stays away from zero on the domain: on $\mathcal{W}$, Assumption 1(b) bounds the labor sum of B_f below by a constant $\underline{\beta} > 0$ uniform in f; with $\bar{a}$ and $\bar{k}$ bounding $|a^{(o,f)}|$ and k_f (Definition 2(ii)–(iii)), take ϵ_0 with $e^{-\epsilon_0 \bar{a}} \geq \frac{1}{2}$ and $\epsilon_0 \bar{k} \leq \underline{\beta}/4$. Then $B_f \geq \underline{\beta}/4$ for all $|\epsilon| \leq \epsilon_0$; for $\epsilon < 0$ the AI term ϵk_f is negative, which is why ϵ_0 must be small.

Next, Assumption 1(b) and the bounds of Definition 2(ii)–(iii) further imply that the derivatives of every order are bounded uniformly in f: each derivative depends on f only through the primitives ξ_f, $\Psi_f^{o,0}$, $a^{(o,f)}$, k_f and through B_f, and all of these are uniformly bounded (B_f also from above, by $e^{\epsilon_0 \bar{a}}\bar{\beta} + \epsilon_0 \bar{k}$ with $\bar{\beta}$ the uniform upper bound on the labor sum). This is what smoothness of the integrals requires: by differentiation under the integral sign, $P^{1-\sigma}$ and L_o^D are C^∞; $P^{1-\sigma}$ is bounded away from zero, so P is C^∞. L_o^S is C^∞ directly, since $D(\mathbf{w}) > 0$ on $\mathcal{W}$.

Finally, at $(\bar{\mathbf{w}}, 0)$ the bracket equals $(c_f^{NA}(\bar{\mathbf{w}}))^{1-\eta}$, and differentiating (34),

$$\frac{d}{d\epsilon} \ln c_f^A(\bar{\mathbf{w}};\epsilon)\Big|_{\epsilon=0} \;=\; \frac{1}{1-\eta} \cdot \frac{\sum_{o\in\mathcal{O}} a^{(o,f)} \left(\bar{w}_o/\Psi_f^{o,0}\right)^{1-\eta} + k_f}{\left(c_f^{NA}(\bar{\mathbf{w}})\right)^{1-\eta}} \;=\; -\,b_f,$$

with b_f as in (30). So $\delta_f(0) = 0$ and $\delta_f'(0) = b_f$, and $\delta_f(\epsilon) = \epsilon b_f + O(\epsilon^2)$ since δ_f is C^∞ by (i). □

A3.5.2 First-order responses of demand and supply

Throughout this appendix, d denotes the first-order change at $(\bar{\mathbf{w}}, 0)$: for a log-wage change $d\ln\mathbf{w} \in \mathbb{R}^n$ and the technology change ϵ, and for a quantity $F(\mathbf{w};\epsilon)$,

$$dF \;\equiv\; \sum_{o\in\mathcal{O}} \frac{\partial F}{\partial \ln w_o}\Big|_{(\bar{\mathbf{w}},0)} d\ln w_o \;+\; \frac{\partial F}{\partial \epsilon}\Big|_{(\bar{\mathbf{w}},0)} \epsilon.$$

Here $d\ln\mathbf{w}$ is an arbitrary log-wage change, and the maps are the demand and supply schedules of A3.5.1, defined at any wages: the lemma below is calculus about these schedules and does not assert market clearing ($d\ln L_o^D = d\ln L_o^S$) or tie $d\ln\mathbf{w}$ to equilibrium. Table A2 collects the notation.

Symbol	Definition	Description
δ_f	$\ln c_f^{NA}(\bar{\mathbf{w}}) - \ln c_f^A(\bar{\mathbf{w}};\epsilon)$	adopter f's log cost reduction at the no-AI wages; $\epsilon b_f + O(\epsilon^2)$
$\bar{s}_f^o$	$(\bar{w}_o/\Psi_f^{o,0})^{1-\eta}/(c_f^{NA}(\bar{\mathbf{w}}))^{1-\eta}$	firm f's baseline cost share of occupation o
ω_f	$\sum_{o'} \bar{s}_f^{o'}\, d\ln w_{o'}$	cost-share-weighted average of the wage changes at firm f
$\bar{\lambda}_f^o$	$\bar{\ell}_f^o/\bar{L}_o$	firm f's share density of baseline occupation-o employment
ρ_o	$\int_{\mathcal{A}} \bar{\lambda}_f^o\, df$	adopters' share of baseline occupation-o employment
τ_o	$\rho_o^{-1}\int_{\mathcal{A}} \bar{\lambda}_f^o\big[\alpha^{(o,f)} + (\sigma-\eta)\delta_f\big]\, df$	matched-DiD estimand for occupation o, (31)
ω_o	$\int \bar{\lambda}_f^o\, \omega_f\, df$	average of ω_f, weighted by occupation-o employment
π_o	$\bar{L}_o/N$	baseline share of workers in occupation o
$\bar{v}_f$	$\xi_f\, \bar{p}_f^{1-\sigma}/\bar{P}^{1-\sigma}$	firm f's baseline revenue share density; $\int \bar{v}_f\, df = 1$
δ^P	$\int_{\mathcal{A}} \bar{v}_f\, \delta_f\, df$	revenue-weighted average of adopters' cost reductions
ω^P	$\int \bar{v}_f\, \omega_f\, df$	revenue-weighted average of the ω_f

Table A2: Notation for Appendix A3.5.

The lemma differentiates each equilibrium map in turn.

Lemma 5 (Responses of demand and supply)**.** *For every firm f, every cell o, and every direction $(d\ln\mathbf{w},\epsilon)$,*

$$\begin{aligned} d\ln c_f &= -\,\delta_f\,\mathbf{1}_{f\in\mathcal{A}} + \omega_f, && (36)\\ d\ln P &= -\,\delta^P + \omega^P, && (37)\\ d\ln y_f &= -\sigma\,d\ln c_f + (\sigma-1)\,d\ln P, && (38)\\ d\ln \ell_f^o &= \alpha^{(o,f)}\mathbf{1}_{f\in\mathcal{A}} - \eta\,d\ln w_o + (\eta-\sigma)\,d\ln c_f + (\sigma-1)\,d\ln P, && (39)\\ d\ln L_o^D &= -\eta\,d\ln w_o + \rho_o\tau_o - (\sigma-1)\,\delta^P + (\sigma-1)\,\omega^P + (\eta-\sigma)\,\omega_o, && (40)\\ d\ln L_o^S &= \zeta\,d\ln w_o - d\ln D. && (41) \end{aligned}$$

Proof. By Lemma 4 the maps are differentiable at $(\bar{\mathbf{w}},0)$; for the integrals P and L_o^D, differentiation under the integral sign is justified by the uniform bounds of Lemma 4.

Cost. At $\epsilon=0$, $c_f^A(\cdot\,;0)=c_f^{NA}$, so for adopters and non-adopters alike $\partial\ln c_f/\partial\ln w_{o'}$ at $(\bar{\mathbf{w}},0)$ is the baseline cost share $\bar{s}_f^{o'}$. In ϵ, $\partial_\epsilon \ln c_f^A(\bar{\mathbf{w}};0)=-b_f$ by (35)–(30), and c_f^{NA} does not depend on ϵ. Hence $d\ln c_f = -\epsilon b_f\,\mathbf{1}_{f\in\mathcal{A}} + \sum_{o'}\bar{s}_f^{o'}\,d\ln w_{o'}$, which is (36) since $\delta_f=\epsilon b_f$ to first order.

Price index. Differentiating (26) with $p_f=\mu c_f$ gives $(1-\sigma)\,d\ln P=\int \bar{v}_f\,(1-\sigma)\,d\ln c_f\,df$, so $d\ln P=\int \bar{v}_f\,d\ln c_f\,df=-\delta^P+\omega^P$, which is (37); the superscript P records that these two averages are the pieces of the price-index response.

Output. Equation (38) is (8) in logs, with $E=1$.

Cell demand. From (6), $\ln\ell_f^o=\ln(\Psi_f^o)^{\eta-1}-\eta\ln w_o+\eta\ln c_f+\ln y_f$. By (4) and Definition 2(ii), $\partial_\epsilon\ln(\Psi_f^{o,A})^{\eta-1}=a^{(o,f)}$, so $d\ln(\Psi_f^o)^{\eta-1}=\alpha^{(o,f)}\mathbf{1}_{f\in\mathcal{A}}$; substituting (38) gives (39).

Aggregate demand. $d\ln L_o^D=\int\bar{\lambda}_f^o\,d\ln\ell_f^o\,df$. Substituting (39) with (36)–(37): the shifter and cost-reduction terms integrate to $\int_{\mathcal{A}}\bar{\lambda}_f^o\big[\alpha^{(o,f)}+(\sigma-\eta)\delta_f\big]df=\rho_o\tau_o$ by (31); the ω_f term integrates to $(\eta-\sigma)\,\omega_o$; the $-\eta\,d\ln w_o$ and price-index terms do not vary across firms and pass through, since $\int\bar{\lambda}_f^o\,df=1$. This is (40).

Supply. From (24), $\ln L_o^S=\ln T_o+\zeta\ln w_o+\ln N-\ln D$, and $\partial\ln D/\partial\ln w_{o'}=\zeta T_{o'}\bar{w}_{o'}^{\zeta}/\bar{D}=\zeta\bar{L}_{o'}/N=\zeta\pi_{o'}$, so $d\ln D=\zeta\sum_{o'}\pi_{o'}\,d\ln w_{o'}$. This is (41). □

A3.5.3 The Jacobian is nonsingular

Proposition 3(a) applies the implicit function theorem to the market-clearing system $Z(\mathbf{w};\epsilon)=0$, where $Z_o\equiv L_o^D-L_o^S$. Lemma 4 established that Z is C^∞ near $(\bar{\mathbf{w}},0)$. Lemma 6 establishes the theorem's remaining condition: the Jacobian of Z with respect to log wages is nonsingular there—no first-order wage movement leaves every labor market clearing.

Lemma 6 (Regularity)**.** *The Jacobian of the market-clearing conditions (vii) of Definition 1 with respect to log wages is nonsingular at every no-AI equilibrium.*

Proof. Step 1: reduction to a matrix kernel. Set $\epsilon = 0$ in Lemma 5, so that δ_f, τ_o, and δ^P vanish and only the log-wage change remains; write $x \equiv d\ln \mathbf{w}$. Subtracting (41) from (40),

$$d\ln L_o^D - d\ln L_o^S \;=\; -(\eta+\zeta)\,x_o + (\eta-\sigma)\,(\Omega x)_o + (\sigma-1)\,\bar{S}^\top x + \zeta\,\pi^\top x \;\equiv\; -(Mx)_o,$$

where $\Omega_{oo'} \equiv \int \bar{\lambda}_f^o\,\bar{s}_f^{o'}\,df$ (so that $\omega_o = (\Omega x)_o$), $\bar{S}_{o'} \equiv \int \bar{v}_f\,\bar{s}_f^{o'}\,df$ (so that $\omega^P = \bar{S}^\top x$), $\mathbf{1} \equiv (1,\ldots,1)^\top$ is the vector of ones, and

$$M \;\equiv\; A - \mathbf{1}q^\top, \qquad A \;\equiv\; (\eta+\zeta)\,I + (\sigma-\eta)\,\Omega, \qquad q \;\equiv\; (\sigma-1)\,\bar{S} + \zeta\,\pi.$$

At the no-AI equilibrium $L_o^D = L_o^S = \bar{L}_o$, so the first-order change of the level $L_o^D - L_o^S$ is $\bar{L}_o\,(d\ln L_o^D - d\ln L_o^S) = -\bar{L}_o\,(Mx)_o$: the Jacobian is $-\mathrm{diag}(\bar{L}_o)\,M$. It therefore suffices to show that $Mx = 0$ implies $x = 0$.

Step 2: if $Mx = 0$, then x is uniform. Call a wage change x *uniform* if all its entries are equal—$x = c\,\mathbf{1}$ for a scalar c, so that every wage moves by the same proportion. This step shows that $Mx = 0$ forces x to be uniform; Step 3 then shows that a uniform x with $Mx = 0$ is zero. The price-index and outside-option channels enter M through the rank-one matrix $\mathbf{1}q^\top$, which moves every market by the same amount $q^\top x$; so if $Mx = 0$, then

$$0 \;=\; Mx \;=\; Ax - \mathbf{1}\,(q^\top x) \qquad\Longrightarrow\qquad Ax \;=\; (q^\top x)\,\mathbf{1},$$

and it suffices that A is invertible and maps $\mathbf{1}$ to a multiple of itself.

The eigenvalues of Ω are real and lie in $[0,1]$. Let

$$\Lambda \;\equiv\; \mathrm{diag}(\bar{w}_o\bar{L}_o), \qquad B_{oo'} \;\equiv\; \int \frac{(\bar{w}_o\bar{\ell}_f^o)(\bar{w}_{o'}\bar{\ell}_f^{o'})}{\bar{C}_f}\,df,$$

where $\bar{C}_f = \bar{c}_f\bar{y}_f$ is firm f's baseline total cost. Since $\bar{\lambda}_f^o = \bar{\ell}_f^o/\bar{L}_o$ and $\bar{s}_f^{o'} = \bar{w}_{o'}\bar{\ell}_f^{o'}/\bar{C}_f$, we have $\Lambda\Omega = B$. B is symmetric, and positive semidefinite because $x^\top Bx = \int \big(\sum_o x_o\bar{w}_o\bar{\ell}_f^o\big)^2/\bar{C}_f\,df \geq 0$. Hence $\Lambda^{1/2}\Omega\Lambda^{-1/2} = \Lambda^{-1/2}B\Lambda^{-1/2}$: Ω is *similar* to a symmetric positive semidefinite matrix, and similar matrices (X and TXT^{-1} for invertible T) have the same eigenvalues, so the eigenvalues of Ω are real and nonnegative. They are at most one: Ω has nonnegative entries and unit row sums ($\Omega\mathbf{1} = \mathbf{1}$, since each row of Ω is a weighted average of firms' cost-share vectors, and cost shares sum to one), so for any eigenpair $\Omega v = \lambda v$, picking an entry i with $|v_i|$ largest gives

$$|\lambda|\,|v_i| \;=\; \Big|\sum_{o'}\Omega_{io'}v_{o'}\Big| \;\leq\; |v_i|\sum_{o'}\Omega_{io'} \;=\; |v_i|.$$

If $\Omega v = \lambda v$, then $Av = \big[(\eta+\zeta) + (\sigma-\eta)\lambda\big]v$, and every eigenvalue of A is of this form because A is a polynomial in Ω. With $\lambda \in [0,1]$, these numbers lie between

$\eta+\zeta$ and $\sigma+\zeta$, both positive, so A is invertible. And $A\mathbf{1} = (\sigma+\zeta)\mathbf{1}$, by $\Omega\mathbf{1} = \mathbf{1}$, so $A^{-1}\mathbf{1} = \mathbf{1}/(\sigma+\zeta)$. Applying A^{-1} to $Ax = (q^\top x)\,\mathbf{1}$ gives $x = \frac{q^\top x}{\sigma+\zeta}\,\mathbf{1}$, a scalar multiple of $\mathbf{1}$: x is uniform.

Step 3: if x is uniform and $Mx = 0$, then $x = 0$. With E fixed, a uniform wage rise raises every cost and price in proportion, so labor demand falls one-for-one. The outside option's value is fixed, so labor supply rises by $\zeta\pi_0$. In the algebra: applying $q^\top$ to $x = \frac{q^\top x}{\sigma+\zeta}\,\mathbf{1}$ and using $q^\top\mathbf{1} = (\sigma-1) + \zeta(1-\pi_0)$,

$$q^\top x \;=\; q^\top x \cdot \frac{(\sigma-1)+\zeta(1-\pi_0)}{\sigma+\zeta} \quad\Longrightarrow\quad q^\top x \cdot \frac{1+\zeta\pi_0}{\sigma+\zeta} \;=\; 0,$$

so $q^\top x = 0$ and $x = 0$. □

A3.5.4 Proof of part (a)

Part (a) makes two claims: a continuously differentiable wage path $\mathbf{w}(\epsilon)$ through $\bar{\mathbf{w}}$ exists, and the allocation it induces is an equilibrium of the ϵ-economy. The implicit function theorem delivers the path (Step 1). Every condition of Definition 1 then holds by construction except adoption optimality, which is the substance of Step 2.

Step 1: the wage path. View $Z(\ln\mathbf{w};\epsilon)$, with $Z_o \equiv L_o^D - L_o^S$, as a function of log wages. By Lemma 4 it is C^∞ on an open set containing $(\ln\bar{\mathbf{w}}, 0)$, and $Z(\ln\bar{\mathbf{w}}; 0) = 0$ because $\bar{\mathbf{w}}$ is a no-AI equilibrium and the 0-economy is the no-AI economy. By Lemma 6 its Jacobian in log wages is nonsingular at that point. By the implicit function theorem there exist $\epsilon_1 \in (0, \epsilon_0)$, a neighborhood V of $\ln\bar{\mathbf{w}}$, and a unique C^∞ map $\ln\mathbf{w} : (-\epsilon_1, \epsilon_1) \to V$ with $\ln\mathbf{w}(0) = \ln\bar{\mathbf{w}}$ and $Z(\ln\mathbf{w}(\epsilon);\epsilon) = 0$; write $\mathbf{w}(\epsilon) \equiv \exp(\ln\mathbf{w}(\epsilon))$.

Step 2: the path is an equilibrium for $\epsilon \in [0, \epsilon_1)$. Given $\mathbf{w}(\epsilon)$, define the allocation and the price index by conditions (i)–(iii), (v), and (vi) of Definition 1, as part (a) states; these conditions then hold by construction, and (vii) is $Z(\ln\mathbf{w}(\epsilon);\epsilon) = 0$. Only (iv), adoption optimality, requires an argument. Firms outside $\mathcal{A}$ cannot adopt. A firm $f \in \mathcal{A}$ has zero adoption cost, so adopting is optimal iff $\Delta_f \equiv (c_f^A)^{1-\sigma} - (c_f^{NA})^{1-\sigma} \geq 0$; and since $\sigma > 1$, Δ_f has the sign of

$$g_f(\epsilon) \;\equiv\; \ln c_f^{NA}(\mathbf{w}(\epsilon)) - \ln c_f^A(\mathbf{w}(\epsilon);\epsilon).$$

It therefore suffices that $g_f(\epsilon) > 0$ for every adopter, so that adoption is strictly optimal. Continuity is not enough: $g_f(0) = 0$, and the bound must hold for all adopters on a common interval. We use a first-order bound with uniform control of the curvature. By the chain rule,

$$g_f'(\epsilon) \;=\; -\,\partial_\epsilon \ln c_f^A\big(\mathbf{w}(\epsilon);\epsilon\big) \;+\; \sum_{o\in\mathcal{O}} \big[s_f^{o,NA} - s_f^{o,A}\big]\big(\mathbf{w}(\epsilon);\epsilon\big)\;\frac{d\ln w_o(\epsilon)}{d\epsilon},$$

where $s_f^{o,d} = \partial \ln c_f^d / \partial \ln w_o$ are the cost shares: wage movements along the path affect both costs through their cost shares, and the technology affects c_f^A alone. At $\epsilon = 0$ the two cost functions coincide, so the difference in cost shares vanishes and only the technology term survives: $g_f'(0) = \delta_f'(0) = b_f \geq \underline{b}$, by Lemma 4(ii) and Definition 2(iv). Shrink ϵ_1 so that $\ln \mathbf{w}(\cdot)$ and its first two derivatives are bounded on $[0, \epsilon_1]$; by the uniform bounds of Lemma 4, g_f is then C^2 with $|g_f''| \leq C$ on $[0, \epsilon_1]$ for a constant C independent of f. Hence $g_f(\epsilon) \geq \underline{b}\,\epsilon - C\epsilon^2/2 > 0$ for $0 < \epsilon < \min\{\epsilon_1,\, 2\underline{b}/C\}$. Shrinking ϵ_1 once more, every firm in $\mathcal{A}$ strictly prefers to adopt, which is (iv). Finally, every equilibrium quantity is a C^∞ function of $(\mathbf{w}, \epsilon)$ composed with the C^∞ path $\mathbf{w}(\epsilon)$, hence continuously differentiable in ϵ. □

A3.5.5 Proof of part (b)

Part (b) claims the decomposition (32). The proof differentiates market clearing along the path: this yields a linear system that the wage response solves (Step 1), and substituting the wage response into labor supply yields (32) (Step 2).

Step 1: the wage response solves a linear system. Part (a) provides a continuously differentiable family of equilibria, so market clearing holds identically in ϵ and can be differentiated. By the chain rule, the first-order change of any quantity $F(\mathbf{w}(\epsilon); \epsilon)$ along the path is dF of Appendix A3.5.2, evaluated in the direction the path moves: $d \ln \mathbf{w}$ is now the first-order change in equilibrium wages, not an arbitrary wage change. Condition (vii) of Definition 1 holds along the path, so $d \ln L_o^D = d \ln L_o^S$ for every o; substituting (40)–(41), and using (37) to write $-(\sigma - 1)\,\delta^P + (\sigma - 1)\,\omega^P = (\sigma - 1)\, d \ln P$, the wage changes satisfy

$$(\eta + \zeta)\, d \ln w_o \;=\; \rho_o\, \tau_o + (\sigma - 1)\, d \ln P + (\eta - \sigma)\, \omega_o + d \ln D, \qquad o \in \mathcal{O}. \tag{42}$$

Substituting $d \ln P = -\delta^P + \bar{S}^\top d \ln \mathbf{w}$, $\omega_o = (\Omega\, d \ln \mathbf{w})_o$, and $d \ln D = \zeta\, \pi^\top d \ln \mathbf{w}$, and collecting every $d \ln \mathbf{w}$ term on the left, (42) reads

$$(M\, d \ln \mathbf{w})_o \;=\; \rho_o\, \tau_o - (\sigma - 1)\, \delta^P,$$

with M exactly as in Lemma 6: a linear system in $d \ln \mathbf{w}$ with a nonsingular coefficient matrix, which therefore determines $d \ln \mathbf{w}$ uniquely.

Step 2: substitute into labor supply. Markets clear along the path, so $d \ln L_o = d \ln L_o^S = \zeta\, d \ln w_o - d \ln D$ by (41). Solving (42) for $d \ln w_o$ and substituting,

$$d \ln L_o \;=\; \frac{\zeta}{\eta + \zeta} \Big[\rho_o \tau_o - (\sigma - \eta)\, \omega_o + (\sigma - 1)\, d \ln P + d \ln D \Big] \;-\; d \ln D,$$

and $\zeta/(\eta + \zeta) - 1 = -\eta/(\eta + \zeta)$ on the $d \ln D$ terms, which is (32) with Γ as defined there. □

A3.5.6 Proof of Example 1

The example claims the three formulas of (33) and the statements about the price index and the labor share. Step 1 computes the estimand directly; Steps 2–3 solve the wage system (42) by guess-and-verify; Step 4 substitutes into labor supply; Step 5 proves the price-index and labor-share statements.

Step 1: the estimand. By hypotheses (a)–(b) and (30), using that cost shares sum to one, $b_f = \tilde{a}_0 + \tilde{k}/(\eta-1)$ for every adopter, so to first order $\delta_f = \epsilon b_f = \tilde{\alpha}_0 + \tilde{\kappa}/(\eta-1) \equiv \delta$, the same for every adopter. The integrand of (31) is then constant across adopters, so the weights integrate to one and

$$\tau_o \;=\; (\eta-1)\,\tilde{\alpha}_0 + (\sigma-\eta)\,\delta \;=\; (\sigma-1)\,\tilde{\alpha}_0 + \frac{\sigma-\eta}{\eta-1}\,\tilde{\kappa},$$

the first formula of (33).

Step 2: adopters' revenue share equals ρ. The price index weights firms by revenue, while hypothesis (c) fixes adopters' share of employment; this step shows the two coincide at the baseline. At the no-AI equilibrium a firm's revenue is μ times its wage bill: $\bar{v}_f = \bar{p}_f\bar{y}_f = \mu\,\bar{c}_f\bar{y}_f$, and by the factor demands (6), $\sum_o \bar{w}_o\bar{\ell}^o_f = \bar{c}_f^{\eta}\bar{y}_f\sum_o(\bar{w}_o/\Psi^{o,0}_f)^{1-\eta} = \bar{c}_f\bar{y}_f$, since no firm uses AI capital at the baseline. Integrating over adopters, using hypothesis (c) and then $\mu\sum_o \bar{w}_o\bar{L}_o = \int \bar{v}_f\,df$,

$$\int_{\mathcal{A}} \bar{v}_f\,df \;=\; \mu\sum_o \bar{w}_o\int_{\mathcal{A}}\bar{\ell}^o_f\,df \;=\; \mu\sum_o \bar{w}_o\,\rho\,\bar{L}_o \;=\; \rho\int \bar{v}_f\,df \;=\; \rho.$$

Step 3: the wage change is uniform, $\hat{w} = -\kappa/(1+\zeta\pi_0)$. Guess a common wage change, $d\ln w_o = \hat{w}$ for every o, and verify it in the wage system (42); the coefficient matrix is nonsingular, so the verified guess is the solution. With a common wage change, $\omega_f = \hat{w}$ for every firm, because cost shares sum to one; hence $\omega_o = \hat{w}$; $d\ln P = \int \bar{v}_f(-\delta\,\mathbf{1}_{f\in\mathcal{A}} + \hat{w})\,df = -\rho\delta + \hat{w}$ by Step 2; and $d\ln D = \zeta\sum_o \pi_o\hat{w} = \zeta(1-\pi_0)\hat{w}$. Substituting into (42),

$$(\eta+\zeta)\,\hat{w} \;=\; \rho\tau_o - (\sigma-1)\,\rho\,\delta + (\sigma-1)\hat{w} + (\eta-\sigma)\hat{w} + \zeta(1-\pi_0)\hat{w}.$$

By Step 1, $\rho\tau_o - (\sigma-1)\rho\delta = \rho(\eta-1)(\tilde{\alpha}_0-\delta) = -\rho\tilde{\kappa} = -\kappa$. The $\hat{w}$ terms on the right sum to $(\eta-1)\hat{w} + \zeta(1-\pi_0)\hat{w}$; moving them left leaves $(1+\zeta\pi_0)\,\hat{w} = -\kappa$, so $\hat{w} = -\kappa/(1+\zeta\pi_0)$, the second formula of (33).

Step 4: the employment change. By (41),

$$d\ln L_o \;=\; \zeta\hat{w} - d\ln D \;=\; \zeta\hat{w} - \zeta(1-\pi_0)\hat{w} \;=\; \zeta\pi_0\,\hat{w} \;=\; -\,\frac{\zeta\pi_0}{1+\zeta\pi_0}\,\kappa,$$

the third formula of (33).

Step 5: the price index and the labor share. By Step 3, $d\ln P = -\rho\delta + \hat{w} = -\rho\delta - \kappa/(1+\zeta\pi_0)$, which is negative because $\delta > 0$ by condition (iv). The wage bill is $\sum_o w_o L_o$, and by Steps 3–4 its log change is $\hat{w} + \zeta\pi_0\hat{w} = (1+\zeta\pi_0)\hat{w} = -\kappa$, the same in every occupation. At the no-AI equilibrium the wage bill equals $E/\mu = 1/\mu$ by Step 2, so to first order it equals $(1-\kappa)/\mu$ with AI; since $E = 1$ is total income by (23), this is labor's share of income. □

A3.6 Numerical simulation

This appendix describes the simulated economy of Appendix A3.4.

A3.6.1 Setup

The economy has two occupations and one seniority. $F = 500$ firms approximate the continuum of Section 2, with $\sum_f \xi_f = 1$ in place of $\int \xi_f\, df = 1$. Population $N = 1$ and expenditure $E = 1$ are normalizations. The primitives are:

- Elasticities: $\eta = 0.5$, so that tasks are complements; $\sigma = 4$; $\zeta = 2$.
- Demand shifters $\xi_f \sim \text{lognormal}(0, 0.7)$, normalized to sum to one. This is the only source of heterogeneity in firm size.
- Production technology, common to all firms: $\Psi_f^{1,0} = e^{0.5}$ and $\Psi_f^{2,0} = 1$. With a common technology every firm has the same cost shares, so adopters share the same cost reduction δ and adoption selects on fixed costs only.
- Fréchet scales $T_1 = 1.5$ and $T_2 = 0.5$; outside option $w_0 = 1$ with $T_0 = (0.4/0.6)(T_1 + T_2)$, which gives an outside-option share of $\pi_0 = 0.368$ at the no-AI equilibrium.
- Adoption costs $\mathcal{F}_f = \phi_f\, \bar{\pi}_f$, where $\bar{\pi}_f$ is the firm's variable profit at the no-AI equilibrium and ϕ_f is lognormal with median 0.05 and log standard deviation 1 (interquartile range 0.025–0.098).
- Adopters' technology, which defines a scenario: $\Psi_f^{o,A} = e^{\tilde{\alpha}_o}\, \Psi_f^{o,0}$, so that $\alpha^{(o,f)} = (\eta - 1)\tilde{\alpha}_o$; and an AI cost term $(r/\Psi_f^{k,A})^{1-\eta}$ in the unit cost (7) set to $\frac{\tilde{\kappa}}{1-\tilde{\kappa}} \sum_o (\bar{w}_o/\Psi_f^{o,A})^{1-\eta}$, so that AI's share of the adopter's cost at the no-AI wages $\bar{\mathbf{w}}$ is $\tilde{\kappa}$. Only the ratio $r/\Psi_f^{k,A}$ matters, and $r = 1$.

At the no-AI equilibrium, wages are $\bar{\mathbf{w}} = (0.877, 1.504)$, so occupation 1 is the cheaper one; the occupations' cost shares are $(\bar{S}_1, \bar{S}_2) = (0.373, 0.627)$ and their shares of market employment are $(h_1, h_2) = (0.505, 0.495)$.

A3.6.2 Solution

For given wages $\mathbf{w}$, unit costs $c_f^{NA}(\mathbf{w})$ and $c_f^A(\mathbf{w})$ follow from (7). The price index and the adopter set are the fixed point of Lemma 3: for a candidate P, firm f adopts if $(\mu^{1-\sigma}/\sigma)\,\xi_f P^{\sigma-1}E\,[(c_f^A)^{1-\sigma}-(c_f^{NA})^{1-\sigma}] \geq \mathcal{F}_f$, and $G(P)$ is the index of the resulting prices; $P-G(P)$ is increasing, and bisection finds its zero. Output, labor demands, and labor supply then follow from conditions (i), (ii), and (v) of Definition 1, and the two market-clearing conditions (vii) are solved for $(\ln w_1, \ln w_2)$ by a Newton-type root finder. The no-AI equilibrium is solved with adoption unavailable. With a finite number of firms, excess demand jumps when a marginal firm's adoption decision flips, so market clearing holds up to the size of one firm's employment; the largest log residual across the scenarios of Table A1 is 3×10^{-4}, an order of magnitude below the smallest employment change reported.

A3.6.3 Outcomes

For each adopter f, the matched-DiD estimand is $\ln \ell_f^o(\mathbf{w}', P') - \ln \ell_f^{o,NA}(\mathbf{w}', P')$, where $(\mathbf{w}', P')$ are the wages and price index of the equilibrium with AI and $\ell_f^{o,NA}$ is the firm's labor demand had it not adopted, evaluated at the same wages and price index. This is the comparison of Proposition 1. τ_o averages it over adopters with weights equal to their baseline occupation-o employment; the firm-total estimand τ uses baseline total employment. $d\ln L_o$ and $d\ln L$ are log changes in occupation and total market employment between the two equilibria. Table A3 reports all outcomes. As a check on the solver, a uniform productivity gain without AI capital ($\tilde{\alpha}_1 = \tilde{\alpha}_2 = 0.30$, $\tilde{\kappa} = 0$) reproduces Example 1: $\tau_o = 0.900$ in both occupations and $d\ln L_o = 0$ to numerical precision.

A4 Robustness to Nested CES Production

The propositions in the main text are stated under the single-level CES production function (3). Up to monotone transformations, CES is essentially the unique homothetic, twice-differentiable production function whose pairwise elasticities of substitution are constant (Uzawa's theorem and its extensions), so strict deviations from CES generally yield non-constant cross-elasticities and break the clean Proposition 1 decomposition. One generalization preserves the structure: *nested CES*, where cells are partitioned into nests with potentially different within-nest and across-nest elasticities. This appendix derives the counterpart of Proposition 1.

A4.1 Setup

Partition the labor types $\mathcal{T}$ into disjoint nests $\{\mathcal{T}_g\}_{g\in\mathcal{G}}$; in the application each seniority is a nest. *AI capital occupies its own singleton nest* $\{k\}$ *at the outer level*:

Scenario	Share adopting	κ	$d\ln w_1$	$d\ln w_2$	$d\ln P$	Occupation 1 τ_1	Occupation 1 $d\ln L_1$	Occupation 2 τ_2	Occupation 2 $d\ln L_2$	τ; $d\ln L$
Gain in the expensive occupation only	0.99	0.000	+0.055	−0.025	−0.231	+0.831	+0.0883	+0.631	−0.0713	+0.735; +0.0125
Uniform gain, with AI capital	0.96	0.100	−0.059	−0.059	−0.140	+0.141	−0.0455	+0.141	−0.0455	+0.141; −0.0455
Gain mainly in the expensive occupation, AI capital absorbs most of it	0.18	0.038	−0.010	−0.028	−0.022	−0.046	+0.0038	−0.321	−0.0316	−0.173; −0.0136
Uniform gain, AI capital absorbs most of it	0.79	0.167	−0.102	−0.102	−0.129	−0.135	−0.0805	−0.135	−0.0805	−0.135; −0.0805

Table A3: Full outcomes for the scenarios of Table A1. κ: AI's share of sector cost in the equilibrium with AI. $d\ln w_o$ and $d\ln P$: log changes in wages and the price index. Totals: firm-total matched-DiD estimand τ and log change in total market employment $d\ln L$.

AI capital competes with seniority *bundles*, not with individual cells inside a bundle. Let $g(t)$ denote the nest of cell t. The two-level production function is:

$$Z_f^g \;=\; \left[\sum_{t\in\mathcal{T}_g}\big(\Psi_f^t\,x_f^t\big)^{(\eta_g-1)/\eta_g}\right]^{\eta_g/(\eta_g-1)} \qquad (\eta_g>0,\;\eta_g\neq 1), \tag{43}$$

$$y_f \;=\; \left[\sum_{g\in\mathcal{G}}\big(Z_f^g\big)^{(\eta_*-1)/\eta_*} + \big(\Psi_f^k\,k_f\big)^{(\eta_*-1)/\eta_*}\right]^{\eta_*/(\eta_*-1)} \qquad (\eta_*>0,\;\eta_*\neq 1), \tag{44}$$

with the AI term present only at adopters. The single-level function (3) is the special case $\eta_g=\eta_*$ for all g.

Lemma 7 (Nested CES factor demands)**.** *Cost minimization at factor prices* $\{w_t\}, r$ *under* (43)–(44) *yields cell-level labor demand*

$$\ell_f^t \;=\; \big(\Psi_f^t\big)^{\eta_{g(t)}-1}\,w_t^{-\eta_{g(t)}}\,\big(c_f^{g(t)}\big)^{\eta_{g(t)}-\eta_*}\,\big(c_f\big)^{\eta_*}\,y_f, \tag{45}$$

where

$$c_f^g \;\equiv\; \Big[\sum_{t\in\mathcal{T}_g}(w_t/\Psi_f^t)^{1-\eta_g}\Big]^{1/(1-\eta_g)} \qquad \textit{and} \qquad c_f \;\equiv\; \Big[\sum_g (c_f^g)^{1-\eta_*}+(r/\Psi_f^k)^{1-\eta_*}\Big]^{1/(1-\eta_*)}$$

are the nest-g and firm-level unit costs (AI term at adopters only). With CES product-market demand of elasticity $\sigma > 1$, $y_f \propto c_f^{-\sigma}$ as in the single-level case.

Proof. Solve the two-level cost minimization sequentially. The inner problem (produce Z^g at minimum cost) is a standard CES and gives, exactly as in (6)–(7) with (Z^g, η_g, c^g) in place of (y, η, c): $\ell^t = (\Psi_f^t)^{\eta_g - 1} w_t^{-\eta_g} (c^g)^{\eta_g} Z^g$ for $t \in \mathcal{T}_g$. The outer problem treats each nest composite as a factor with price c^g (by the inner stage, one unit of Z^g costs c^g) and gives $Z^g = (c^g)^{-\eta_*} c^{\eta_*} y$. Substituting the second into the first yields (45). $\square$

Relative to the single-level demand (6), the extra term $(c_f^{g(t)})^{\eta_{g(t)} - \eta_*}$ carries nest-level cost information; it vanishes when $\eta_{g(t)} = \eta_*$.

A4.2 Adopter–twin comparison under nesting

Proposition 4 (Proposition 1 under nested CES)**.** *Let f be an adopter and f' a non-adopting twin facing the same factor prices, the same price index P, and the same expenditure E. For every cell t in nest g,*

$$\ln \ell_f^{t,A} - \ln \ell_{f'}^{t,NA} \;=\; \frac{\eta_g - 1}{\eta - 1}\, \alpha^{(t,f)} \;+\; (\eta_g - \eta_*)\, \ln \frac{c_f^{g,A}}{c_{f'}^{g,NA}} \;+\; (\eta_* - \sigma)\, \ln \frac{c_f^{A}}{c_{f'}^{NA}}, \quad (46)$$

where $c_f^{g,A}$ and $c_{f'}^{g,NA}$ are the nest-g unit costs of Lemma 7 at the two firms, and c_f^A, $c_{f'}^{NA}$ their firm-level unit costs.

Proof. By Lemma 7,

$$\ell_f^{t,A} = (\Psi_f^{t,A})^{\eta_g - 1}\, w_t^{-\eta_g}\, (c_f^{g,A})^{\eta_g - \eta_*}\, (c_f^{A})^{\eta_*}\, y_f,$$
$$\ell_{f'}^{t,NA} = (\Psi_{f'}^{t,0})^{\eta_g - 1}\, w_t^{-\eta_g}\, (c_{f'}^{g,NA})^{\eta_g - \eta_*}\, (c_{f'}^{NA})^{\eta_*}\, y_{f'}.$$

The wage terms are common and cancel in the ratio. Twins share $\Psi^{t,0}$ and ξ, so $(\Psi_f^{t,A}/\Psi_{f'}^{t,0})^{\eta_g - 1} = \exp\!\big(\frac{\eta_g - 1}{\eta - 1}\alpha^{(t,f)}\big)$ by (4), and $y_f/y_{f'} = (c_f^A/c_{f'}^{NA})^{-\sigma}$ by the last statement of Lemma 7 with common P and E. Taking logs of the ratio gives (46). $\square$

Relative to (12), the nest-level cost term is new. Its coefficient $\eta_g - \eta_*$ nets two responses to a cheaper nest bundle. Through the within-nest elasticity, demand for each cell per unit of the bundle falls by η_g percent, because the cell has become relatively more expensive than the bundle. Through the across-nest elasticity, the firm's use of the cheaper bundle rises by η_* percent. A cheaper nest therefore raises demand for its cells if and only if $\eta_* > \eta_g$. The term vanishes when $\eta_g = \eta_*$, and (46) is (12) when $\eta_g = \eta_* = \eta$ for all g. Contrasting two cells in the same nest cancels the second and third terms, so within-nest contrasts identify shifter differences up to the factor $(\eta_g - 1)/(\eta - 1)$; contrasting cells in different nests leaves the difference of their nest-level terms. With seniority as the nest, the within-seniority quartile contrasts of Section 2.4 are unaffected, and junior-versus-senior contrasts carry the nest-level terms.

A5 Computing the Scale and Bias Terms

This appendix gives the computation behind Figure 7 and Table 4. Let $L_{f,s}$ be employment in seniority group s at affiliate f and $\omega_{f,c,s} = L_{f,c,s}/L_{f,s}$ the share of that group in exposure quartile c, so that $\ln L_{f,c,s} = \ln L_{f,s} + \ln \omega_{f,c,s}$. Write $\beta_q^{L_s}$ and $\beta_q^{\omega_{c,s}}$ for the coefficients of (15) when the outcome is log group employment and the cell share, respectively, and $\bar{\omega}_{c,s}$ for the cell's baseline share in October 2022. Expanding $\ln \omega_{f,c,s}$ to first order around $\bar{\omega}_{c,s}$, the effect of adoption on log cell employment is

$$\beta_q^{L_{c,s}} \;\approx\; \beta_q^{L_s} \;+\; \frac{\beta_q^{\omega_{c,s}}}{\bar{\omega}_{c,s}}. \tag{47}$$

The expansion is around the pooled baseline share rather than each affiliate's own, which keeps affiliates with no baseline employment in a cell in the calculation, at the cost of an approximation when baseline shares differ across affiliates. A quartile cell aggregates occupations, so define its shifter as $\alpha_{Q,s} \equiv \ln \sum_{o \in Q} \theta_{o,s}\, e^{\alpha(o,s,f)}$, where $\theta_{o,s}$ is occupation o's baseline share of the cell's employment; summing (11) over the cell's occupations gives $\ln L^A_{Q,s} - \ln L^{NA}_{Q,s} = \alpha_{Q,s} + (\sigma - \eta)\ln(c^{NA}_{f'}/c^A_f)$, so (13) holds for cells, and $\alpha_{Q1,s} = 0$ holds when every occupation in the bottom quartile has a zero shifter. Differencing the top and bottom exposure quartiles within group s removes the scale term, so the bias of AI toward the most-exposed cell is

$$\alpha_{Q4,s} - \alpha_{Q1,s} \;\approx\; \frac{\beta_q^{\omega_{Q4,s}}}{\bar{\omega}_{Q4,s}} \;-\; \frac{\beta_q^{\omega_{Q1,s}}}{\bar{\omega}_{Q1,s}}, \tag{48}$$

and, when AI leaves the task content of the least-exposed occupations unchanged ($\alpha_{Q1,s} = 0$), the scale term of (12) is the implied effect on log employment in the bottom quartile,

$$(\sigma - \eta)\ln\!\left(c^{NA}_{f'}/c^A_f\right) \;\approx\; \beta_q^{L_s} \;+\; \frac{\beta_q^{\omega_{Q1,s}}}{\bar{\omega}_{Q1,s}}. \tag{49}$$

We evaluate (48) and (49) in every event month for Figure 7 and at the March 2026 endpoint ($q = 40$) for Table 4, for juniors, seniors, and all workers. The baseline shares $\bar{\omega}_{c,s}$ are October 2022 employment-weighted shares among matched treated affiliates: the bottom quartile holds 21.3 percent of juniors and 6.8 percent of seniors, and the top quartile 35.2 and 31.6 percent. Because the bottom-quartile shares are small, especially among seniors, a share change of a fraction of a percentage point in that cell translates into a percentage change of several points, which is why the top-minus-bottom share gap of Figure 4 cannot be used on its own: the two share changes enter (48) with different weights. Figure A1 therefore reports the two share effects separately for each group. Standard errors come from a regression that, for each event month, stacks the log-employment and share long differences with outcome-specific treatment indicators, instruments, and match-by-outcome fixed effects, so that the covariance between the coefficients is estimated. The combinations in (48) and (49) are then delta-method standard errors with the baseline shares treated as known, clustered by the metropolitan area of the parent's headquarters. Under Proposition 1 the scale term is common to every type within a firm; the last column of Table 4 reports the difference between the senior and junior scale terms, whose standard error comes from the same stacked regression.

Figure A1: Event study plots for the effect of AI adoption on the bottom- and top-quartile exposure shares within seniority groups

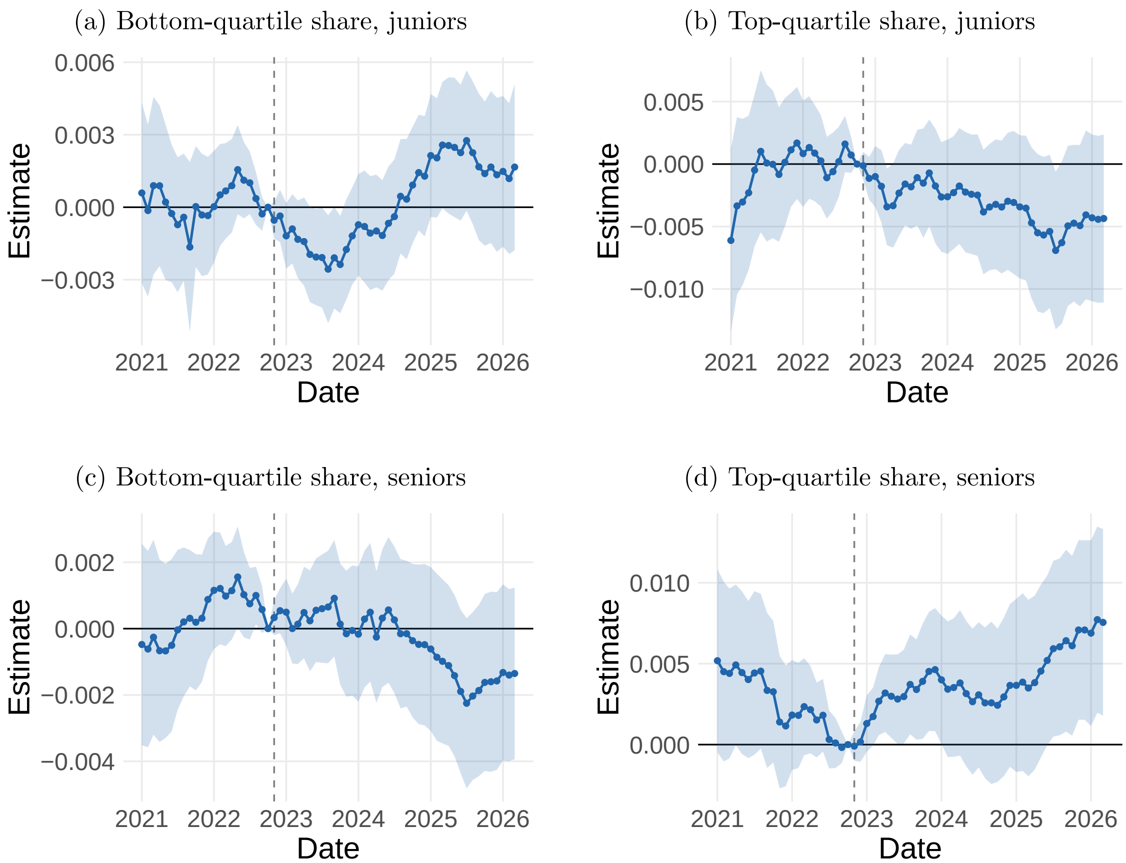

Notes: The figure plots the share inputs to (48) and (49). Each point is a two-stage least squares coefficient β_q from Equation (15) in which the outcome is the share of the seniority group's employment in occupations in the bottom quartile (Panels A and C) or the top quartile (Panels B and D) of AI exposure, entered as the long difference between the month on the horizontal axis and October 2022. Panels A and B are shares of junior employment and Panels C and D shares of senior employment, where junior covers Revelio seniority levels 1–2 and senior levels 3–7. The difference between the top- and bottom-quartile effects within each group is the outcome plotted in Figure 4. Sample, instrument, matching, and weights as in Figure 3. Shaded bands are 95% confidence intervals, clustered by the metropolitan area of the parent's headquarters. The dashed vertical line marks ChatGPT's launch in November 2022.

A6 Additional Details on the AI Adoption Metric

This appendix documents the two stages of the adoption measure of Section 3.3.

Stage 1: keyword screen. We query the text of every job advertisement in the Revelio listings universe for the 41 sample countries, from January 2021 through March 2026, with a case-insensitive regular expression built from the 32 terms in Table A4. Each term is matched as a whole word, that is, with a space on either side, so that "LLM" does not match "LLMs" inside a longer token and "GPT" does not match "GPTW". An advertisement is a candidate if its description matches at least one term. The screen returns 2,092,460 candidates across the countries in the raw data.

Stage 2: LLM verification. Every candidate is passed to Llama 3.1 8B Instruct (Grattafiori et al., 2024) with the prompt in Table A5, at temperature zero and with the description truncated to its first 4,000 characters. The model answers "yes" or "no". The prompt asks whether the advertisement indicates that the job involves using generative AI tools, counting optional or preferred mentions ("experience with ChatGPT is a plus") as "yes", and instructs the model to answer "no" when the technology is mentioned only to prohibit it, when a keyword is a false positive (a law degree, a person's name, Capgemini, a red-amber-green status, an airline copilot), when only conventional machine learning is mentioned, or when the advertisement discusses only whether such tools may be used in applications. Of the 2,092,460

Table A4: Keyword list for the first-stage screen

Group	**Terms**
Model and product names	ChatGPT; GPT; Claude; Gemini; Copilot; Llama; Mistral; Stable diffusion; Midjourney; Dall-e
Core vocabulary	large language model; LLM; generative AI; Gen AI; prompt engineering; prompt design; transformer-based model
Developer tooling and frameworks	LangChain; LlamaIndex; HuggingFace Transformers; RetrievalQA; RAG; retrieval-augmented generation; vector embeddings; vector database
APIs and infrastructure	OpenAI API; Anthropic Claude API; Azure OpenAI; Google Vertex AI Generative; Pinecone; Weaviate; Milvus

Notes: The 32 terms of the first-stage screen, grouped as in Section 3.3. Matching is case-insensitive and requires the term to appear as a whole word. An advertisement is a candidate for the second stage if its description contains at least one term.

Table A5: Prompt for the second-stage LLM verification

```
We are measuring firm-level adoption of generative AI from job descriptions.
As a first pass, job descriptions were selected from Revelio Labs if they
contained at least one GenAI-related keyword (e.g., "ChatGPT", "LLM", "generative
AI", "LangChain"). This prompt serves as a second pass to remove false
positives. Job descriptions may be truncated.

Does this job description indicate the need to use generative AI?

Answer "yes" if the job involves using generative AI tools, including:

Tools/APIs: ChatGPT, GPT-3, GPT-4, GPT-5, OpenAI, Azure OpenAI, Copilot, Claude,
Gemini, Anthropic
Frameworks: LangChain, LlamaIndex, HuggingFace Transformers, RetrievalQA
Concepts: large language model, LLM, generative AI, Gen AI, GenAI, prompt
engineering, prompt design
RAG: retrieval-augmented generation, RAG (in AI context), vector embeddings,
vector database
Infrastructure: Pinecone, Weaviate, Milvus, Google Vertex AI
Note: "Experience with [GenAI tool] is a plus" or similar optional/preferred
mentions still count as "yes", as long as the tools are generative AI tools.

Answer "no" if:

GenAI is mentioned negatively ("don't use", "prohibited", "not allowed")
Keywords are false positives (e.g., LLM = law degree, Claude = person's name,
Gemini = Capgemini, RAG = Red-Amber-Green status, Copilot = airplane pilot).
These are examples; flag any keyword that clearly refers to something other than
generative AI.
Only traditional ML/AI is mentioned (machine learning, neural networks, deep
learning) without GenAI
The job only discusses whether generative AI tools like ChatGPT can or cannot be
used in applications, without indicating that the job itself involves using these
tools

Answer only "yes" or "no".

Job: {description}
```

Notes: The prompt is reproduced verbatim from the verification script; {description} is replaced by the first 4,000 characters of the advertisement. The model is Meta Llama 3.1 8B Instruct, called through an inference API with temperature zero and a ten-token response limit. The prompt's own keyword list omits some screen terms (Llama, Mistral, Stable diffusion, Midjourney, Dall-e, transformer-based model), so advertisements flagged only on those terms are judged under the general instruction rather than an explicit rule.

candidates, 1,365,826 (65 percent) are verified, 706,061 are rejected, and 20,573 (1 percent) returned an API error and are treated as rejected. Verified advertisements in the 41 sample countries number 1,362,406, of which 99 percent are posted on or after November 1, 2022.

Figure A2: AI share of job postings over time

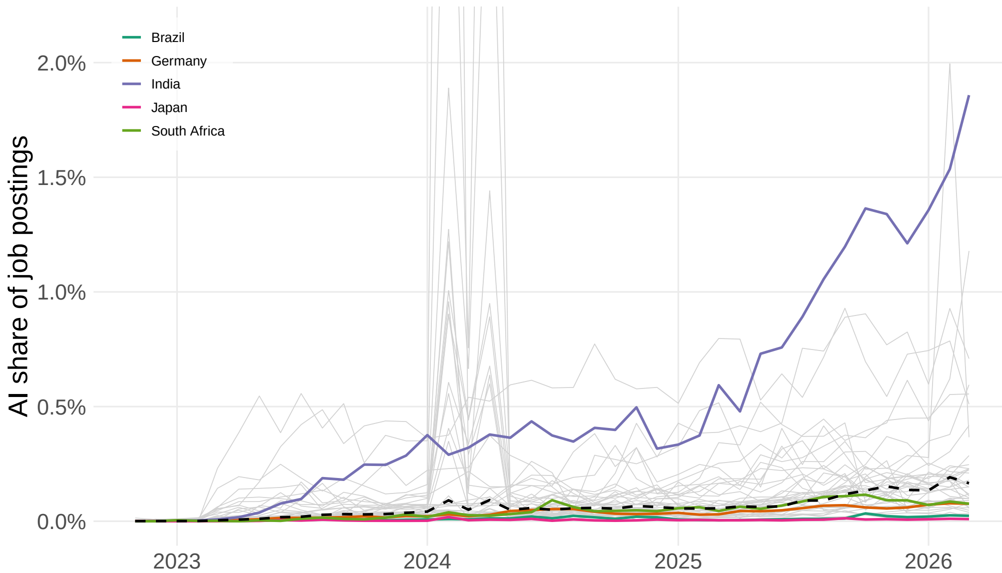

Notes: The figure tracks the share of job advertisements that involve generative AI over time. For each month from November 2022 through March 2026, the series report LLM-verified AI job advertisements as a percentage of all job advertisements first posted that month. Colored lines show India, Brazil, Japan, Germany, and South Africa; gray lines show the remaining countries. The dashed black line pools all countries except the United States, whose posting data end in October 2025 and would otherwise introduce a composition break in the pooled series. Gray lines are truncated at the top of the axis: in a country with few postings in a month, a single hiring burst can push that month's share far above the range shown. For the posting universe and the country sample, see Figure 1.

Adoption over time. Figure A2 plots the monthly AI share of postings by country, and Figure A3 the share of employment at treated affiliates by sector, month by month, in the descriptive sample of Section 3. Treatment accumulates fastest in the information sector, where companies employing more than two thirds of the sector's positions had posted an AI job by March 2026, and slowest in government and other services.

Figure A3: Share of employment at treated affiliates, by sector

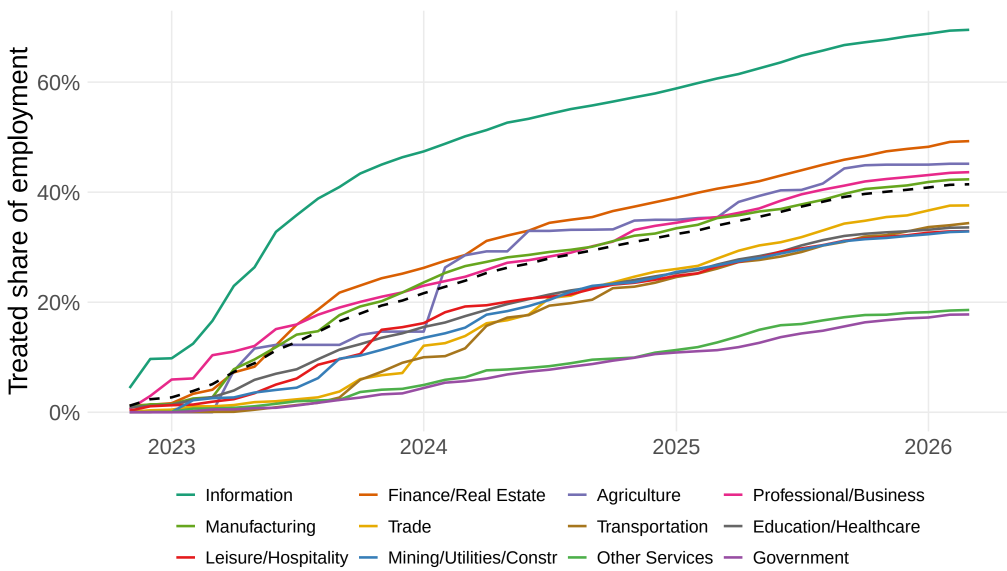

Notes: The figure plots, for each month from November 2022 through March 2026, the share of October 2022 positions held by affiliates that are treated by that month, that is, whose company had posted at least one LLM-verified AI job advertisement anywhere in the world on or before the end of the month. The sample is the affiliates satisfying the sample restrictions of Section 3 across the 41 analysis countries, pooled; each affiliate is weighted by its positions in October 2022. Sectors aggregate Revelio's industry classification into the twelve groups used throughout the paper, and the dashed black line pools all sectors. United States AI-posting data end in October 2025, so treatment there stops accumulating after that month.

A7 AI Adoption Descriptives

This appendix reports posting counts and AI shares by country (Table A6), affiliate counts, baseline characteristics, and treatment shares by country (Table A7), and the AI job postings of adopting companies (Table A8).

Table A6: AI share of job postings by country

Country	Job postings	AI job postings	AI share (%)
India	40,408,907	236,822	0.59
Israel	1,352,853	7,119	0.53
Slovakia	274,484	1,180	0.43
Luxembourg	527,278	1,909	0.36
Lithuania	250,496	801	0.32
Poland	9,069,188	28,289	0.31
Spain	33,001,402	72,673	0.22
Singapore	6,893,611	14,660	0.21
Hungary	1,365,437	2,888	0.21
United States*	398,397,264	690,893	0.17
Canada	23,746,017	40,591	0.17
Denmark	1,755,037	2,810	0.16
Portugal	7,862,269	10,685	0.14
Finland	1,767,811	2,314	0.13
Ireland	3,220,413	4,061	0.13
Greece	1,333,445	1,666	0.12
Czech Republic	2,976,608	3,315	0.11
Norway	2,322,857	2,496	0.11
United Arab Emirates	4,185,813	4,382	0.10
Australia	30,803,582	30,117	0.10
Saudi Arabia	1,592,424	1,494	0.09
New Zealand	2,343,417	1,899	0.08
Egypt	1,446,662	1,074	0.07
South Africa	5,994,396	3,847	0.06
United Kingdom	56,970,613	36,348	0.06
Indonesia	4,545,686	2,726	0.06
Switzerland	8,118,464	4,771	0.06
Sweden	5,108,164	2,836	0.06
South Korea	3,420,881	1,829	0.05
Turkey	2,107,903	1,050	0.05
Italy	60,830,641	26,749	0.04
Colombia	12,149,012	4,974	0.04
Germany	71,203,479	26,519	0.04
Austria	4,287,420	1,489	0.03
Mexico	44,892,154	14,470	0.03
Belgium	6,528,233	2,086	0.03
Netherlands	22,703,983	6,840	0.03
Chile	4,960,535	936	0.02
France	89,430,778	15,896	0.02
Brazil	145,060,543	18,224	0.01
Japan	121,699,258	6,038	0.005
All countries	1,246,909,418	1,341,766	0.11

Notes: The table reports, for each country, the number of distinct job advertisements posted between November 2022 and March 2026, the number of those that are LLM-verified AI job advertisements, and the resulting AI share. The unit is a distinct company–advertisement pair, dated by the day it first appears, and the universe covers all postings in the country with none of the sample restrictions on affiliates or companies of Section 3. A country enters the analysis sample only when enough distinct affiliates posted an LLM-verified AI job advertisement there, which excludes Kenya, Nigeria, Latvia, Estonia, Slovenia, Costa Rica, Ethiopia, and Iceland; Russia and Iran are excluded because LinkedIn is officially blocked there, leaving coverage that is not comparable with the rest of the sample. *Shorter data window, applied to both the numerator and the denominator (United States: November 2022 to October 2025). Figure 1 plots the AI shares reported here.

Table A7: Treatment and sample characteristics by country

Country	Affiliates	Treated	Treated share (%)	Median employment	Mean AI exposure	Junior share (%)
Singapore	3,135	1,417	45.2	103	0.517	47.7
Slovakia	500	200	40.0	100	0.578	59.2
Japan	1,783	709	39.8	105	0.539	46.9
South Korea	1,264	490	38.8	114	0.534	51.0
Hungary	1,104	424	38.4	104	0.570	60.4
Israel	1,860	665	35.8	97	0.599	48.9
Luxembourg	489	170	34.8	99	0.537	51.7
Ireland	2,429	789	32.5	104	0.497	62.2
Czech Republic	2,034	599	29.4	96	0.547	59.7
Poland	5,009	1,470	29.3	105	0.560	61.0
Lithuania	571	163	28.5	96	0.561	55.8
Portugal	2,700	694	25.7	107	0.506	65.5
Greece	1,238	311	25.1	98	0.517	65.3
Austria	1,877	458	24.4	95	0.529	53.2
Spain	9,081	2,007	22.1	107	0.501	61.6
Belgium	4,230	903	21.3	100	0.506	61.8
Switzerland	4,128	871	21.1	103	0.499	54.3
New Zealand	1,809	372	20.6	107	0.498	58.8
India	36,558	7,487	20.5	114	0.573	60.9
Colombia	3,781	774	20.5	119	0.527	63.4
Germany	15,965	3,262	20.4	99	0.530	52.0
South Africa	4,275	869	20.3	112	0.503	67.2
Mexico	7,435	1,507	20.3	109	0.519	63.0
Finland	1,872	369	19.7	108	0.506	59.4
United Arab Emirates	5,354	1,039	19.4	107	0.469	65.8
Denmark	3,273	623	19.0	109	0.487	61.8
Australia	11,323	2,018	17.8	105	0.484	59.2
Italy	9,784	1,734	17.7	104	0.502	66.3
Chile	3,117	552	17.7	116	0.492	62.0
Sweden	4,803	847	17.6	112	0.494	60.8
Egypt	3,395	583	17.2	115	0.518	64.8
Canada	16,696	2,853	17.1	103	0.487	64.5
Norway	2,288	385	16.8	108	0.481	54.9
United Kingdom	28,408	4,651	16.4	104	0.494	56.0
Turkey	5,253	823	15.7	108	0.503	68.1
Saudi Arabia	3,562	527	14.8	126	0.479	65.3
Netherlands	12,178	1,718	14.1	105	0.479	64.0
Indonesia	5,907	801	13.6	124	0.502	72.3
France	19,410	2,626	13.5	106	0.497	59.1
Brazil	22,686	2,353	10.4	118	0.500	73.9
United States*	123,074	9,933	8.1	124	0.488	58.1
All countries	395,638	61,046	15.4	112	0.505	60.5

Notes: The table reports, for each country, the number of affiliates that satisfy the sample restrictions of Section 3, the number and share of them that are treated, that is, whose company posted at least one LLM-verified AI job advertisement in any country between November 2022 and March 2026, the median number of positions in October 2022, and unweighted affiliate-level means of two baseline characteristics: the workforce's average AI exposure, defined as the occupation-level GPT-4 exposure score of Eloundou et al. (2024), and the junior share of employment, defined as Revelio seniority levels 1–2. Both characteristics are measured in October 2022, the month before ChatGPT's launch. A country enters the analysis sample only when enough distinct companies posted an LLM-verified AI job advertisement there, which excludes Kenya, Nigeria, Latvia, Estonia, Slovenia, Costa Rica, Ethiopia, and Iceland; Russia and Iran are excluded because LinkedIn is officially blocked there, leaving coverage that is not comparable with the rest of the sample. *United States: its AI-posting data end in October 2025, so its treated count and treated share cover a shorter window than the countries measured through March 2026 and are understated; the treatment-share rankings in the corresponding figures therefore exclude it.

Table A8: AI job postings at adopting companies

Panel A: Job postings per adopting company	Mean	P25	Median	P75	P90
All job postings	7,817	263	812	2,970	11,101
AI job postings	46.9	1.0	3.0	11.0	36.0
AI share of postings (%)	1.82	0.18	0.56	1.65	4.08

Panel B: AI versus other job postings	AI	Other	Difference
Average exposure of the occupation group	0.621	0.485	0.136
In a high-exposure occupation group (%)	57.9	26.4	31.6
In an occupation group new to the affiliate (%)	6.4	3.3	3.1
Number of postings	1,212,401	199,485,782	

Panel C: Occupation of the posting	Share of postings (%)		Adopting	AI share of	Group
	AI	Other	companies (%)	the group (%)	exposure
Computer & Mathematical	56.8	18.2	79.6	1.86	0.746
Educational Instruction & Library	18.5	2.2	3.6	4.89	0.403
Management	5.0	12.9	19.8	0.23	0.467
Business & Financial Operations	4.0	10.2	18.3	0.24	0.546
Arts, Design, Entertainment & Media	3.4	2.3	14.0	0.87	0.538
Architecture & Engineering	3.3	6.0	15.0	0.33	0.505
Life, Physical & Social Science	3.0	2.2	5.7	0.81	0.506
Sales & Related	2.3	9.7	6.6	0.14	0.524
All other occupations	3.7	36.2	8.9	0.06	0.416
Total	100.0	100.0	100.0	0.60	0.621

Notes: The table describes the job postings of the 26,174 companies that posted at least one LLM-verified AI job advertisement between November 2022 and March 2026, 7.9 percent of the 332,898 companies meeting the sample restrictions of Section 3 in at least one of the 41 sample countries. A company is a Revelio company identifier, the level at which adoption is defined, and its postings are pooled across the sample countries in which it is eligible; a posting is a distinct company–advertisement pair dated by its first posting date. United States postings, AI and total alike, are counted through October 2025, when its AI verification data end. Pooled across adopting companies, AI postings are 0.60 percent of all postings. They are concentrated: the ten largest AI posters account for 40 percent of them, so the quantiles in Panel A and the *adopting companies* column of Panel C — the share of adopting companies with at least one AI posting in the group, which therefore sums to more than 100 — weight every company equally and describe the typical adopter better than the posting-weighted shares do. Panel C's company percentages use the 25,999 companies with at least one AI posting carrying a usable occupation code. Panels B and C are at the posting level and cover the postings carrying a usable occupation code (98.8 percent of AI and 98.1 percent of other postings); occupations are the two-digit SOC major groups, and a group's exposure is the October 2022 employment-weighted average GPT-4 exposure score (Eloundou et al., 2024) of the occupations in it (AI-posting-weighted across groups in the last two rows). A high-exposure group is one whose average exposure exceeds the top-quartile cutoff used for the paper's exposure outcomes (0.58): Computer & Mathematical and Office & Administrative Support. A posting is in a group new to the affiliate if the company posted no advertisement in that group in the same country between January 2021 and October 2022; this row is the only statistic in the table measured at the affiliate level, every other statistic pools a company's postings across countries. The sample keeps countries where more than 50 distinct companies posted a verified AI job and excludes Russia and Iran, where LinkedIn is blocked.

A8 Revelio Labs Representativeness

This section benchmarks Revelio Labs data against official labor force statistics. For each country and year we count the positions active in that year, map their O*NET codes to the nine major groups of the International Standard Classification of Occupations, and compare the counts and shares against Eurostat for 21 countries, and against the International Labour Organization for 14 others. The benchmark covers 35 of the 41 countries in our sample, and we use the most recent year each source reports.[31]

Coverage varies by a factor of thirty across countries (Table A9). Active Revelio positions amount to 67 percent of official employment in the United States and 63 percent in the United Kingdom, and 13 countries reach at least half. The median country reaches 29 percent. At the other end, Revelio captures 24 percent of employment in Germany and Brazil, 7.9 percent in India, and 2.1 percent in Japan, where domestic career platforms are responsible for most professional hiring. Singapore exceeds official employment outright, at 330 percent, because regional headquarters there list workers whom the national survey counts elsewhere or not at all.

Coverage is uneven across occupations as well. In the median country, managers hold 24 percent of Revelio positions against 5 percent of official employment. Professionals and technicians together hold 51 percent against 38 percent.[32] The four manual groups, which cover agriculture, craft and trades, plant and machine operators, and elementary occupations, hold 5.3 percent of Revelio positions against 31 percent of official employment. Figure A4 shows the pattern country by country. The correspondence between the two share distributions is close in countries where coverage is high, with a correlation of 0.73 in the United States, 0.87 in the United Kingdom, and 0.85 in Ireland, but low when coverage is thin.

Figure A5 plots year-over-year growth in Revelio against year-over-year growth in the official series, pooling countries, years, and occupation groups from 2022 through 2024. Revelio grows faster than the official series in eight of the nine groups, by 0.5 percentage points a year for professionals and by 2.7 points for technicians. Within countries, the correlation between national growth rates and Revelio Labs growth

[31] Eurostat publishes employed persons by two-digit occupation in its Labour Force Survey, and the ILO publishes the same series through its Data Explorer. Both classify workers by ISCO-08, so we aggregate their two-digit codes to the nine major groups and aggregate the O*NET codes in Revelio to the same nine groups. Armed forces occupations fall out of the comparison, because O*NET is a civilian classification. Canada, the Netherlands, New Zealand, Saudi Arabia, South Africa, and South Korea report occupation on national classifications that we do not harmonize here.

[32] Part of the excess of professional workers is an artifact of the crosswalk. O*NET has no counterpart to the ISCO group of technicians and associate professionals, so our mapping matches healthcare support occupations to that group and matches the remaining candidates to professionals. Reporting professionals and technicians as one block cuts the median gap from 28 to 15 percentage points. The managerial and manual gaps do not depend on this choice.

Table A9: Revelio coverage of official employment, by country

Country	Coverage (%)	Share correlation	Country	Coverage (%)	Share correlation
Singapore	330	0.80	Israel	29	0.78
Luxembourg	91	0.81	Austria	27	0.50
Denmark	76	0.67	Greece	25	0.65
United States	67	0.73	Lithuania	24	0.81
Switzerland	65	0.72	Brazil	24	0.12
United Kingdom	63	0.87	Czech Republic	24	0.40
Sweden	62	0.79	Germany	23	0.48
Australia	59	0.83	Poland	19	0.70
Ireland	59	0.85	Turkey	16	0.08
Belgium	56	0.78	Hungary	15	0.41
Norway	55	0.77	Colombia	15	−0.08
France	53	0.65	Slovakia	14	0.37
United Arab Emirates*	51	0.34	Mexico	11	−0.20
Portugal	47	0.69	Egypt	11	−0.07
Finland	45	0.60	India	8	−0.32
Italy	40	0.26	Indonesia*	5	−0.32
Spain	38	0.48	Japan*	2	0.16
Chile*	29	0.34			
Median, all 35 countries	29	0.60			

Notes: The table compares the Revelio workforce with official labor force statistics for the 35 sample countries that report employment by occupation on a harmonized classification. Coverage is the number of Revelio positions active in the reference year, as a percentage of official employment in the same year, summing over the nine ISCO-08 major groups. Revelio counts positions rather than people, so a worker who holds two jobs enters the numerator twice. The share correlation is the correlation between the Revelio and the official vector of employment shares across the same nine groups, so it measures whether the occupational mix in Revelio matches the mix in the economy, and it is invariant to the level of coverage. Official employment comes from the Eurostat Labour Force Survey for 21 countries and from the ILO Data Explorer for 14 countries. Armed forces occupations are excluded from both series. The reference year is 2024 unless marked. *Latest year the official source reports (United Arab Emirates, Chile, and Japan: 2022; Indonesia: 2023). Figure A4 plots the underlying occupation shares.

rates is also relatively weak.

Figure A4: Occupation employment shares in Revelio and official statistics

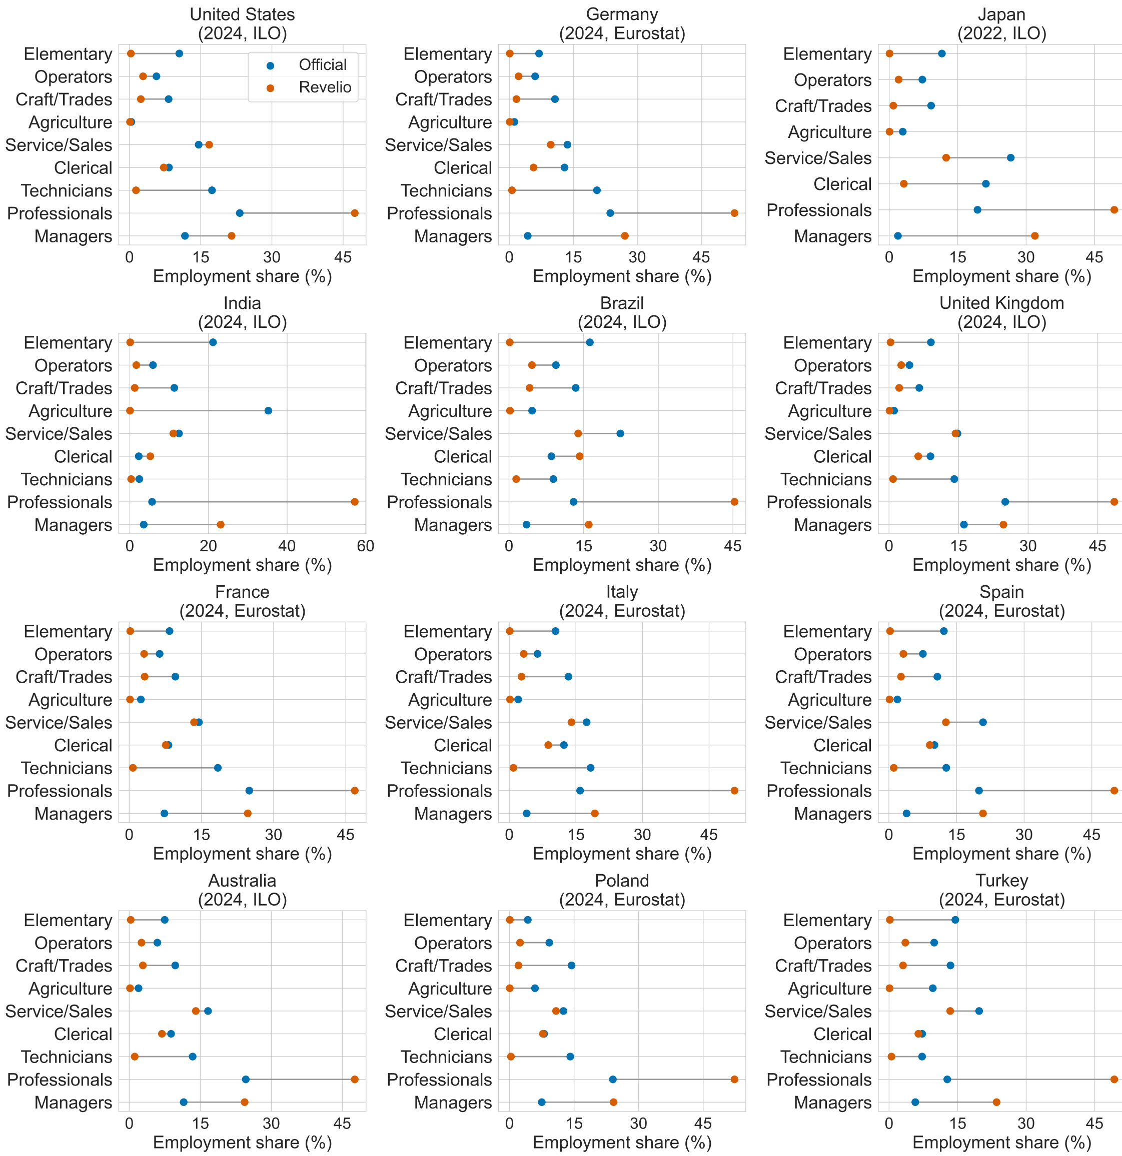

Notes: Each panel plots, for one country, the share of employment in each of the nine ISCO-08 major groups. Blue markers give the official share and orange markers the Revelio share, and the connecting line measures the gap between them. Revelio shares count the positions active in the year shown in the panel heading, classified by the O*NET code Revelio assigns to the position and aggregated to ISCO major groups. Official shares come from Eurostat or from the ILO, as the panel heading records. Armed forces occupations are excluded from both series. Table A9 lists the source and reference year for every country.

Figure A5: Employment growth in Revelio and official statistics, by occupation

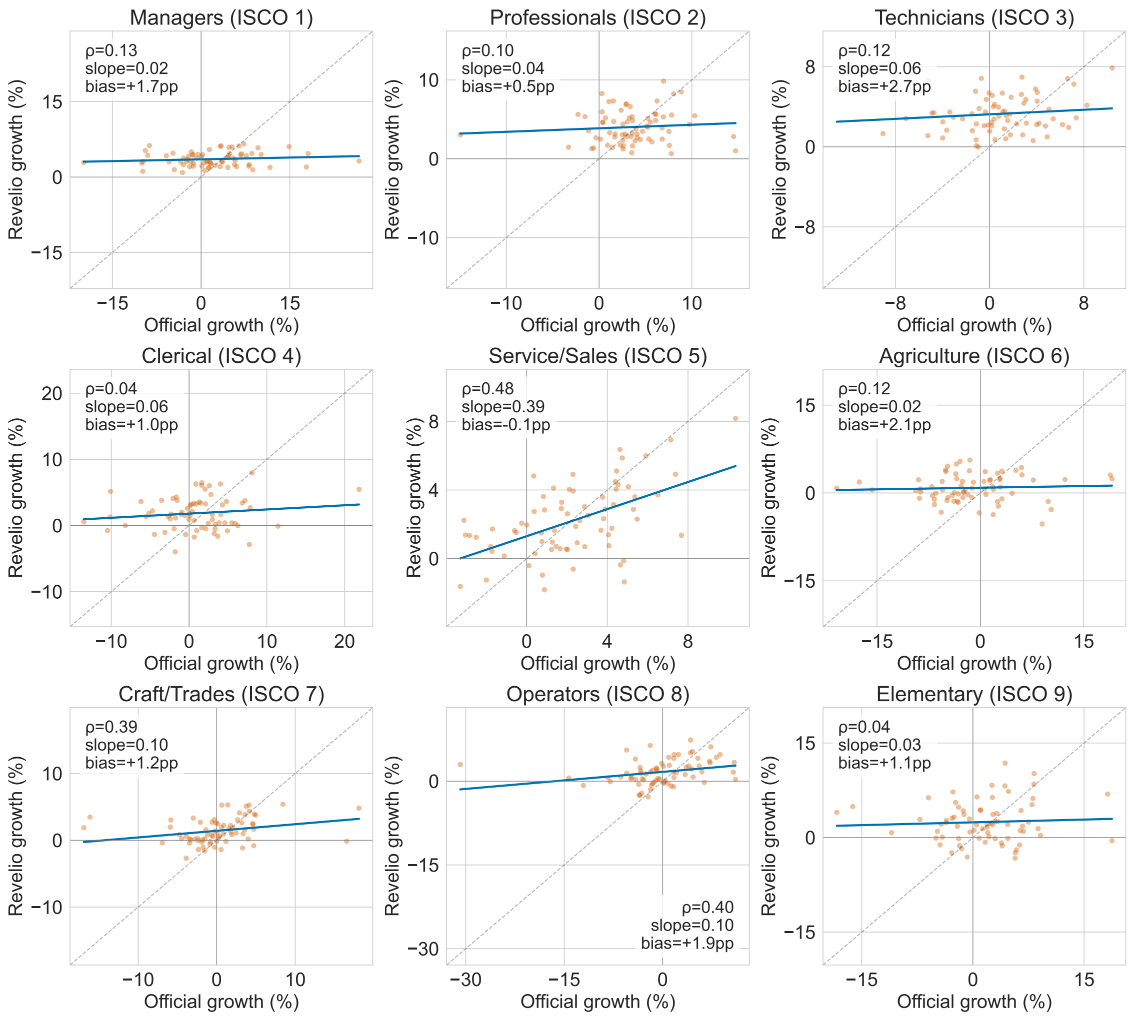

Notes: Each panel plots year-over-year employment growth in Revelio against year-over-year employment growth in the official series for one ISCO-08 major group. A marker is a country-year, pooling the 24 countries for which both sources report at least three consecutive years between 2021 and 2024. The solid line is the ordinary least squares fit and the dashed line is the 45-degree line along which the two sources would agree. Each panel reports the Spearman rank correlation, the fitted slope, and the mean gap in percentage points between Revelio growth and official growth. For the construction of the two series, see Figure A4.

A9 The Matched Sample

This appendix documents the matched sample of Section 4. Table A10 reports how the matching step selects the sample and how well it balances the matching variables and the outcomes.

Table A10: The matched sample

Panel A: Sample sizes						
Treated affiliates in the full sample						37,993
Control affiliates in the full sample						40,551
Matched sets (treated affiliates matched)						26,102 (68.7%)
dropped by the calipers						11,819 (31.1%)
dropped for lack of controls in the cell						72 (0.2%)
Distinct control affiliates matched						27,697 (68.3%)
Controls per treated affiliate, mean and median						3.57 and 5
Treated companies, full sample and matched						10,689 and 9,473
Control companies, full sample and matched						29,938 and 21,405
Panel B: Baseline means, October 2022						
	Full sample			Matched sample		
	Treated	Control	Std. diff.	Treated	Control	Std. diff.
Matching variables						
Log employment	5.317	4.755	0.575	5.012	4.966	0.053
Average AI exposure	0.555	0.501	0.523	0.551	0.550	0.016
Junior share of employment	0.532	0.631	-0.557	0.543	0.547	-0.025
Outcomes						
Log junior employment	4.609	4.245	0.337	4.334	4.300	0.034
Log senior employment	4.461	3.620	0.772	4.138	4.085	0.057
Q4−Q1 exposure share, total	0.224	0.036	0.545	0.216	0.204	0.036
Q4−Q1 exposure share, junior	0.248	0.039	0.531	0.239	0.229	0.026
Q4−Q1 exposure share, senior	0.225	0.067	0.549	0.218	0.197	0.074
Affiliates	37,993	40,551		26,102	27,697	

Notes: The table describes the matched sample of Section 4. The full sample is the set of affiliates that enter the matching step: eligible affiliates whose parent is headquartered in another country, with a defined headquarters-metro instrument and non-missing baseline employment and AI exposure. Treated affiliates are those whose company posted at least one AI job advertisement after November 2022; control affiliates belong to corporate groups that never posted one, and their companies posted at least one job advertisement of any kind anywhere in our data after that date. Within each country and sector, a treated affiliate is matched with replacement to up to five controls by Mahalanobis distance on the three matching variables, with a caliper of half a standard deviation on each; it drops out when no control lies within the calipers, when its country–sector cell has fewer than two controls with complete matching variables, or when its own matching variables are incomplete. Every control carries weight one over the number of controls in its set, so a control matched to several treated affiliates counts once per set. Panel A percentages are shares of the full-sample treated affiliates (control affiliates for the distinct-controls row). Panel B reports October 2022 means for treated and control affiliates in the full sample, each affiliate once, and in the matched sample at these weights, with the normalized difference: the difference in means divided by the square root of the average of the two variances. Log employment and the junior share at baseline coincide with the first and third matching variables, so they appear once. For the outcome definitions, see Table 3.

A10 Occupation-Level Event Studies

This appendix reports the event studies behind the occupation-level endpoint estimates in the left panel of Figure 5. Each panel plots the effect of AI adoption on one two-digit occupation group's share of affiliate employment. All 22 groups clear the first-stage F screen, so none are omitted. Figure A10 repeats the endpoint estimates of Figure 5 with Bonferroni-corrected confidence intervals.

Figure A6: Event study plots for the effect of AI adoption on affiliate-level occupation shares, SOC 11–21

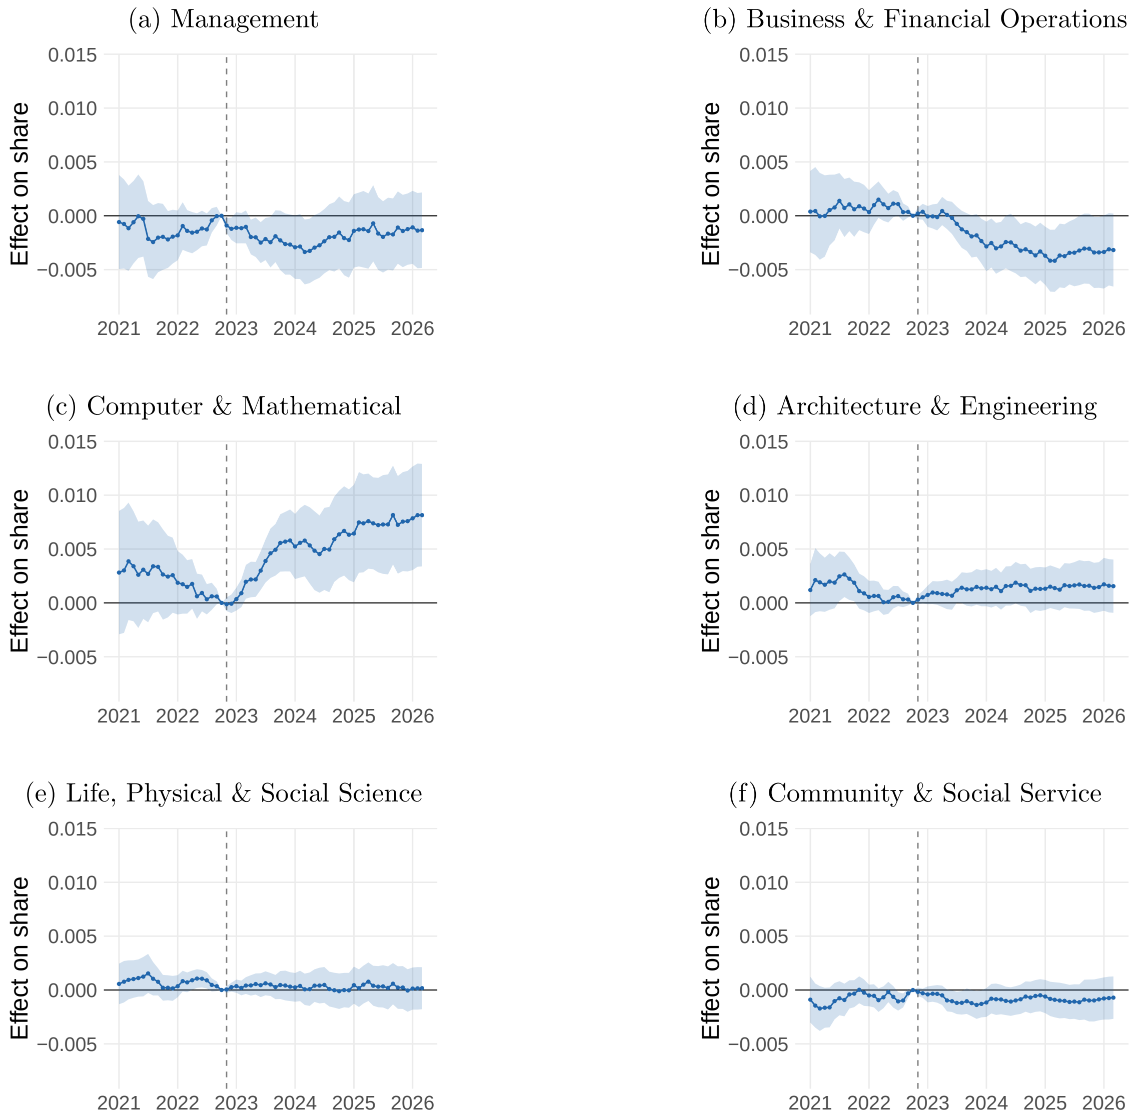

Notes: Each panel plots the estimated effect of AI adoption on one occupation group's share of affiliate employment, month by month and in share units. The coefficients are the two-stage least squares estimates β_q from Equation (15), with the outcome measured as the long difference from October 2022, the month before ChatGPT's launch. Occupations are the two-digit SOC major groups, shown in SOC order across Figures A6–A9, whose panels share a common vertical axis. Each occupation is estimated on its own matched sample: the affiliates that employed it in October 2022, re-matched with the occupation's baseline share added to the main design's matching variables. Shaded bands are 95% confidence intervals, clustered by the metropolitan area of the parent's headquarters. The dashed vertical line marks ChatGPT's launch in November 2022. For the treatment definition, instrument, and matched sample, see Figure 3.

Figure A7: Event study plots for the effect of AI adoption on affiliate-level occupation shares, SOC 23–33

(a) Legal

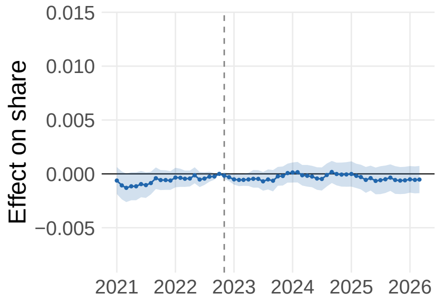

(b) Educational Instruction & Library

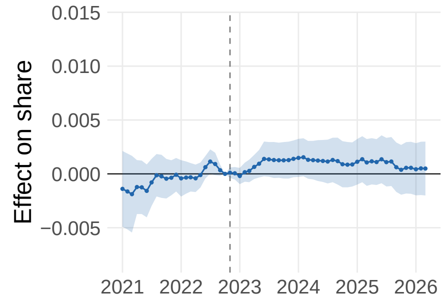

(c) Arts, Design, Entertainment & Media

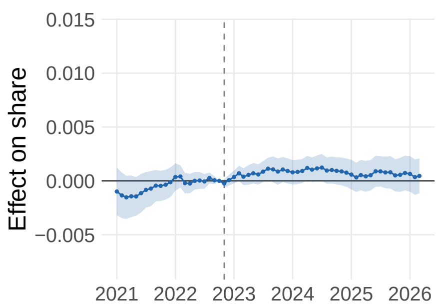

(d) Healthcare Practitioners & Technical

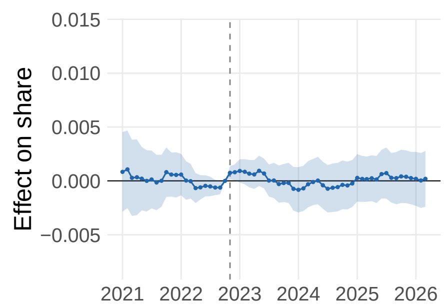

(e) Healthcare Support

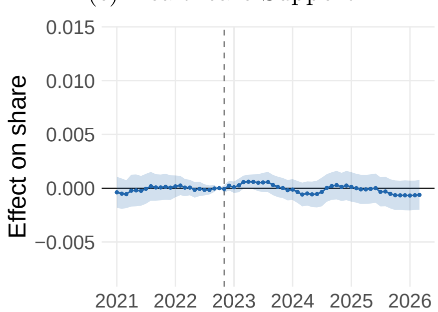

(f) Protective Service

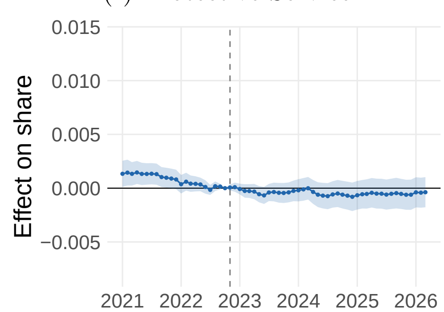

Notes: Each panel plots the estimated effect of AI adoption on one occupation group's share of affiliate employment, month by month and in share units. For all other definitions, see Figure A6.

Figure A8: Event study plots for the effect of AI adoption on affiliate-level occupation shares, SOC 35–45

Notes: Each panel plots the estimated effect of AI adoption on one occupation group's share of affiliate employment, month by month and in share units. For all other definitions, see Figure A6.

Figure A9: Event study plots for the effect of AI adoption on affiliate-level occupation shares, SOC 47–53

Notes: Each panel plots the estimated effect of AI adoption on one occupation group's share of affiliate employment, month by month and in share units. For all other definitions, see Figure A6.

Figure A10: Effect of AI adoption on affiliate-level occupation shares, Bonferroni-corrected intervals

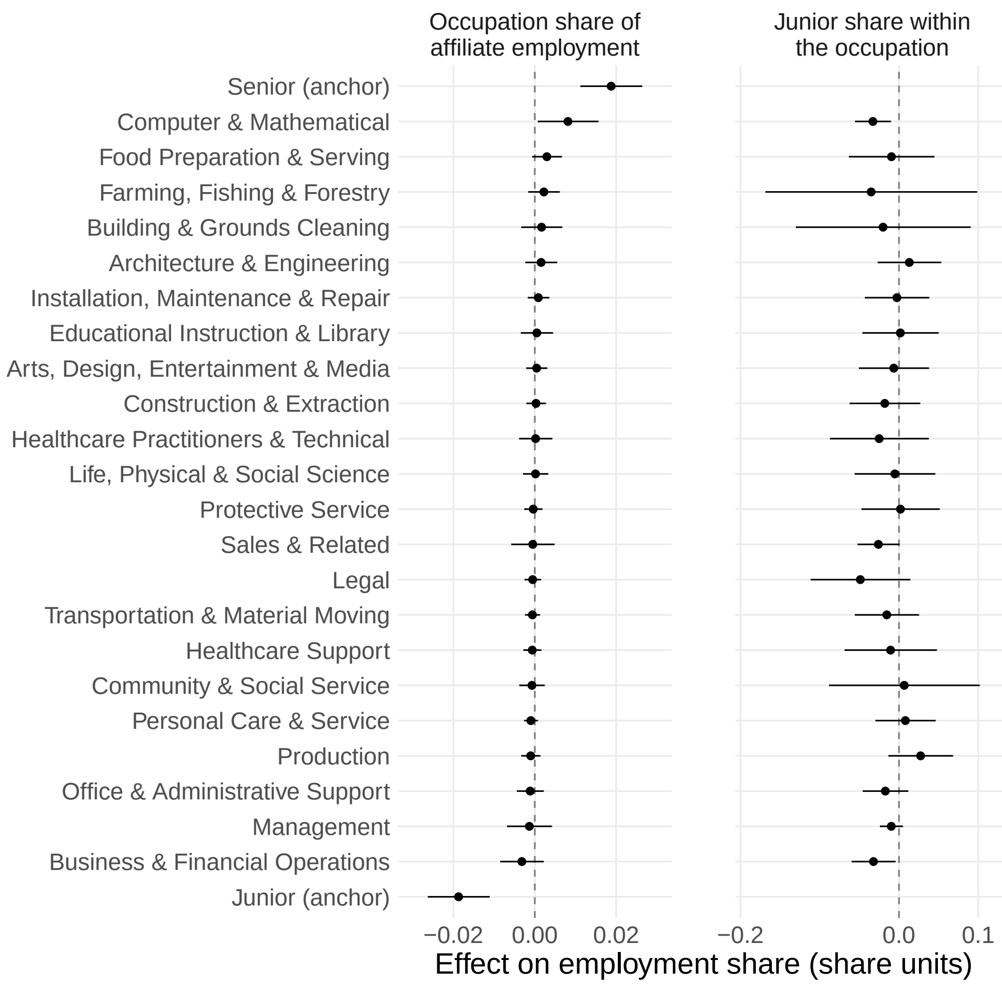

Notes: The figure repeats Figure 5 with the confidence intervals widened by a Bonferroni correction across the occupation groups, which holds the family-wise error rate at 5% over the whole set of occupations rather than at each interval separately. The point estimates are unchanged. The junior and senior anchors keep their uncorrected intervals, as they do not belong to the occupation family. For all other definitions, see Figure 5.

A11 Country-Level Correlations

This appendix reports the correlations behind the country-level results in Section 5.2. Each figure plots the country-level 2SLS endpoint estimates for the four headline outcomes against one country characteristic; the notes to Figure A11 define the sample, weighting, and markers.

Figure A11: Effect of AI adoption on employment and workforce composition, by country GDP per capita

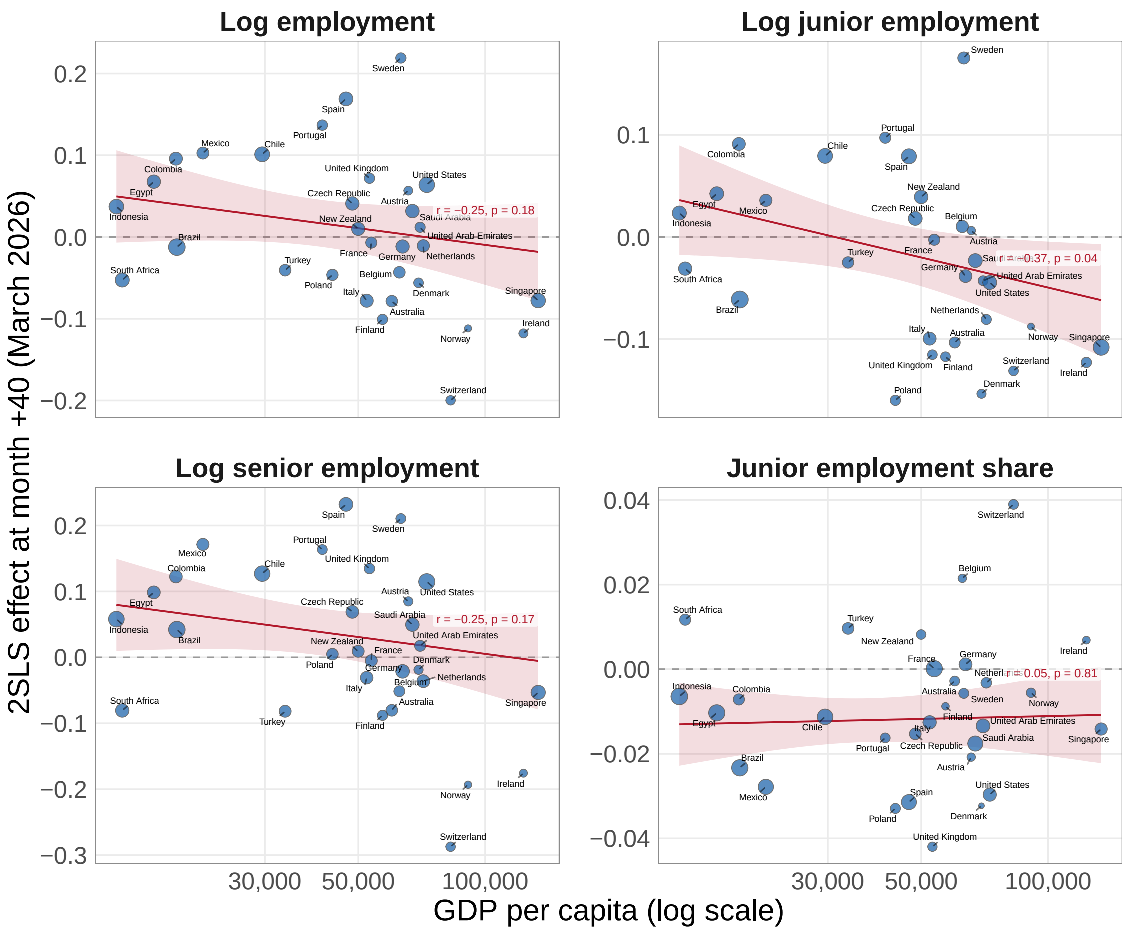

Notes: The figure asks whether the estimated effect of AI adoption varies with how rich a country is. Each panel plots the country-level two-stage least squares endpoint estimates of Figure 6 against GDP per capita, measured in 2022 at purchasing power parity in constant 2021 international dollars (World Bank, 2026) and plotted on a log scale, which is also the scale on which the reported correlation is computed. The top row shows log total employment (left) and log junior employment (right); the bottom row shows log senior employment (left) and the junior share of employment (right). Countries in which the instrument is too weak to support a country-specific estimate are omitted. Point size is proportional to the precision of the country estimate, measured as the inverse of its squared standard error, relative to the average precision in the panel, so that sizes are comparable within but not across panels. The fitted line and shaded band come from a precision-weighted regression, and the annotation beside the line reports the precision-weighted correlation and the p-value of the slope of that regression.

Figure A12: Effect of AI adoption on employment and workforce composition, by country digital readiness

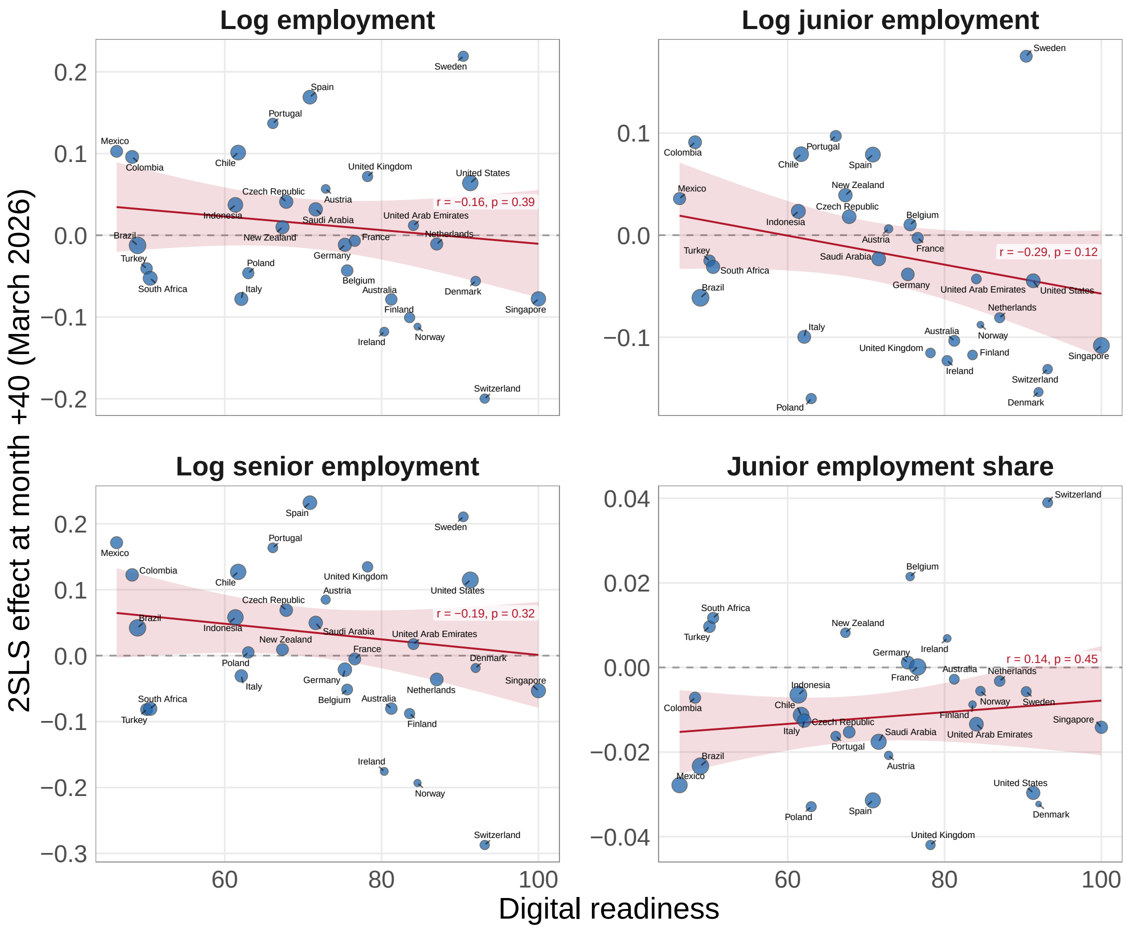

Notes: Each panel plots the country-level two-stage least squares endpoint estimates against the overall score of the 2024 IMD World Digital Competitiveness Ranking (International Institute for Management Development, 2024), a composite measure of a country's capacity to adopt and exploit digital technologies. For all other definitions, see Figure A11.

Figure A13: Effect of AI adoption on employment and workforce composition, by country ICT employment share

Notes: Each panel plots the country-level two-stage least squares endpoint estimates against employment in information and communication (ISIC Rev. 4 section J) as a share of total employment, measured in the latest year available up to 2022 (International Labour Organization, 2026). For all other definitions, see Figure A11.

Figure A14: Effect of AI adoption on employment and workforce composition, by country labor market rigidity

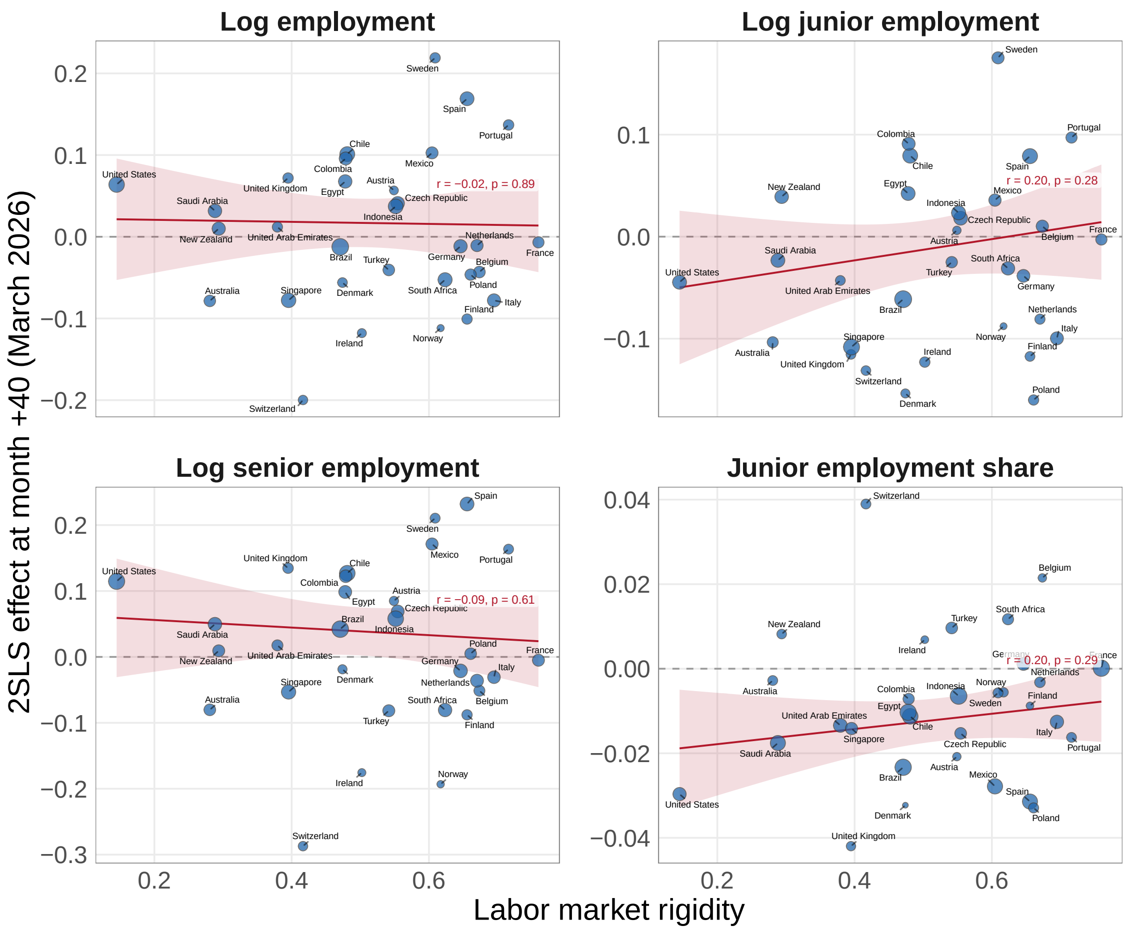

Notes: Each panel plots the country-level two-stage least squares endpoint estimates against labor market rigidity, measured by the CBR Labour Regulation Index in 2022 (Adams et al., 2023). The index reads the statutory and case law of each country and codes indicators of worker protection covering alternative forms of employment, working time, dismissal, employee representation, and industrial action. Each indicator runs from zero to one and the index is their unweighted mean, so higher values indicate more protective labor law. For all other definitions, see Figure A11.

Figure A15: Effect of AI adoption on employment and workforce composition, by country ICT service exports

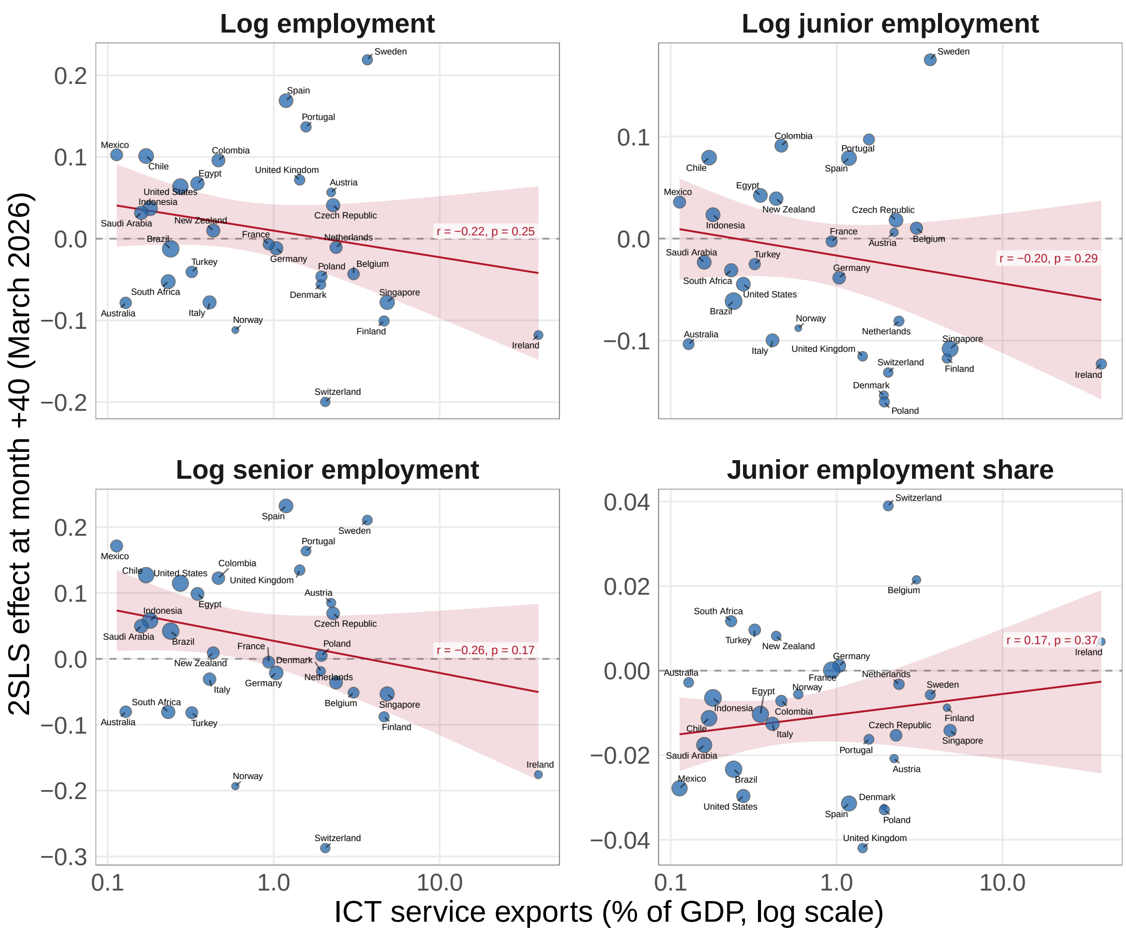

Notes: Each panel plots the country-level two-stage least squares endpoint estimates against exports of telecommunications, computer, and information services as a share of GDP in 2022 (World Bank, 2026). The characteristic is plotted on a log scale, which is also the scale on which the reported correlation is computed, so that Ireland, by far the largest exporter in the sample, does not compress the remaining countries into a single cluster. For all other definitions, see Figure A11.

A12 Additional Heterogeneity Figures

This appendix collects the figures behind Section 5.2.2: the splits by baseline AI exposure (Figures A16–A18), by baseline size (Figures A19–A21), by technology versus non-technology sector (Figures A22–A24), and by the occupation of the company's first AI job (Figures A25–A27).

Figure A16: Event study plots for the effect of AI adoption on employment and workforce composition, by baseline AI exposure

(a) Log employment

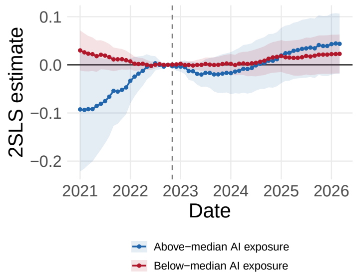

(b) Junior share of employment

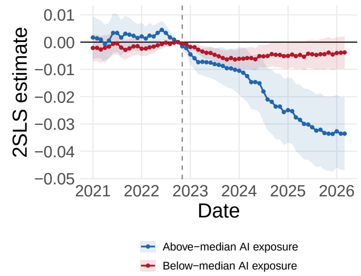

(c) Log junior employment

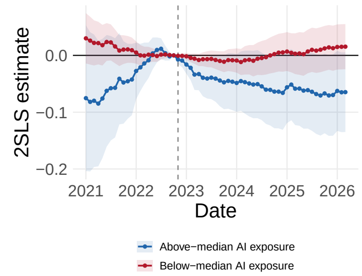

(d) Log senior employment

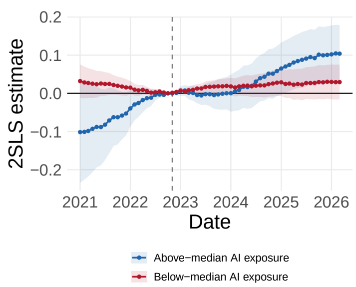

(e) High- minus low-exposure share, junior

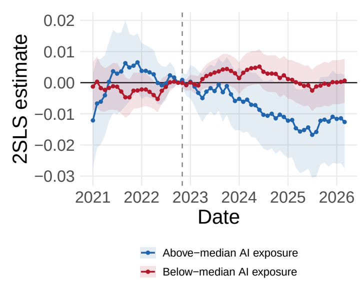

(f) High- minus low-exposure share, senior

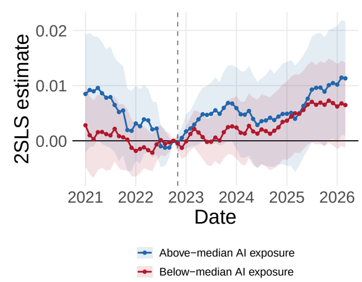

(g) High- minus low-exposure share, overall

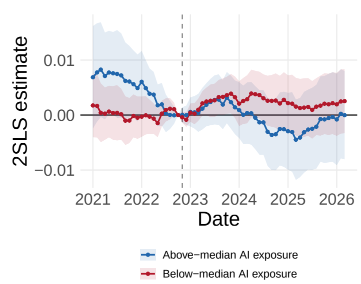

Notes: Each panel overlays the pooled event study estimated separately on matched sets whose treated affiliate sits above (blue) or below (red) the median baseline AI exposure of the matched treated affiliates, where baseline exposure is the average GPT-4 exposure score of the affiliate's October 2022 workforce. Whole matched sets move with their treated affiliate, so a control is never separated from the adopter it was matched to. Treatment, instrument, matching, eligibility, outcomes, and clustering are held at their main-specification values. For all other definitions, see Figures 3 and 4.

Figure A17: Effect of AI adoption on affiliate-level occupation shares, above-median-exposure affiliates

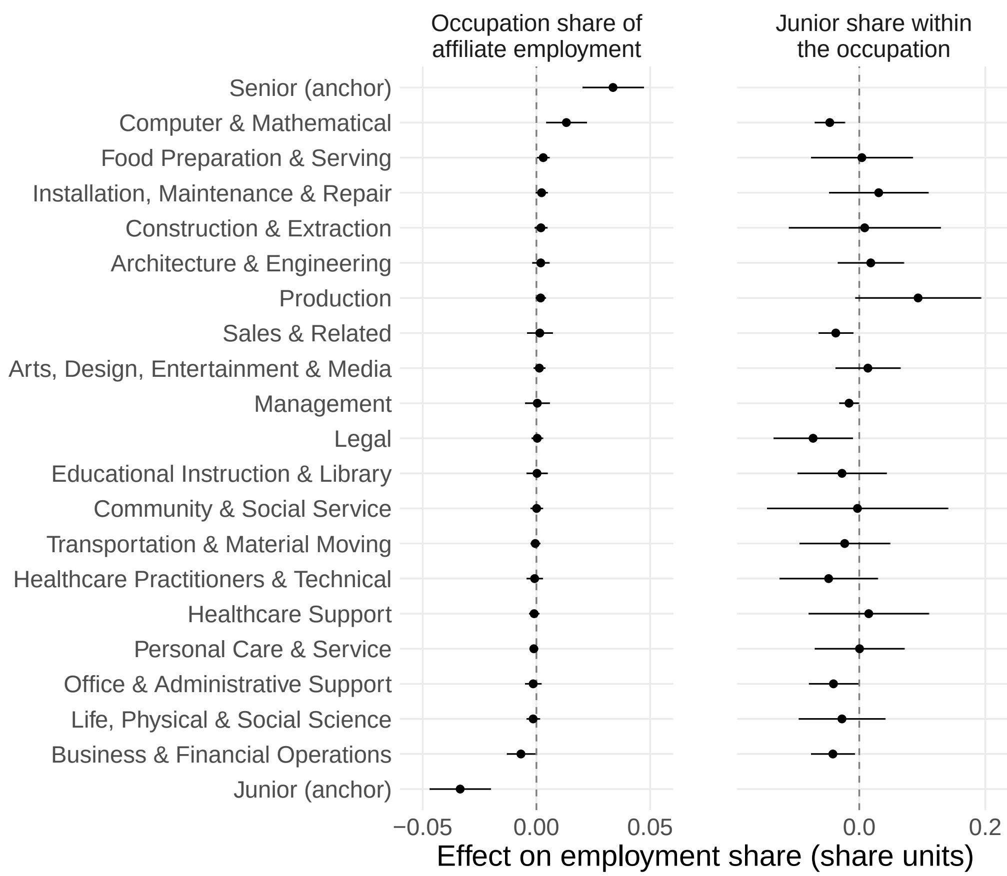

Notes: The figure repeats Figure 5 on matched occupation samples restricted to matched sets whose treated affiliate sits above the median baseline AI exposure of the matched treated affiliates, the above-median group of Figure A16. Occupations whose first stage is too weak in this subsample are omitted. For all other definitions, see Figure 5.

Figure A18: Effect of AI adoption on affiliate-level occupation shares, below-median-exposure affiliates

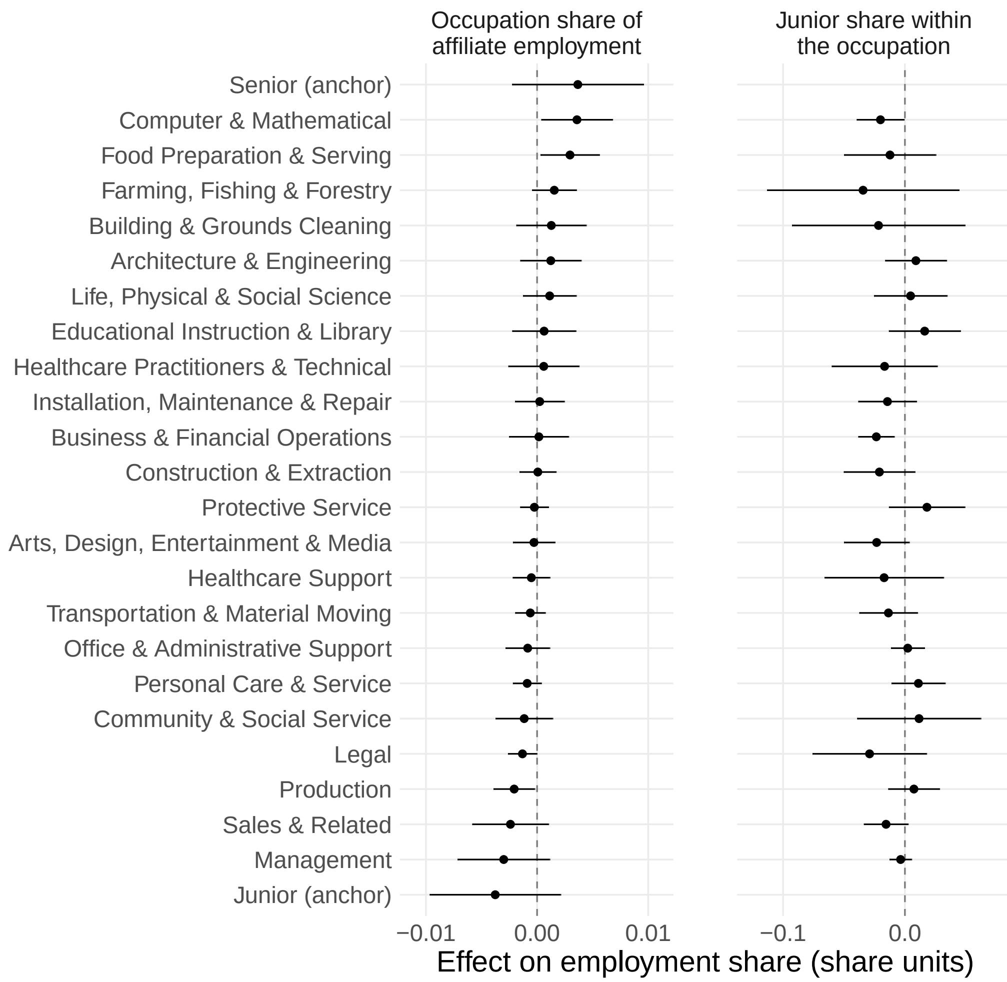

Notes: The figure repeats Figure 5 on matched occupation samples restricted to matched sets whose treated affiliate sits below the median baseline AI exposure. For all other definitions, see Figure A17.

Figure A19: Event study plots for the effect of AI adoption on employment and workforce composition, by baseline affiliate size

(a) Log employment

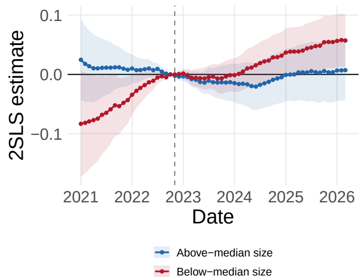

(b) Junior share of employment

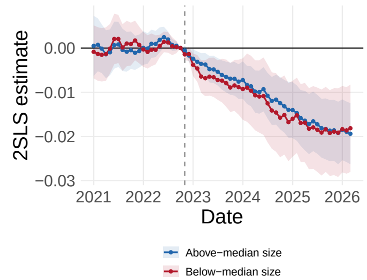

(c) Log junior employment

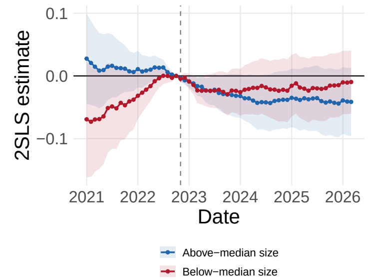

(d) Log senior employment

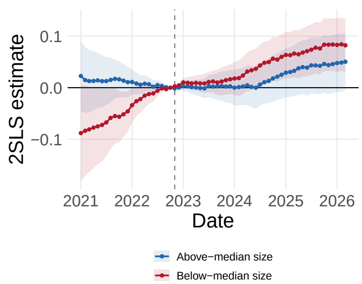

(e) High- minus low-exposure share, junior

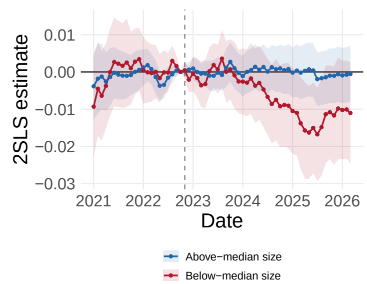

(f) High- minus low-exposure share, senior

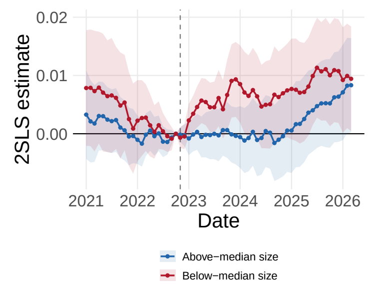

(g) High- minus low-exposure share, overall

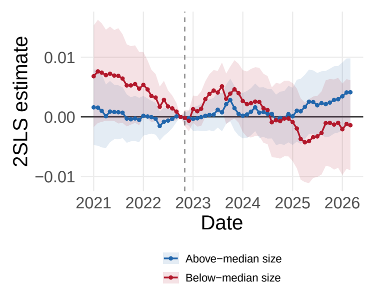

Notes: Each panel overlays the pooled event study estimated separately on matched sets whose treated affiliate sits at or above (blue) or below (red) the pooled median October 2022 employment of the matched treated affiliates, 122 positions. Whole matched sets move with their treated affiliate, so a control is never separated from the adopter it was matched to. Treatment, instrument, matching, eligibility, outcomes, and clustering are held at their main-specification values. For all other definitions, see Figures 3 and 4.

Figure A20: Effect of AI adoption on affiliate-level occupation shares, smaller adopters

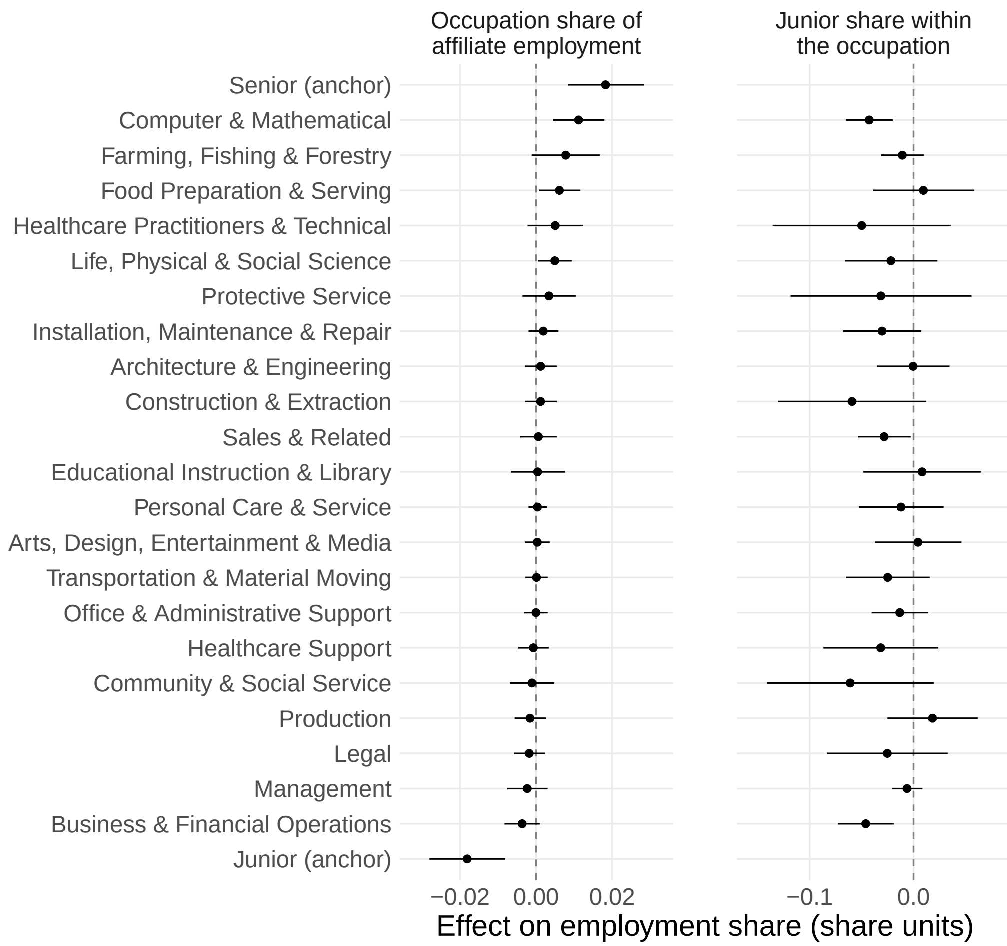

Notes: The figure repeats Figure 5 on matched occupation samples restricted to matched sets whose treated affiliate sits below the pooled median October 2022 employment of the matched treated affiliates, 122 positions, the smaller adopters of Figure A19. Occupations whose first stage is too weak in this subsample are omitted. For all other definitions, see Figure 5.

Figure A21: Effect of AI adoption on affiliate-level occupation shares, larger adopters

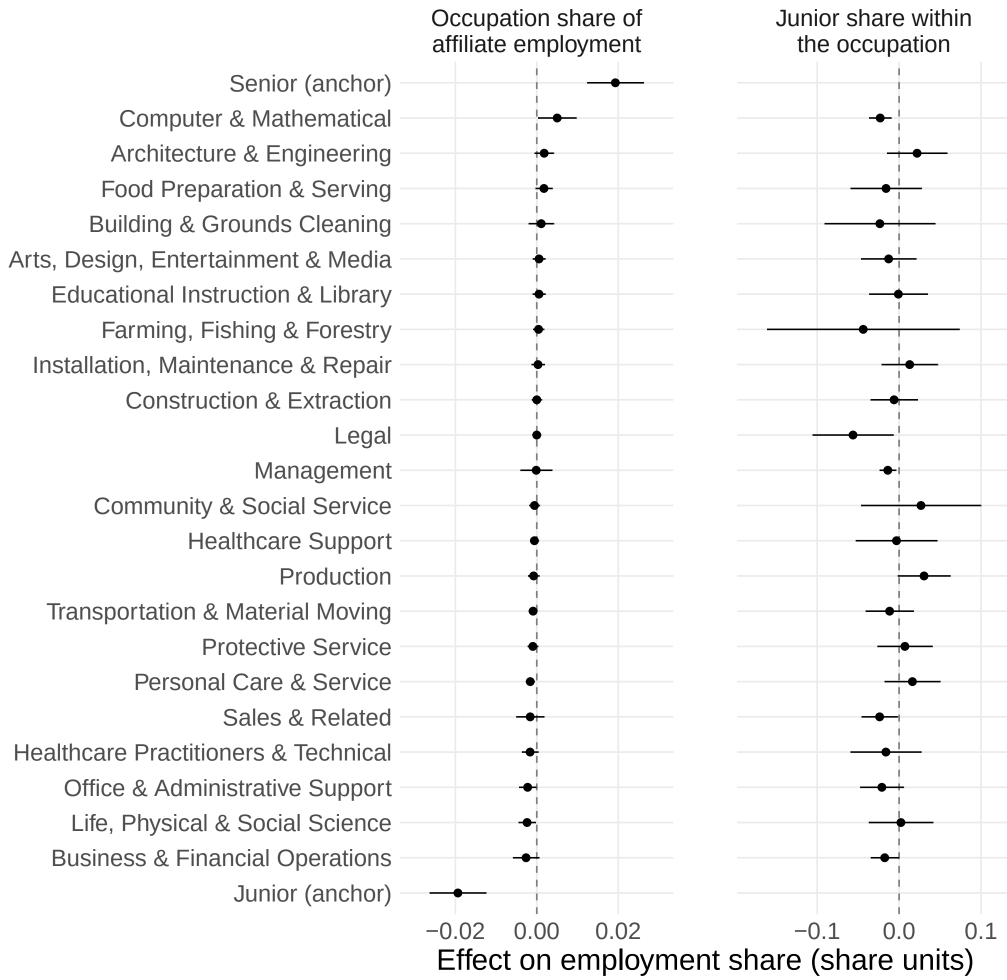

Notes: The figure repeats Figure 5 on matched occupation samples restricted to matched sets whose treated affiliate sits at or above the pooled median October 2022 employment of the matched treated affiliates, 122 positions. For all other definitions, see Figure A20.

Figure A22: Event study plots for the effect of AI adoption on employment and workforce composition, technology versus non-technology affiliates

(a) Log employment

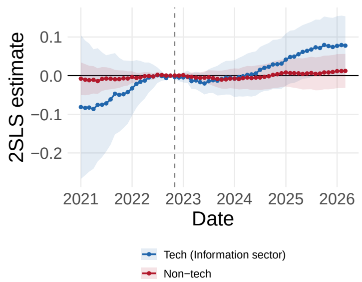

(b) Junior share of employment

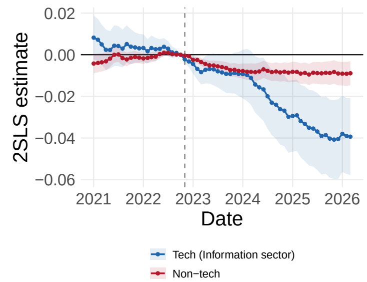

(c) Log junior employment

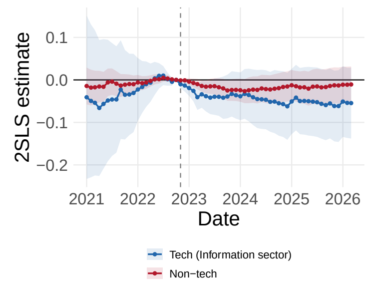

(d) Log senior employment

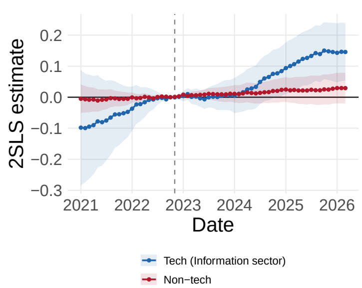

(e) High- minus low-exposure share, junior

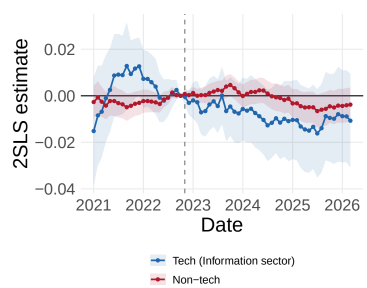

(f) High- minus low-exposure share, senior

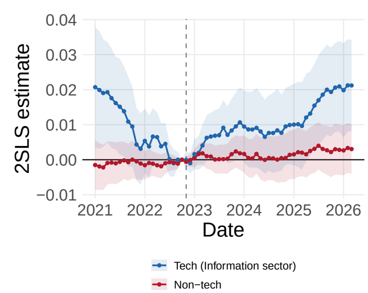

(g) High- minus low-exposure share, overall

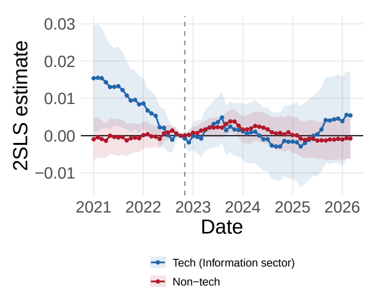

Notes: Each panel overlays the pooled event study estimated separately on technology affiliates (blue), defined as the RICS "Information" sector, and on all other affiliates (red). The matched sample is split by the treated affiliate's sector; because matching is exact on sector, the split partitions whole matched sets. Treatment, instrument, matching, eligibility, outcomes, and clustering are held at their main-specification values. For all other definitions, see Figures 3 and 4.

Figure A23: Effect of AI adoption on affiliate-level occupation shares, technology affiliates

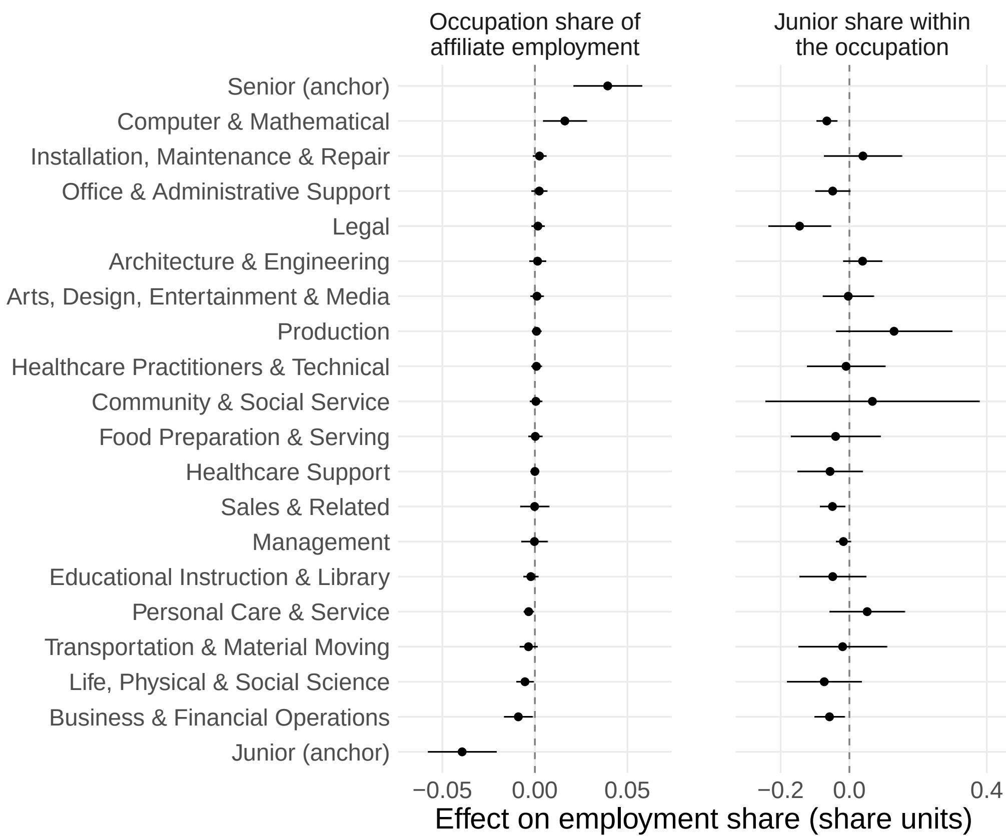

Notes: The figure repeats Figure 5 on matched occupation samples restricted to technology affiliates, defined as the RICS "Information" sector; because matching is exact on sector, the restriction keeps whole matched sets. Occupations whose first stage is too weak in this subsample are omitted. For all other definitions, see Figure 5.

Figure A24: Effect of AI adoption on affiliate-level occupation shares, non-technology affiliates

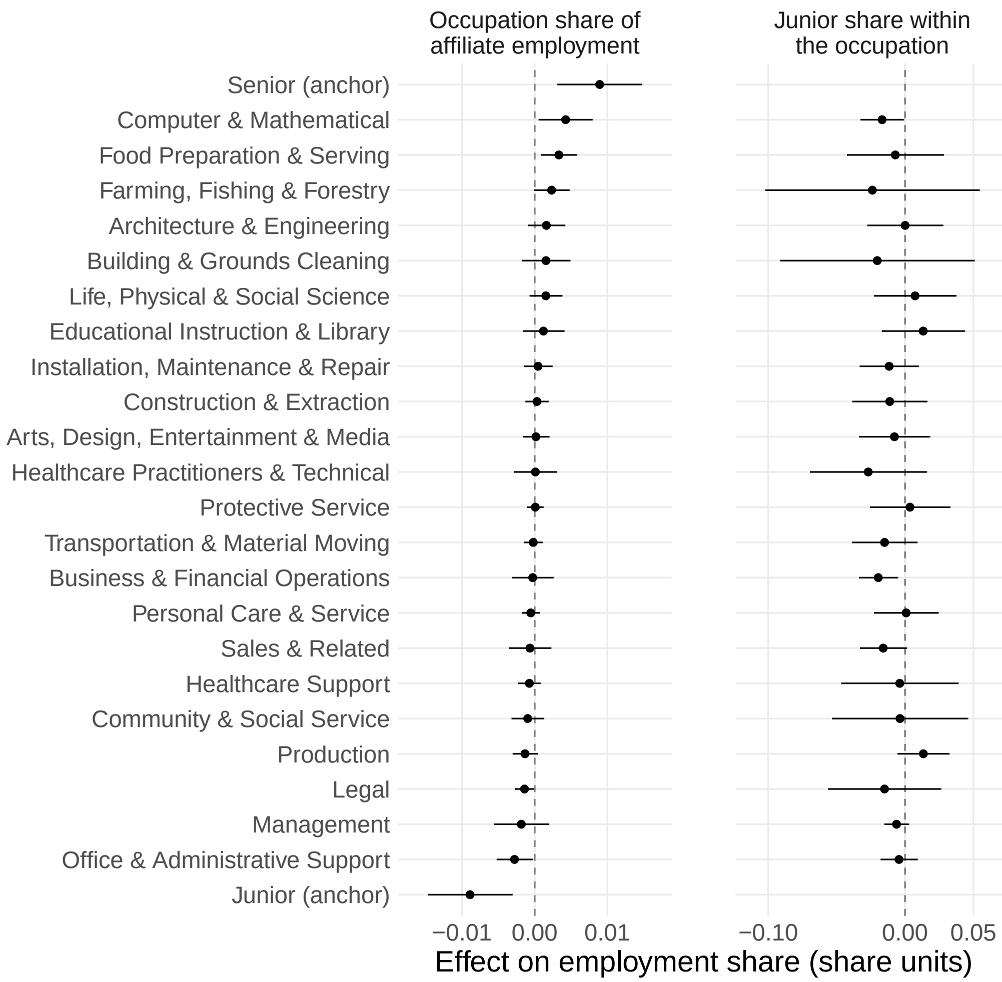

Notes: The figure repeats Figure 5 on matched occupation samples restricted to affiliates outside the technology sector. For all other definitions, see Figure A23.

Figure A25: Junior share of employment, by occupation of the first AI job

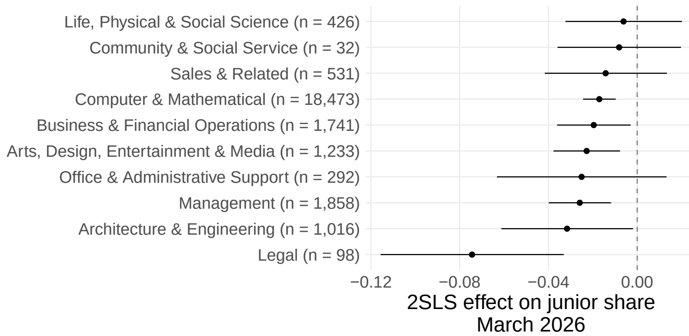

Notes: The figure asks whether the estimated effects depend on the occupation category of a company's first classifiable AI job advertisement. Each row reports the pooled two-stage least squares endpoint estimate at March 2026, re-estimated on the matched sets whose treated affiliate's company's first classifiable AI job advertisement anywhere in our data after ChatGPT's launch falls in the given two-digit occupation category; the label beside each row reports the number of matched treated affiliates behind it. Whole matched sets move with their treated affiliate, and treatment, instrument, matching, eligibility, outcomes, and clustering are held at their main-specification values; categories in which the instrument is too weak to support an estimate are omitted. Whiskers are 95% confidence intervals. For all other definitions, see Figure 3.

Figure A26: Log junior employment, by occupation of the first AI job

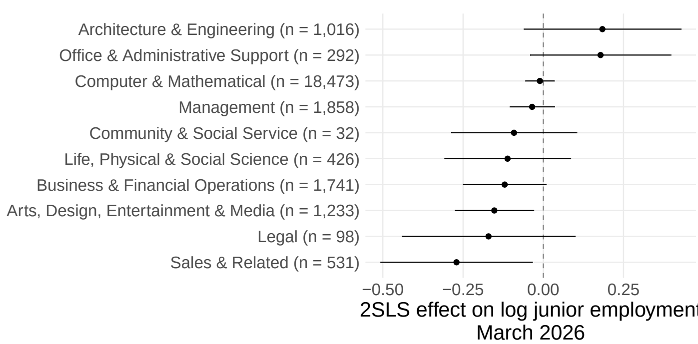

Notes: The figure repeats Figure A25 for log junior employment. For all definitions, see Figure A25.

Figure A27: Log senior employment, by occupation of the first AI job

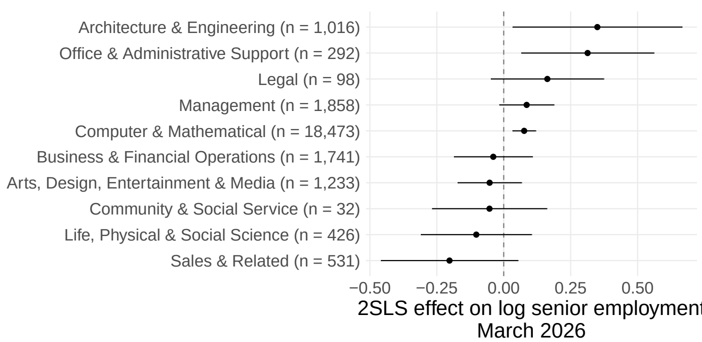

Notes: The figure repeats Figure A25 for log senior employment. For all definitions, see Figure A25.

A13 Instrument-Validity Figures

This appendix collects the figures behind the instrument-validity checks of Section 5.3.1, in the order in which the checks are described there (Figures A28–A32), and the split of the main event studies by the treated affiliate's predicted propensity to adopt (Figure A33).

Figure A28: Event study plots for the effect of AI adoption on employment and workforce composition, peer growth control and sector-purged instrument

(a) Log employment

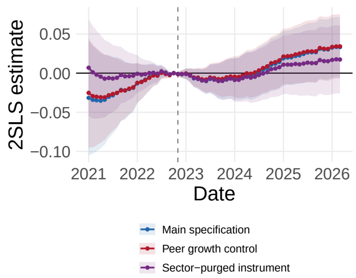

(b) Junior share of employment

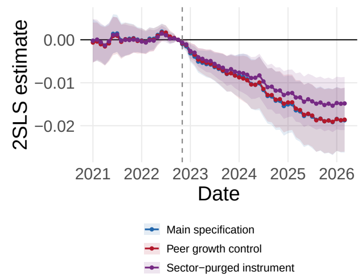

(c) Log junior employment

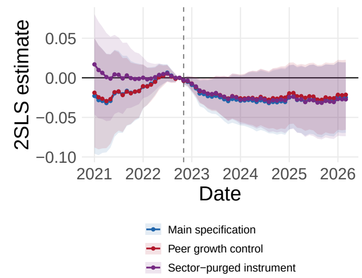

(d) Log senior employment

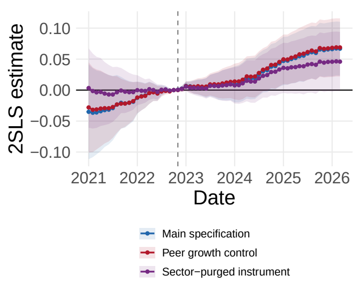

(e) High- minus low-exposure share, junior

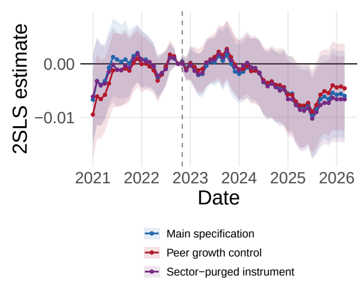

(f) High- minus low-exposure share, senior

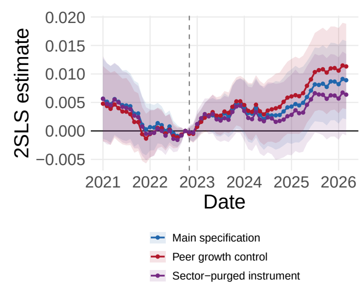

(g) High- minus low-exposure share, overall

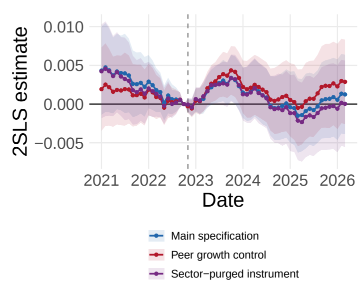

Notes: Each panel overlays the pooled event study under two modifications that target the exclusion restriction. The peer-growth-control series (red) adds to both stages the leave-one-parent-out mean employment growth of multinational parents headquartered in the same metropolitan area, measured across each parent group's eligible entities over the post-launch period relative to the pre-launch period, which absorbs any generic boom around the headquarters that could raise both peer AI adoption and the parent's resources. The sector-purged series (purple) rebuilds the instrument from metro peer parents *outside* the focal parent's own sector, so that within-industry competition and same-industry shocks cannot travel through the instrument. Everything else is held at the main-specification values. For all other definitions, see Figures 3 and 4.

Figure A29: Event study plots for the effect of AI adoption on employment and workforce composition, alternative instrument

(a) Log employment

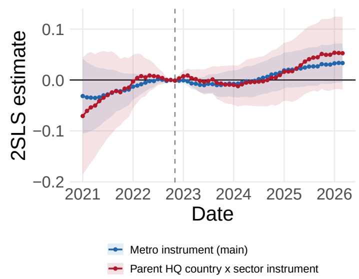

(b) Junior share of employment

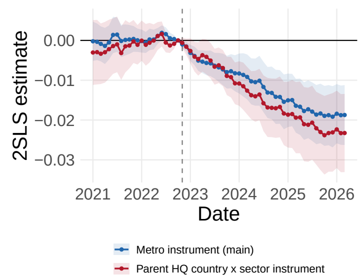

(c) Log junior employment

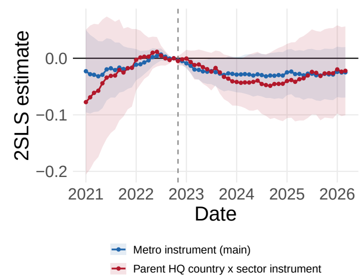

(d) Log senior employment

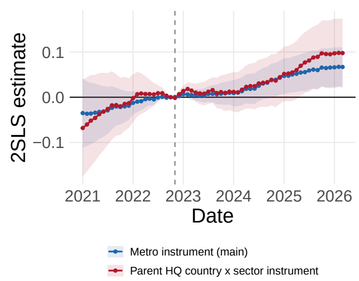

(e) High- minus low-exposure share, junior

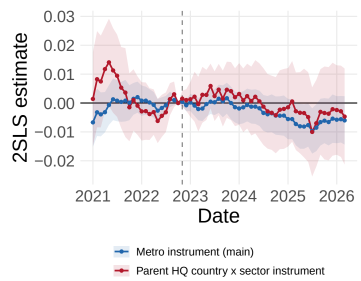

(f) High- minus low-exposure share, senior

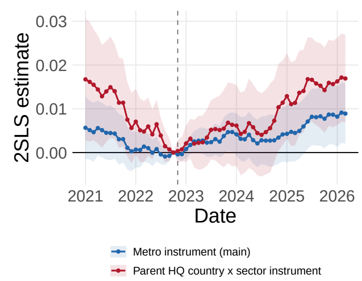

(g) High- minus low-exposure share, overall

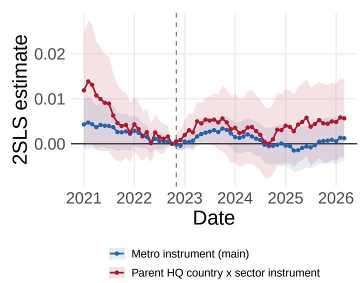

Notes: Each panel overlays the pooled event study under two versions of the instrument. The main specification (blue) instruments AI adoption with the leave-out AI adoption rate among multinational parents headquartered in the same metropolitan area as the affiliate's parent; the alternative (red) uses the leave-out rate among parents in the same headquarters country and RICS sector. Treatment, matching, eligibility, and outcomes are held fixed, and only the instrument and the level at which standard errors are clustered change. Shaded bands are 95% confidence intervals, clustered at the level of whichever instrument is in use. For all other definitions, see Figures 3 and 4.

Figure A30: Event study plots for the effect of AI adoption on employment and workforce composition, alternative adoption definition

(a) Log employment

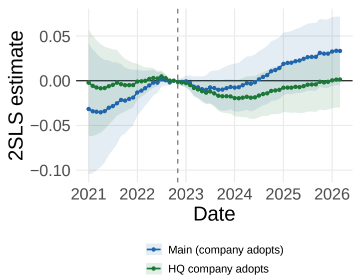

(b) Junior share of employment

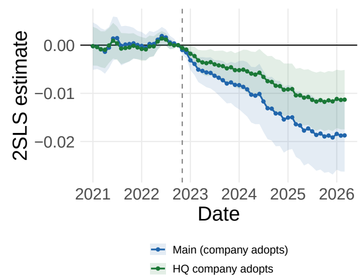

(c) Log junior employment

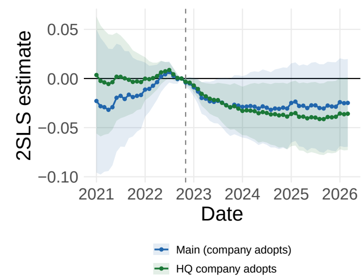

(d) Log senior employment

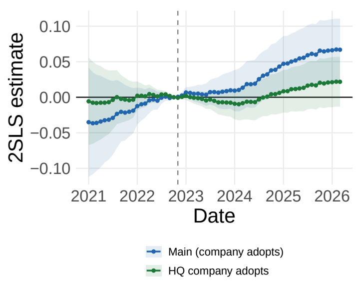

(e) High- minus low-exposure share, junior

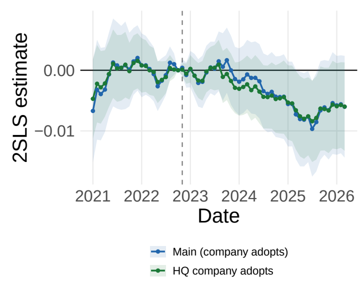

(f) High- minus low-exposure share, senior

(g) High- minus low-exposure share, overall

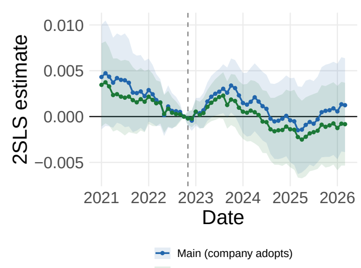

Notes: Each panel overlays the pooled event study under two definitions of treatment. *Main* treats an affiliate as treated if its company posted an AI job advertisement from November 2022 onward, in any country; *HQ adopts* treats every affiliate of a parent as treated if the headquarters entity posted one. Both series draw their controls from the same pool, namely affiliates whose parent group never posted an AI job advertisement anywhere, so only the treated set moves: under *HQ adopts*, affiliates whose own company advertised an AI job under a headquarters entity that did not are dropped rather than recoded as controls. Eligibility, matching, instrument, outcomes, and clustering are held at their main-specification values. For all other definitions, see Figures 3 and 4.

Figure A31: Effect of peer AI adoption on company-level AI adoption, by subgroup

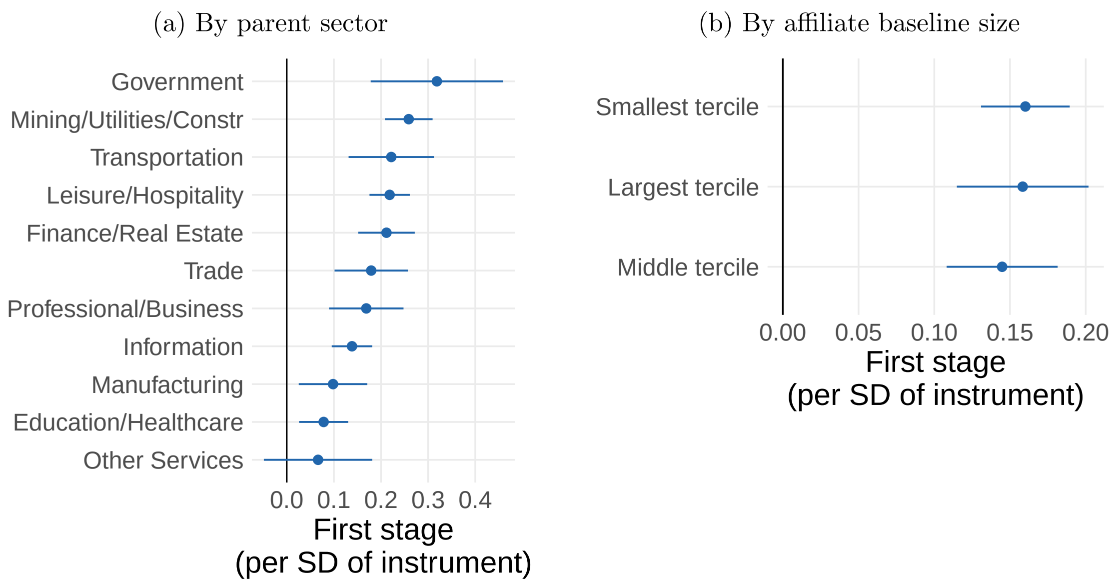

Notes: The figure reports the first stage of the instrument separately within subgroups, as a check on the monotonicity assumption, which requires that peer adoption around the parent never makes an affiliate less likely to be treated and would be violated by a significantly negative coefficient in any subgroup. Each row plots the coefficient from regressing the affiliate's treatment indicator on the standardized instrument, with matched-set fixed effects and design weights, estimated within parent sectors in Panel A and within terciles of baseline affiliate employment in Panel B. Whiskers are 95% confidence intervals, clustered by the metropolitan area of the parent's headquarters. Subgroups with too few treated affiliates or with a single instrument cluster are omitted. For all other definitions, see Figure 3.

Figure A32: Effect of AI adoption on employment and workforce composition, leaving out one metropolitan area at a time

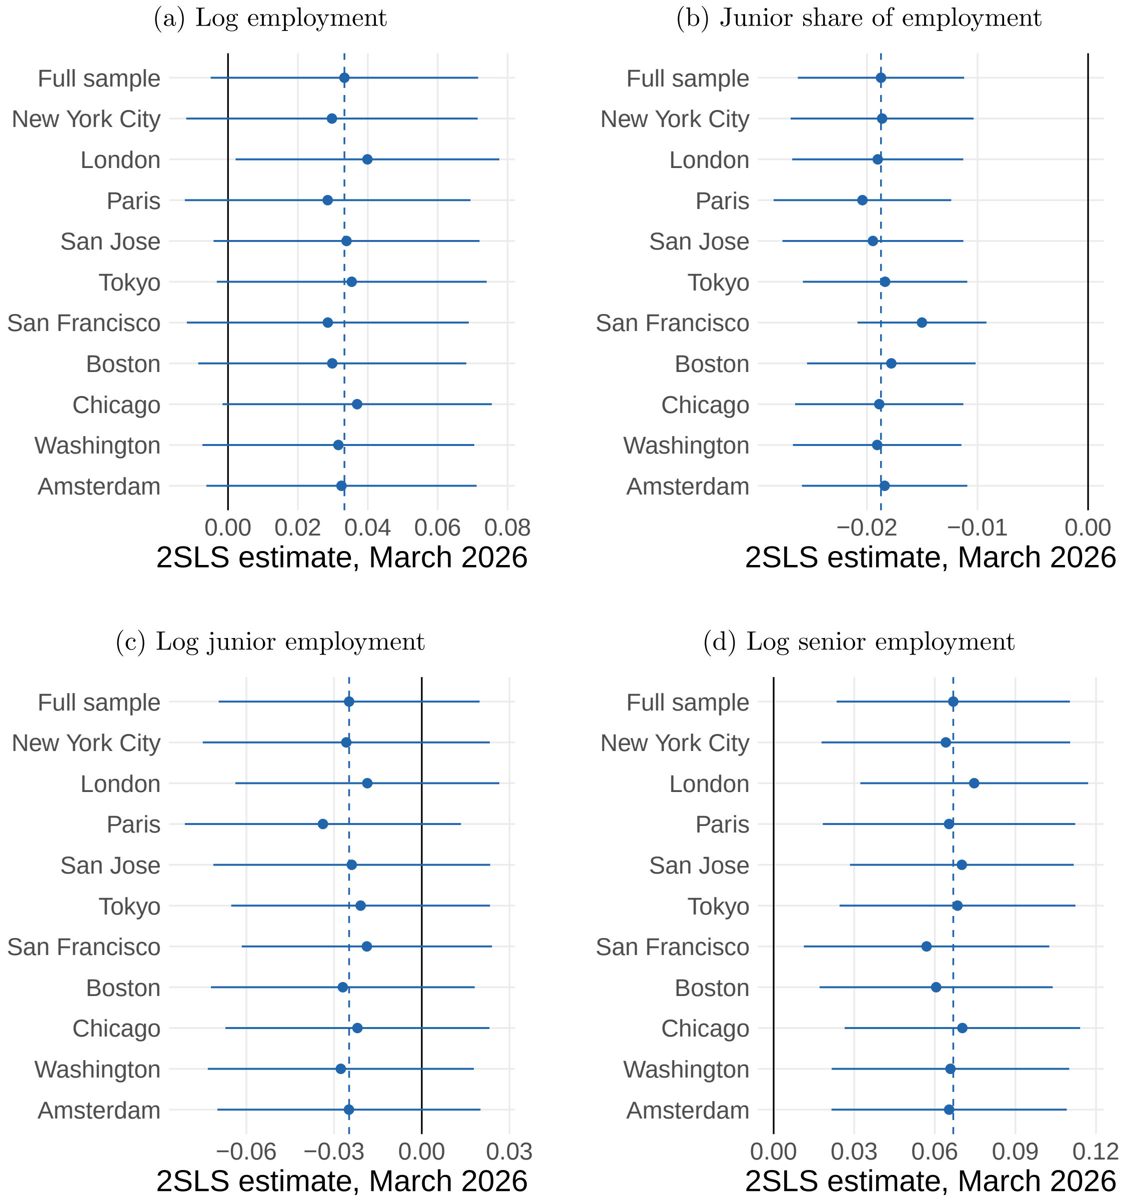

Notes: The figure asks whether any single location drives the results. Each row re-estimates the pooled two-stage least squares endpoint at March 2026 after excluding every affiliate whose parent is headquartered in one of the metropolitan areas that contribute the most treated affiliates, one metropolitan area at a time. The dashed vertical line marks the full-sample estimate and the solid vertical line marks zero; whiskers are 95% confidence intervals, clustered by the metropolitan area of the parent's headquarters. For all other definitions, see Figure 3.

Figure A33: Event study plots for the effect of AI adoption on employment and workforce composition, by predicted adoption propensity

(a) Log employment

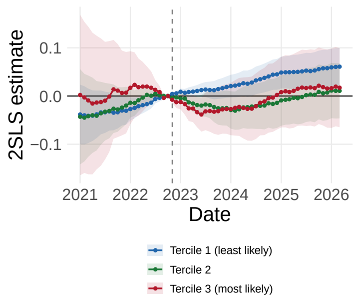

(b) Junior share of employment

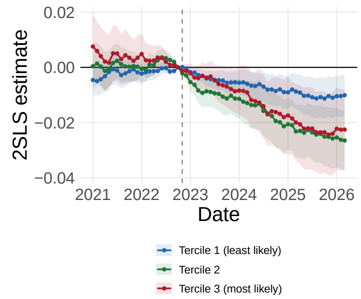

(c) Log junior employment

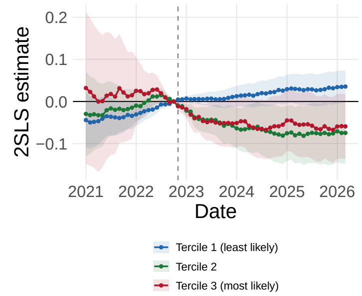

(d) Log senior employment

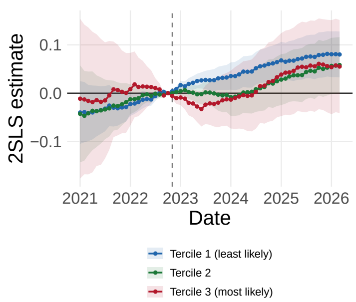

(e) High- minus low-exposure share, junior

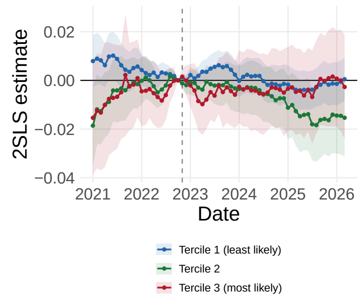

(f) High- minus low-exposure share, senior

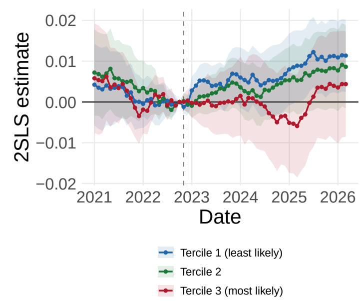

(g) High- minus low-exposure share, overall

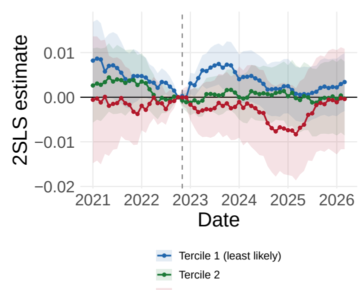

Notes: Each panel overlays the pooled event study with matched sets split by the treated affiliate's predicted probability of treatment. The propensity is the fitted value from a linear probability model of adoption on baseline log employment, average AI exposure, and junior employment share with sector and country fixed effects, estimated on the eligible cross-border universe; the cut points are terciles of the propensity distribution among adopters. Whole matched sets move with their treated affiliate, and treatment, instrument, matching, eligibility, outcomes, and clustering are held at their main-specification values. For all other definitions, see Figures 3 and 4.

A14 Matched Difference-in-Differences and Synthetic Difference-in-Differences

This appendix provides details for the matched difference-in-differences and synthetic difference-in-differences robustness checks. Both analyses keep the treatment definition, the November 2022 treatment date, the eligibility criteria of Section 3, and the outcomes of the main analysis constant. Because neither design uses the instrument, neither requires the restriction to foreign affiliates. To maintain comparability, we nevertheless estimate both designs on the foreign-affiliate sample of the main analysis first. We then extend the sample to the full universe of eligible affiliates of Section 3, domestic and foreign alike. Together, the two analyses separate differences that come from the estimation approach from differences that come from the sample. The latter is also informative about how far our main results generalize beyond foreign affiliates.

The matched difference-in-differences applies the matching procedure of Section 4 to each sample and then compares treated affiliates with their matched controls directly, without the instrument.

The synthetic difference-in-differences replaces the nearest-neighbor matches with estimated weights on control affiliates and on pre-treatment months (Arkhangelsky et al., 2021). For every country and outcome, we form a balanced affiliate-by-month panel from January 2021 through March 2026 that contains the treated affiliates and all eligible control affiliates. Affiliates with a missing outcome value in any month are dropped for that outcome, as the estimator requires a complete panel. The estimator solves

$$\hat{\tau} \;=\; \underset{\tau,\,\eta,\,\theta}{\arg\min} \;\sum_{f}\sum_{t}\bigl(y_{ft} - \eta_f - \theta_t - \tau\, D_f P_t\bigr)^2\, \hat{\omega}_f\, \hat{\lambda}_t, \tag{50}$$

where t measures calendar months, P_t indicates months from November 2022 onward, D_f indicates that affiliate f is treated, that is, that its company posted at least one AI job advertisement from November 2022 onward, and η_f and θ_t are affiliate and month fixed effects. The unit weights $\hat{\omega}_f$ and time weights $\hat{\lambda}_t$ are estimated in a preliminary step, following the regularization of Arkhangelsky et al. (2021). The unit weights, equal across treated affiliates, choose control affiliates whose weighted average tracks the average treated affiliate as closely as possible over the pre-treatment months, up to a constant. The time weights, equal across post-treatment months, select the pre-treatment months whose outcomes best predict the post-treatment average among control affiliates. The estimate $\hat{\tau}$ is the average treatment effect on treated affiliates over the 41 post-treatment months. In other words, the estimator runs the familiar difference-in-differences after reweighting affiliates and months so that treated affiliates and their synthetic control already move together before November 2022.

Table A11: Effect of AI adoption on employment and workforce composition, matched and synthetic difference-in-differences

	Log employment			Junior share	Q4−Q1 exposure share		
	Total	Junior	Senior		Total	Junior	Senior
Panel A: Matched difference-in-differences							
AI adoption	0.079***	0.045***	0.106***	−0.013***	0.004***	0.001	0.009***
	(0.002)	(0.002)	(0.003)	(0.0004)	(0.0003)	(0.001)	(0.0004)
Affiliate × match FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Match × month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Countries	41	41	41	41	41	41	41
Panel B: Synthetic difference-in-differences							
AI adoption	0.027***	−0.011	0.043***	−0.013***	0.001	−0.002***	0.006***
	(0.007)	(0.007)	(0.009)	(0.001)	(0.001)	(0.001)	(0.001)
Countries	41	41	41	41	41	41	41

Notes: The table reports the estimated effect of AI adoption on each column outcome under two identification strategies that do not use the instrument, both estimated on the full eligible universe of affiliates, domestic and foreign alike; an affiliate is treated if its company posted at least one AI job advertisement after November 2022. Both panels report the March 2026 endpoint, the same month as the two-stage least squares estimates. Panel A reports the endpoint of the pooled matched difference-in-differences event study, which compares adopting affiliates with their matched controls, with standard errors clustered by ultimate parent company. Panel B reports the endpoint of the synthetic difference-in-differences gap trajectory of Equation (50): country trajectories are pooled month by month by random-effects meta-analysis, and each country's month-specific standard error is a leave-one-parent-out jackknife that holds the estimated unit and time weights fixed. Appendix A14 gives full estimation details for both designs. For the outcome definitions and the corresponding two-stage least squares estimates, see Table 3. */**/***: $p < 0.10/0.05/0.01$.

For inference, we use a leave-one-out jackknife at the parent level. We drop one corporate group at a time, all of its affiliates together, recompute the estimate with the weights held fixed and the remaining unit weights renormalized, and take the jackknife variance across the resulting estimates. This mirrors the parent-level clustering of the matched design.

Figure A34: Event study plots for the effect of AI adoption on employment and workforce composition, matched difference-in-differences

(a) Log employment

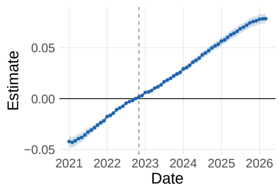

(b) Junior share of employment

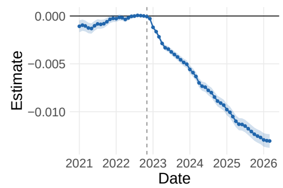

(c) Log junior employment

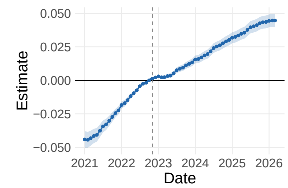

(d) Log senior employment

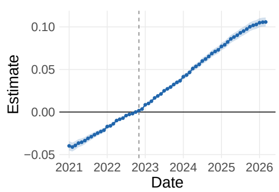

(e) High- minus low-exposure share, junior

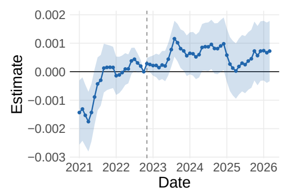

(f) High- minus low-exposure share, senior

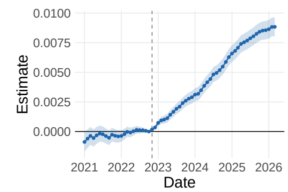

(g) High- minus low-exposure share, overall

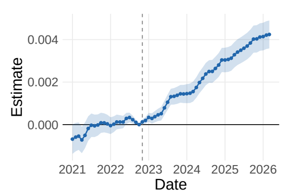

Notes: The figure plots the matched difference-in-differences event study, which compares treated affiliates with their matched controls directly rather than instrumenting treatment. Each point is the estimated gap between treated affiliates and their controls in that month, relative to October 2022, estimated by weighted least squares with affiliate-by-matched-set and matched-set-by-month fixed effects, so that every comparison stays within a matched set and calendar month. The sample is the full eligible universe of affiliates of Section 3, domestic and foreign alike; an affiliate is treated if its company posted at least one AI job advertisement from November 2022 onward, and each is matched within country and sector to control affiliates on baseline employment, average AI exposure, and junior employment share, as described in Section 4. Panels A–D plot log affiliate employment, the junior share of employment, log junior employment, and log senior employment, where junior covers Revelio seniority levels 1–2 and senior levels 3–7. Panels E–G plot the difference between the share of employment in occupations in the top quartile of AI exposure and the share in the bottom quartile, within junior employment, within senior employment, and across the whole workforce. Shaded bands are 95% confidence intervals, clustered by parent. The dashed vertical line marks ChatGPT's launch in November 2022.

Figure A35: Event study plots for the effect of AI adoption on employment and workforce composition, matched difference-in-differences on the foreign-affiliate sample

(a) Log employment

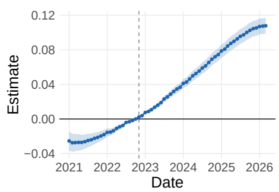

(b) Junior share of employment

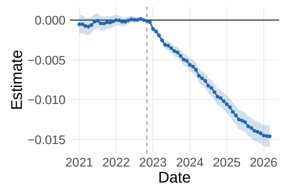

(c) Log junior employment

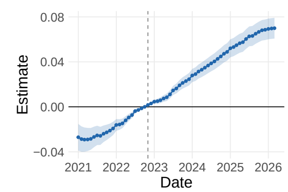

(d) Log senior employment

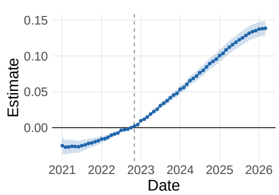

(e) High- minus low-exposure share, junior

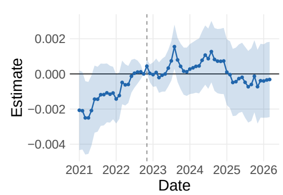

(f) High- minus low-exposure share, senior

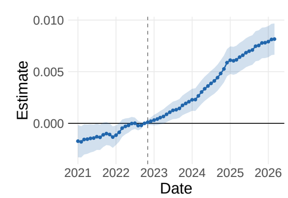

(g) High- minus low-exposure share, overall

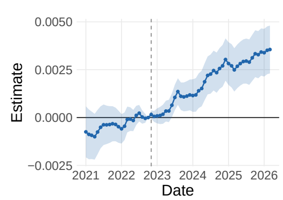

Notes: The figure repeats Figure A34 on the foreign-affiliate sample of the two-stage least squares analysis, namely the affiliates whose parent is headquartered in another country and for which the instrument is defined. For all other definitions, see Figure A34.

Figure A36: Event study plots for the effect of AI adoption on employment and workforce composition, synthetic difference-in-differences

(a) Log employment

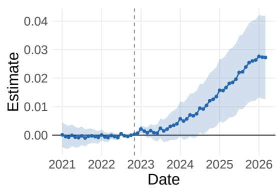

(b) Junior share of employment

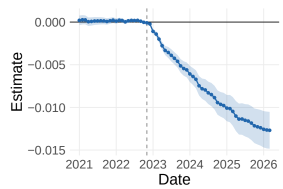

(c) Log junior employment

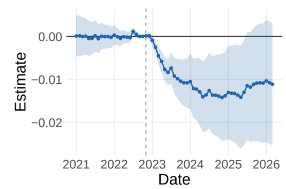

(d) Log senior employment

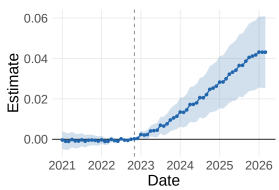

(e) High- minus low-exposure share, junior

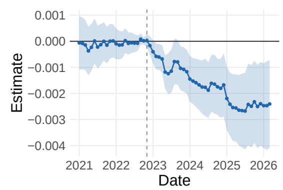

(f) High- minus low-exposure share, senior

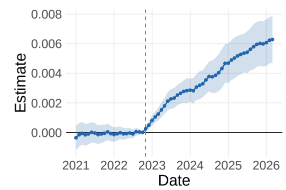

(g) High- minus low-exposure share, overall

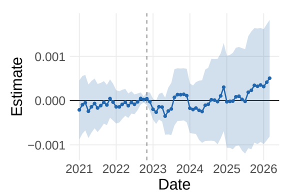

Notes: The figure plots the synthetic difference-in-differences design of Arkhangelsky et al. (2021), which replaces nearest-neighbor matching with estimated weights on control affiliates and on pre-treatment months. Each point is the gap between treated affiliates and their synthetic control in that month, net of the time-weighted pre-treatment gap, so that the series is centered on zero before November 2022. The sample is the full eligible universe of affiliates of Section 3, domestic and foreign alike; an affiliate is treated if its company posted at least one AI job advertisement from November 2022 onward, and controls are affiliates in corporate groups that never posted one. Panels A–D plot log affiliate employment, the junior share of employment, log junior employment, and log senior employment, where junior covers Revelio seniority levels 1–2 and senior levels 3–7. Panels E–G plot the same gap in the difference between the share of employment in occupations in the top quartile of AI exposure and the share in the bottom quartile, within junior employment, within senior employment, and across the whole workforce. Country trajectories are pooled month by month by random-effects meta-analysis. Shaded bands are 95 percent confidence intervals from each country's leave-one-parent-out jackknife of the full trajectory, which holds the estimated unit and time weights fixed at their full-sample values. Because the unit weights are chosen to fit the pre-treatment months, the pre-treatment bands describe the quality of that fit rather than a test of parallel trends. The dashed vertical line marks ChatGPT's launch in November 2022.

Figure A37: Event study plots for the effect of AI adoption on employment and workforce composition, synthetic difference-in-differences on the foreign-affiliate sample

(a) Log employment

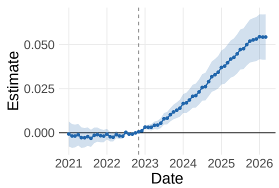

(b) Junior share of employment

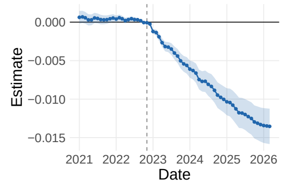

(c) Log junior employment

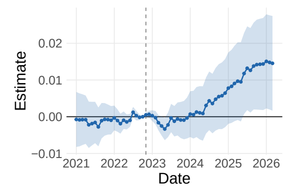

(d) Log senior employment

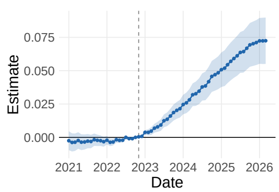

(e) High- minus low-exposure share, junior

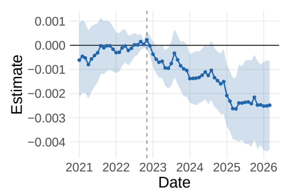

(f) High- minus low-exposure share, senior

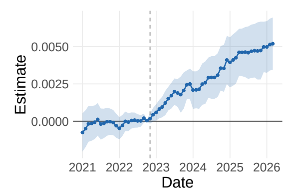

(g) High- minus low-exposure share, overall

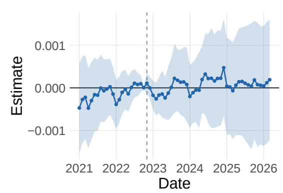

Notes: The figure repeats Figure A36 on the foreign-affiliate sample of the two-stage least squares analysis, namely the affiliates whose parent is headquartered in another country and for which the instrument is defined. Bands and pooling are as in Figure A36. For all other definitions, see Figure A36.

A15 Additional Robustness Figures

This appendix collects the figures behind the remaining checks of Section 5.3.2, in the order in which they are described there (Figures A38–A43).

Figure A38: Event study plots for the effect of AI adoption on employment and workforce composition, adoption intensity

(a) Log employment

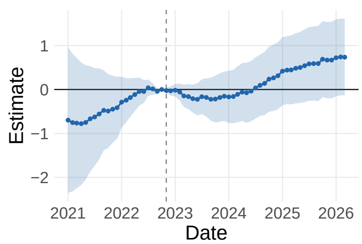

(b) Junior share of employment

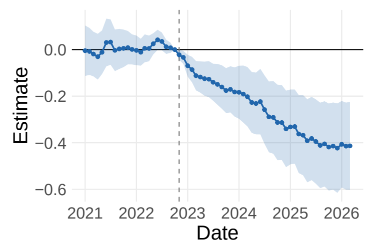

(c) Log junior employment

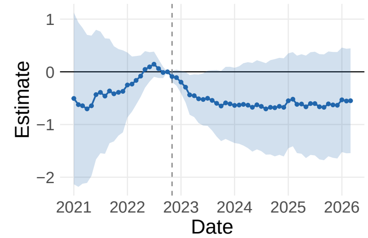

(d) Log senior employment

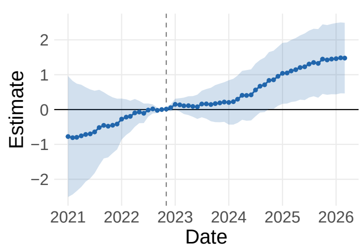

(e) High- minus low-exposure share, junior

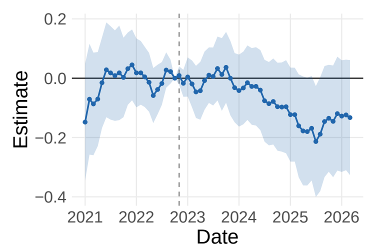

(f) High- minus low-exposure share, senior

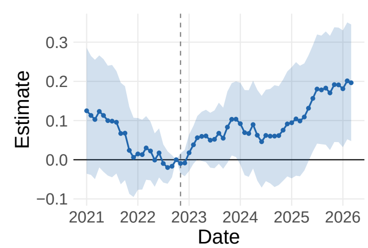

(g) High- minus low-exposure share, overall

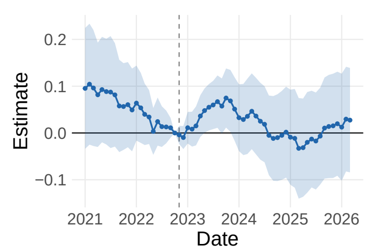

Notes: Each panel replaces the binary adoption indicator of the main design with a continuous measure of how heavily the company hires for AI work, so that the coefficient reads as the effect of a one-unit increase in the log dose. The dose is the log of one plus the ratio of the AI job advertisements the company posted anywhere in the world from November 2022 onward to all the job advertisements it posted anywhere in the world before ChatGPT's launch, from the start of the listings data through October 2022. The denominator is fixed before treatment begins and so cannot respond to adoption itself, and the log compresses the heavy right tail of AI-native companies that posted little before the launch. The dose is measured for the company as a whole and assigned identically to each of its affiliates; matched controls carry a dose of zero. The dose is instrumented by the same standardized peer adoption rate as in the main design, and the matched sample, fixed effects, and clustering are unchanged. For all other definitions, see Figures 3 and 4.

Figure A39: Event study plots for the effect of AI adoption on employment and workforce composition, same-category control postings

(a) Log employment

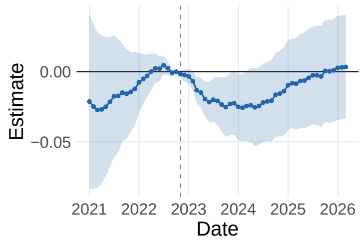

(b) Junior share of employment

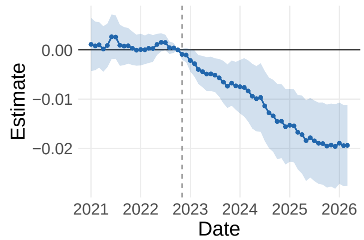

(c) Log junior employment

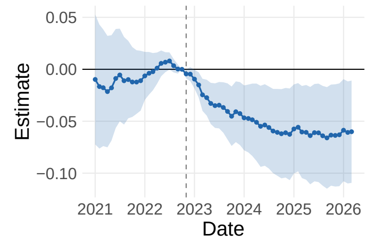

(d) Log senior employment

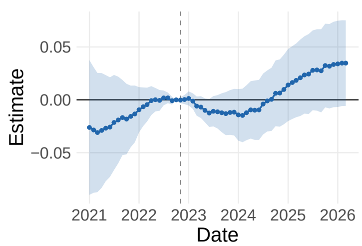

(e) High- minus low-exposure share, junior

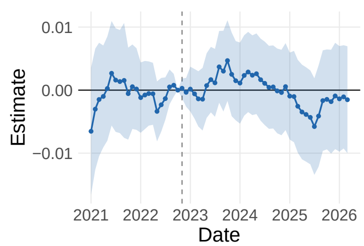

(f) High- minus low-exposure share, senior

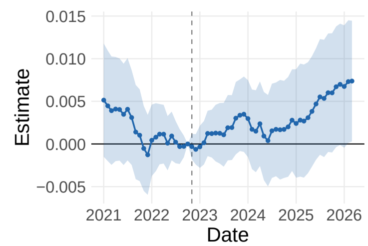

(g) High- minus low-exposure share, overall

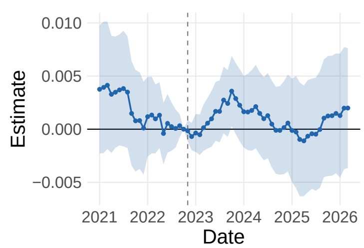

Notes: Each panel re-estimates the pooled event study after restricting the control pool to affiliates whose companies posted at least one job advertisement in the affiliate's country from November 2022 onward in the two-digit occupation category of the treated affiliate's company's first classifiable AI posting anywhere in our data. Because an advertised AI job is itself a posting in its occupation category, the restriction ensures that controls are observably hiring in that category as well. Treated affiliates are re-matched against the restricted pool; treatment, instrument, eligibility, outcomes, and clustering are held at their main-specification values. For all other definitions, see Figures 3 and 4.

Figure A40: Event study plots for the effect of AI adoption on employment and workforce composition, by year of the company's first AI advertisement

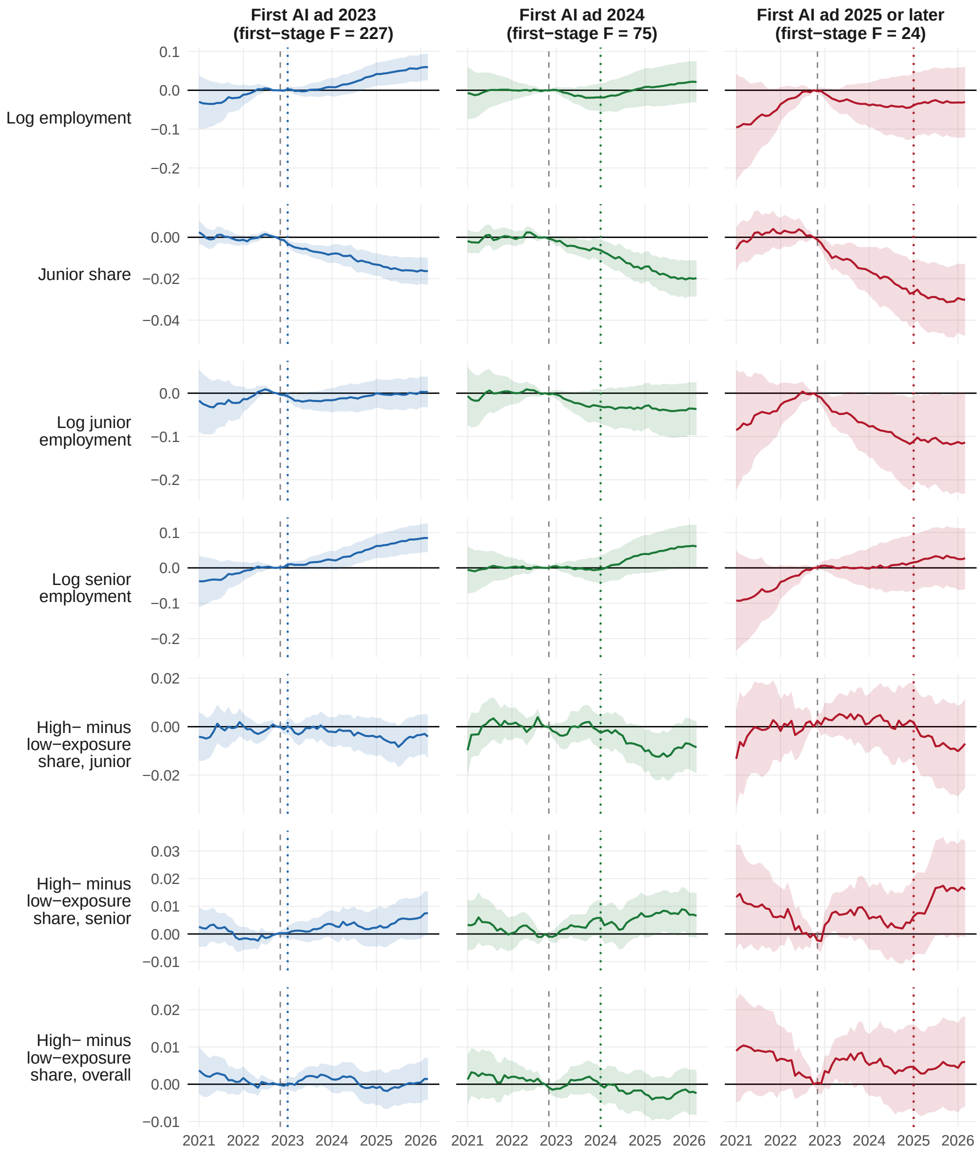

Notes: Each column reports the pooled event study for affiliates whose company's first AI advertisement appeared in 2023, in 2024, or in 2025 or later, and each row reports one outcome, with a common vertical scale across the three cohorts. Each cohort retains its matched never-adopter controls; whole matched sets follow the treated affiliate. Treatment, instrument, matching, eligibility, outcomes, and clustering follow the main specification. The dashed line marks ChatGPT's launch; the dotted line marks 1 January of the cohort's first-advertisement year. The last three rows show the difference between employment shares in the highest and lowest exposure quartiles for juniors, seniors, and all workers. Shaded bands are 95% confidence intervals. For all other definitions, see Figures 3 and 4.

Figure A41: Event study plots for the effect of AI adoption on employment and workforce composition, early versus late adopters

Notes: Each panel re-estimates the pooled event study on affiliates of adopting companies alone. Treated affiliates are those whose companies' first AI job advertisement appeared between November 2022 and December 2024, and controls are drawn exclusively from affiliates of late-adopting companies, whose first AI job advertisement appeared in January 2025 or later, so that both groups select into adoption and only the timing differs. Affiliates of early-adopting companies are re-matched against the late-adopter pool; matching, instrument, eligibility, outcomes, and clustering are held at their main-specification values, and the event clock stays centered on ChatGPT's launch. Because late adopters begin adopting in January 2025, coefficients from that month onward compare earlier with later adoption rather than adoption with none. For all other definitions, see Figures 3 and 4.

Figure A42: Event study plots for the effect of AI adoption on employment and workforce composition, placebo launch date

Notes: Each panel re-estimates the pooled event study under a placebo launch date of January 2022. Sample, treatment, matching, eligibility, instrument, outcomes, and clustering are held at their main-specification values, and only the event clock moves: the baseline month becomes December 2021, and the months from ChatGPT's actual launch onward are dropped, so the coefficients to the right of the dashed line cover only the months that separate the placebo date from the actual launch. The dashed vertical line marks the placebo launch in January 2022 and the dotted vertical line marks ChatGPT's actual launch in November 2022. For all other definitions, see Figures 3 and 4.

Figure A43: Event study plots for the effect of AI adoption on employment and workforce composition, no posting requirement for controls

(a) Log employment

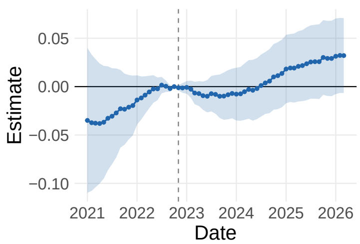

(b) Junior share of employment

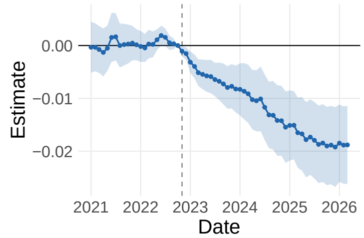

(c) Log junior employment

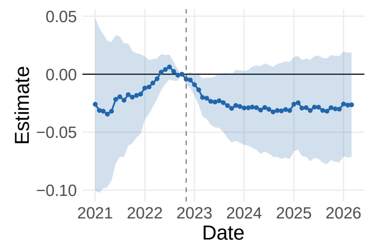

(d) Log senior employment

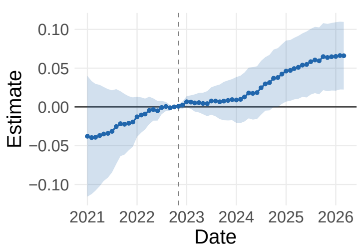

(e) High- minus low-exposure share, junior

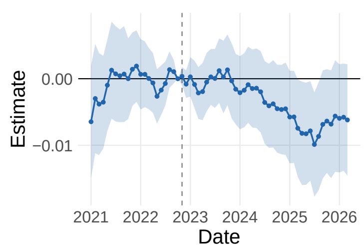

(f) High- minus low-exposure share, senior

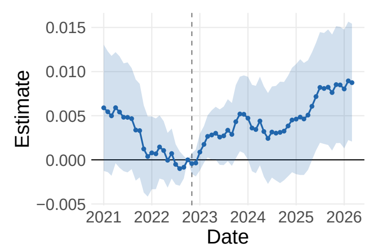

(g) High- minus low-exposure share, overall

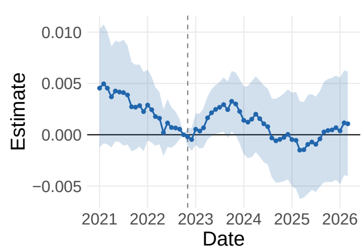

Notes: Each panel re-estimates the pooled event study after dropping the requirement that each control affiliate's company post at least one job advertisement anywhere in our data from November 2022 onward, a condition that would otherwise select the control pool on post-launch hiring activity. Treated affiliates are re-matched against the enlarged control pool; treatment, instrument, eligibility, outcomes, and clustering are held at their main-specification values. For all other definitions, see Figures 3 and 4.

Figure A44: Effect of AI adoption on affiliate-level occupation shares, peer growth control

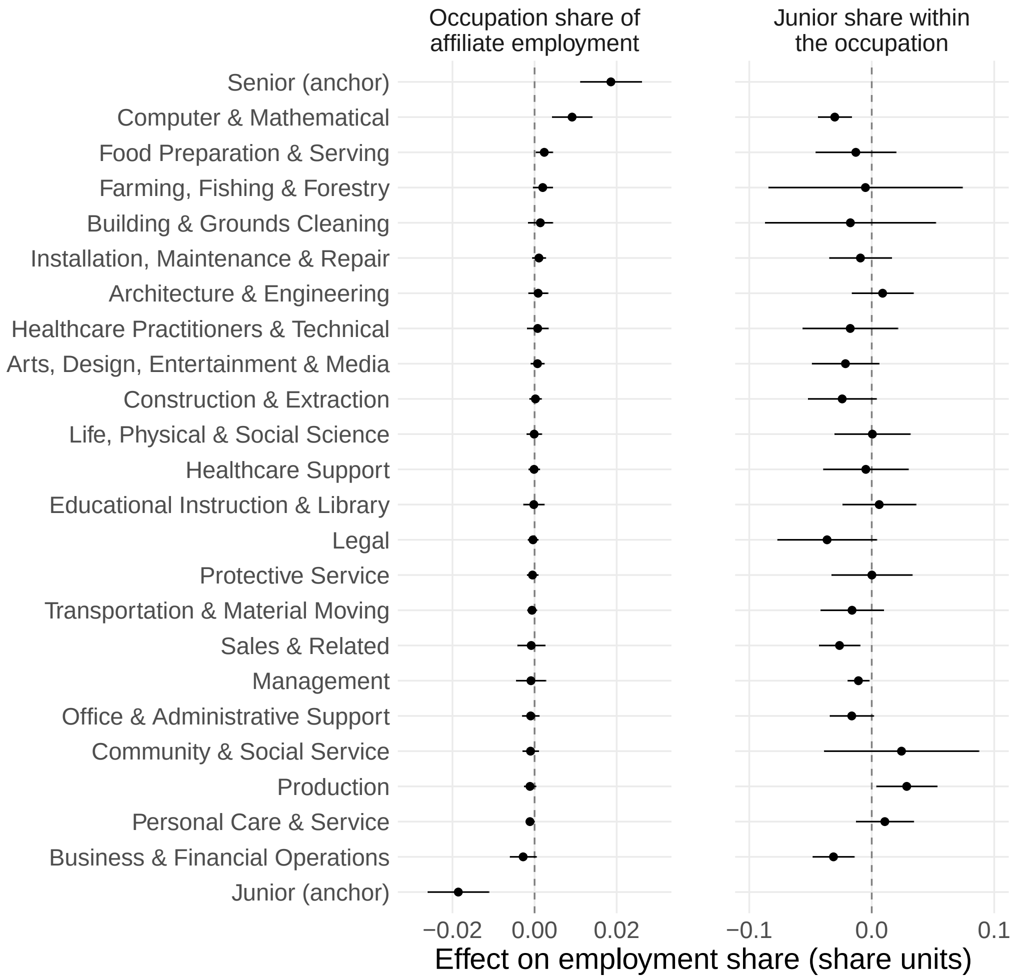

Notes: The figure repeats Figure 5 adding to both stages the leave-one-parent-out mean employment growth of multinational parents headquartered in the same metropolitan area, the peer-growth control of Figure A28. Occupations whose first stage is too weak under this specification are omitted. For all other definitions, see Figure 5.

A16 Occupation-Level Results for the Robustness Checks

This appendix repeats the occupation-level analysis of Figure 5 for each robustness check of Section 5.3, in the order in which the checks are described there (Figures A44–A54).

Figure A45: Effect of AI adoption on affiliate-level occupation shares, country × sector instrument

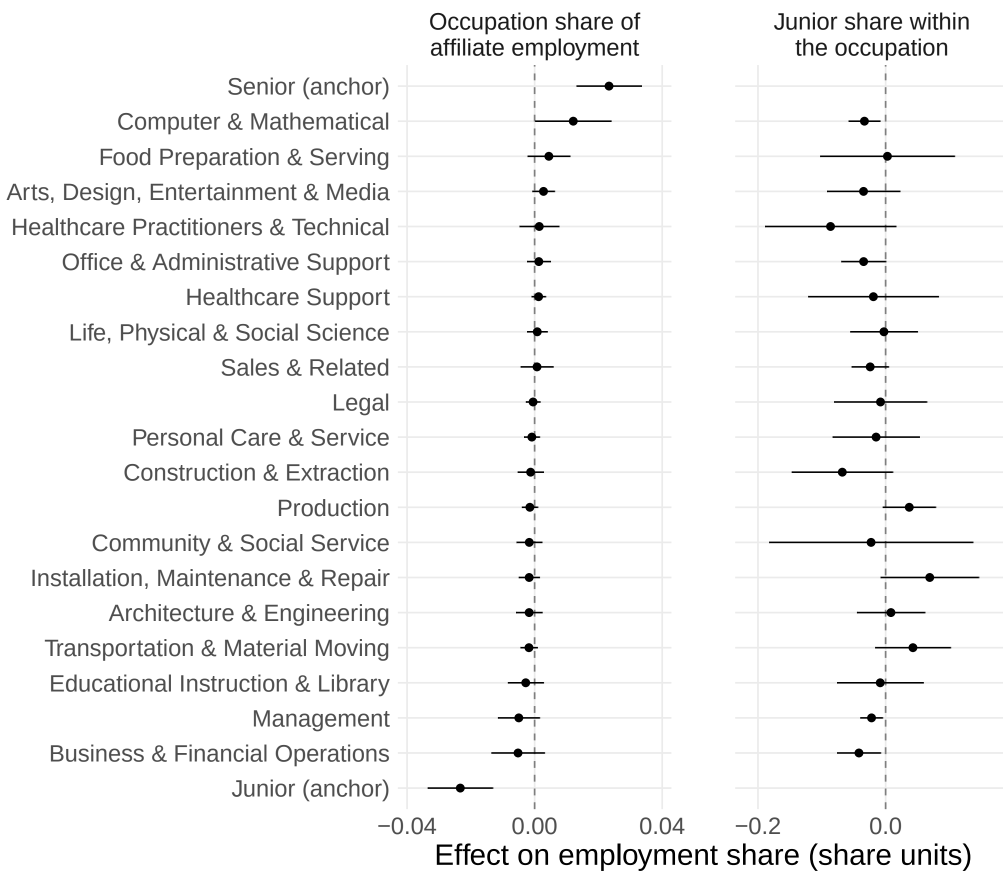

Notes: The figure repeats Figure 5 with adoption instrumented by the leave-out AI adoption rate among multinational parents in the same headquarters country and sector as the focal parent, rather than in the same headquarters metropolitan area, on the same matched occupation samples. Standard errors are clustered at the country-by-sector level, and occupations whose first stage is too weak under this instrument are omitted. For all other definitions, see Figure 5.

Figure A46: Effect of AI adoption on affiliate-level occupation shares, sector-purged instrument

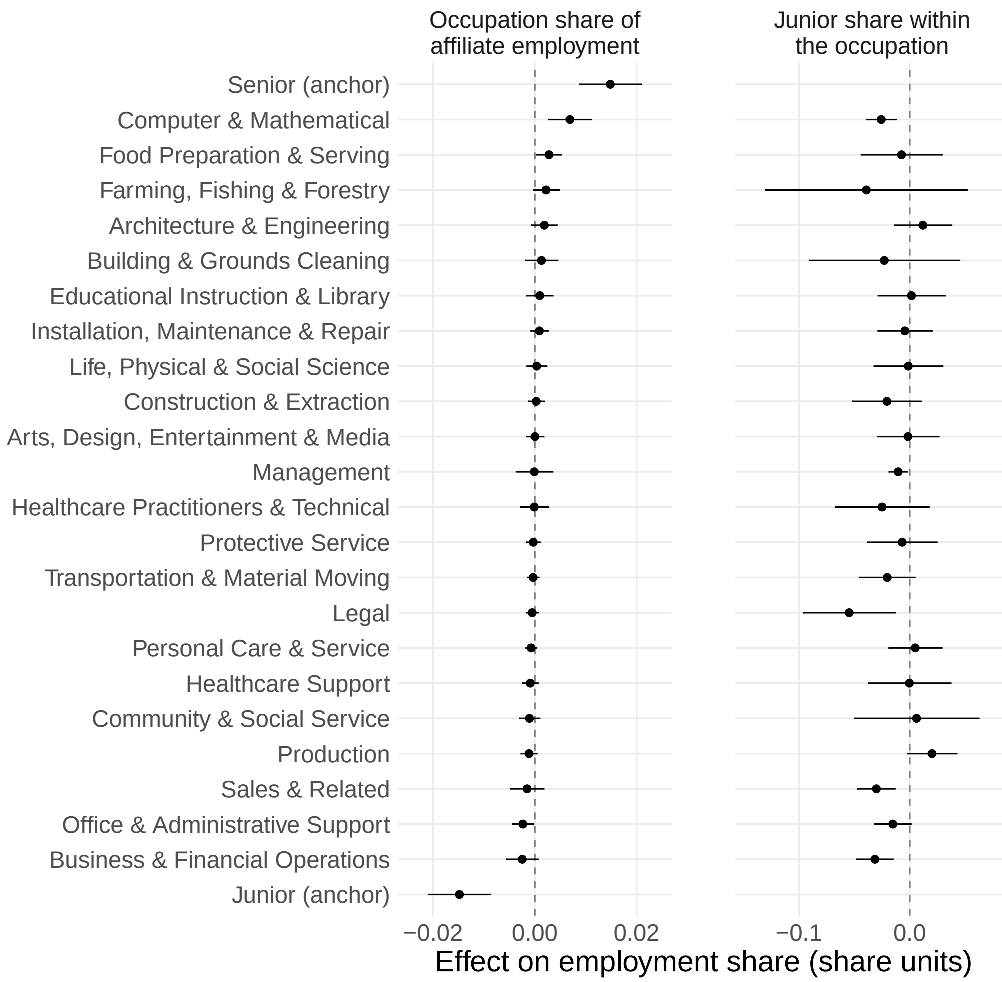

Notes: The figure repeats Figure 5 with the instrument rebuilt from metro peer parents *outside* the focal parent's own sector, the sector-purged instrument of Figure A28. Occupations whose first stage is too weak under this instrument are omitted. For all other definitions, see Figure 5.

Figure A47: Effect of AI adoption on affiliate-level occupation shares, leaving out one metropolitan area at a time

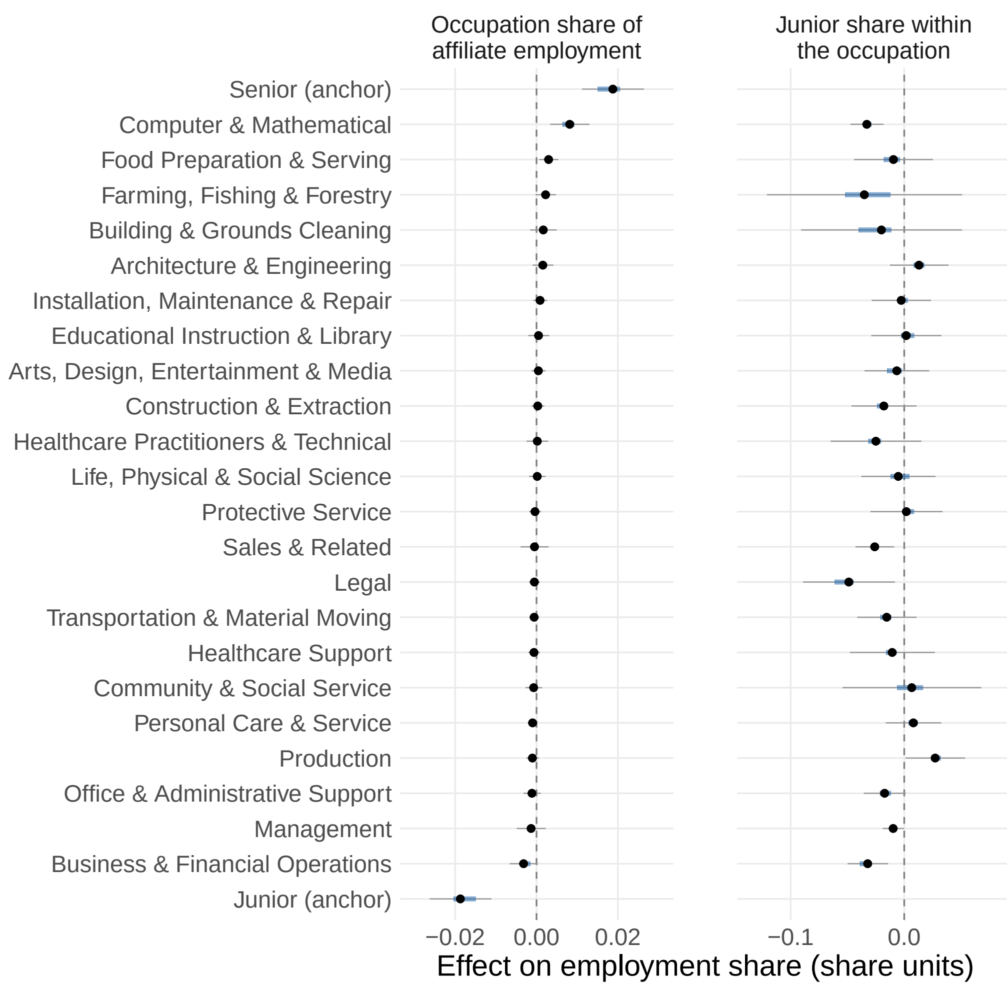

Notes: For each occupation, the point marks the full-sample endpoint estimate of Figure 5, the thin gray whisker its 95% confidence interval, and the thick shaded band the range of the estimates obtained by excluding, one at a time, each of the metropolitan areas that contribute the most treated affiliates, as in Figure A32. Occupations whose full-sample first stage is too weak are omitted. For all other definitions, see Figure 5.

Figure A48: Effect of AI adoption on affiliate-level occupation shares, matched difference-in-differences

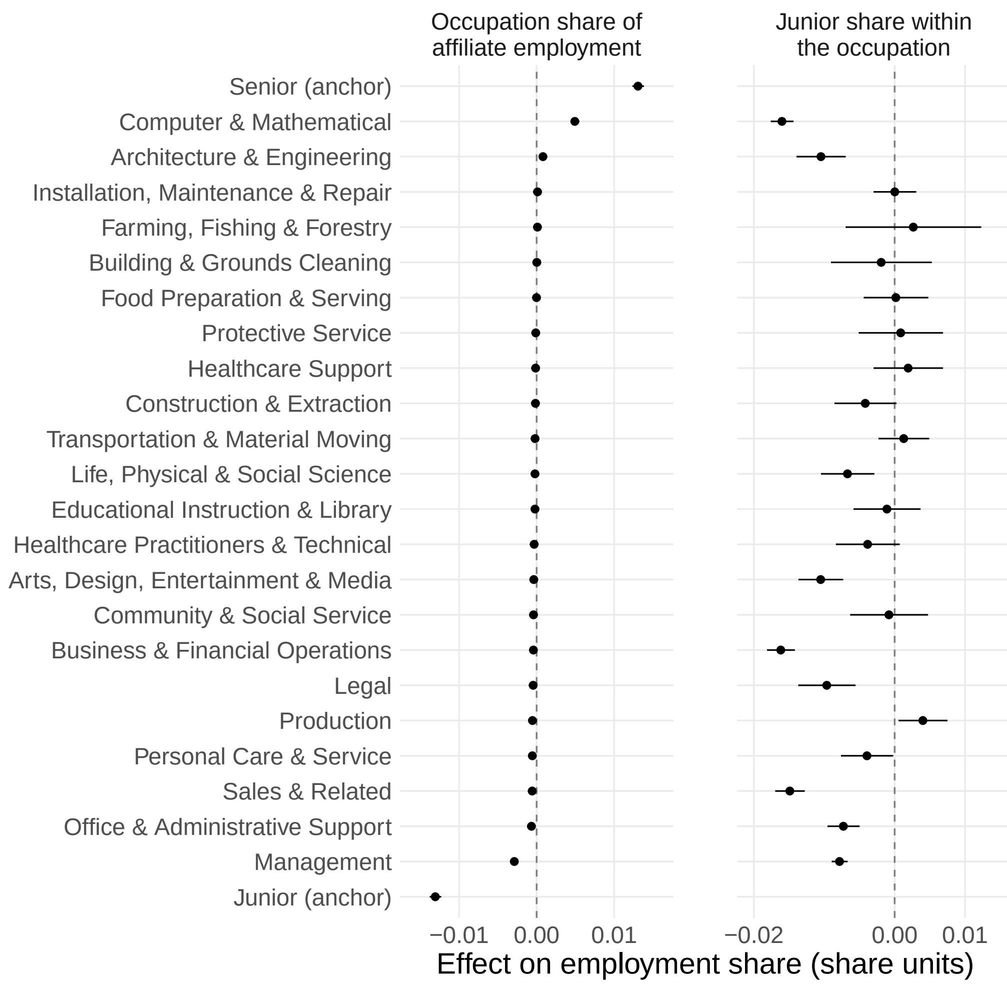

Notes: The figure repeats Figure 5 under the matched difference-in-differences design of Figure A34, which compares adopters with their matched controls directly rather than instrumenting adoption and is estimated on the full eligible universe of affiliates. For all other definitions, see Figure 5.

Figure A49: Effect of AI adoption on affiliate-level occupation shares, matched difference-in-differences on the foreign-affiliate sample

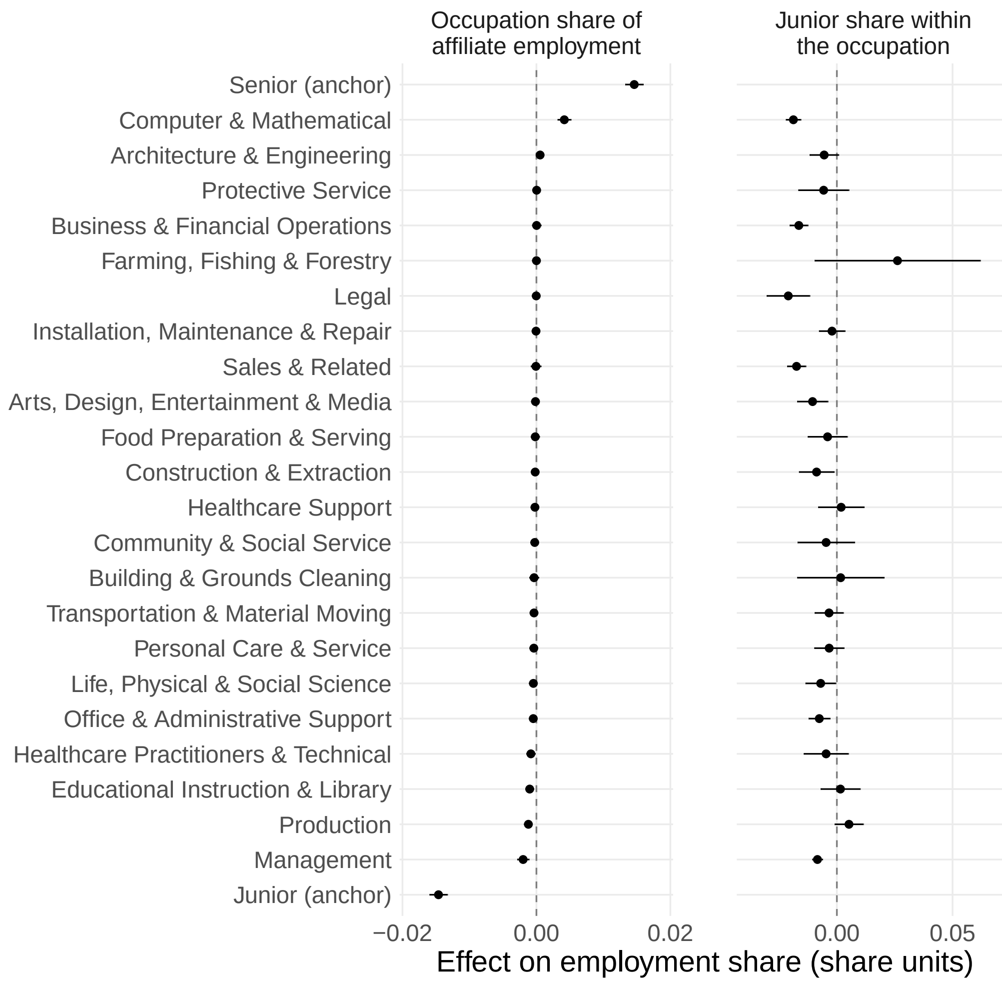

Notes: The figure repeats Figure A48 on the foreign-affiliate sample of the two-stage least squares analysis, namely the affiliates whose parent is headquartered in another country and for which the instrument is defined, the sample of Figure A35. For all other definitions, see Figure 5.

Figure A50: Effect of AI adoption on affiliate-level occupation shares, synthetic difference-in-differences

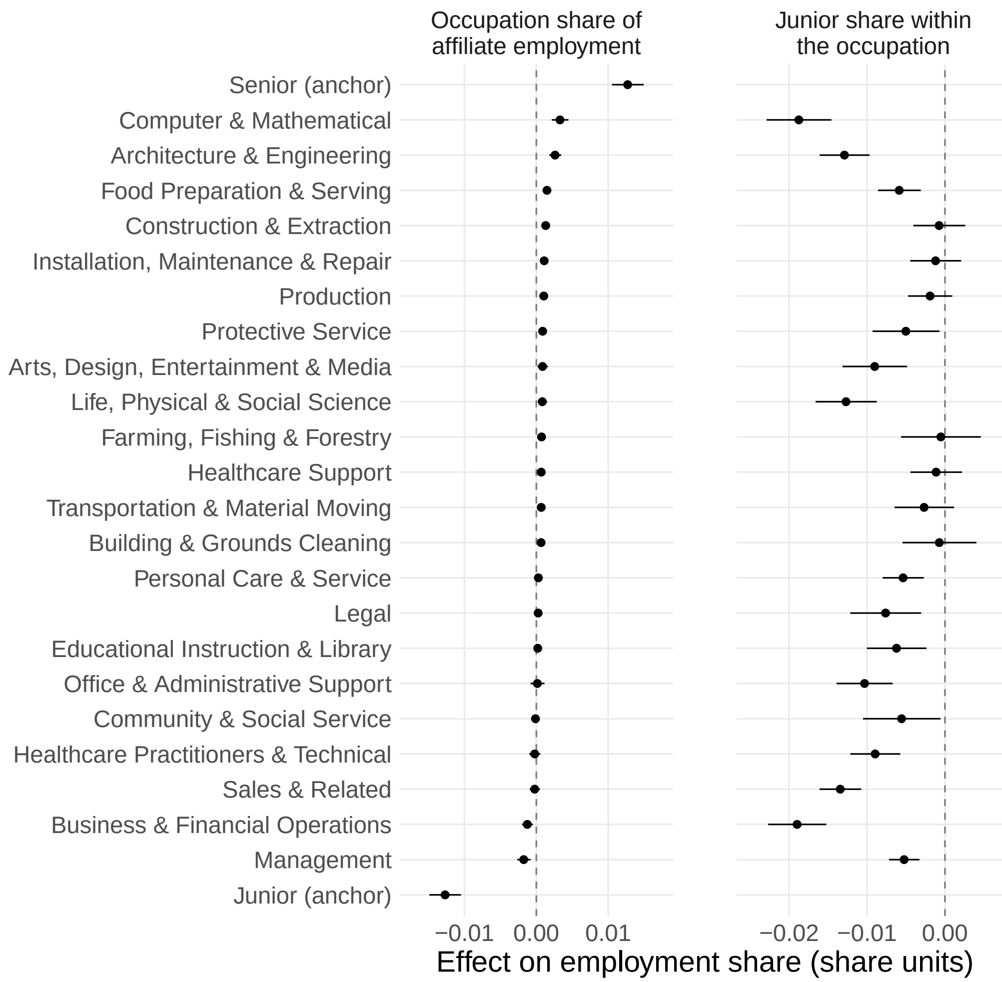

Notes: The figure repeats Figure 5 under the synthetic difference-in-differences design of Figure A36 (Arkhangelsky et al., 2021), estimated on the full eligible universe of affiliates. Each estimate is the March 2026 endpoint of the pooled gap trajectory for that occupation, with country trajectories pooled month by month by random-effects meta-analysis and standard errors from the leave-one-parent-out jackknife of Figure A36. For all other definitions, see Figure 5.

Figure A51: Effect of AI adoption on affiliate-level occupation shares, synthetic difference-in-differences on the foreign-affiliate sample

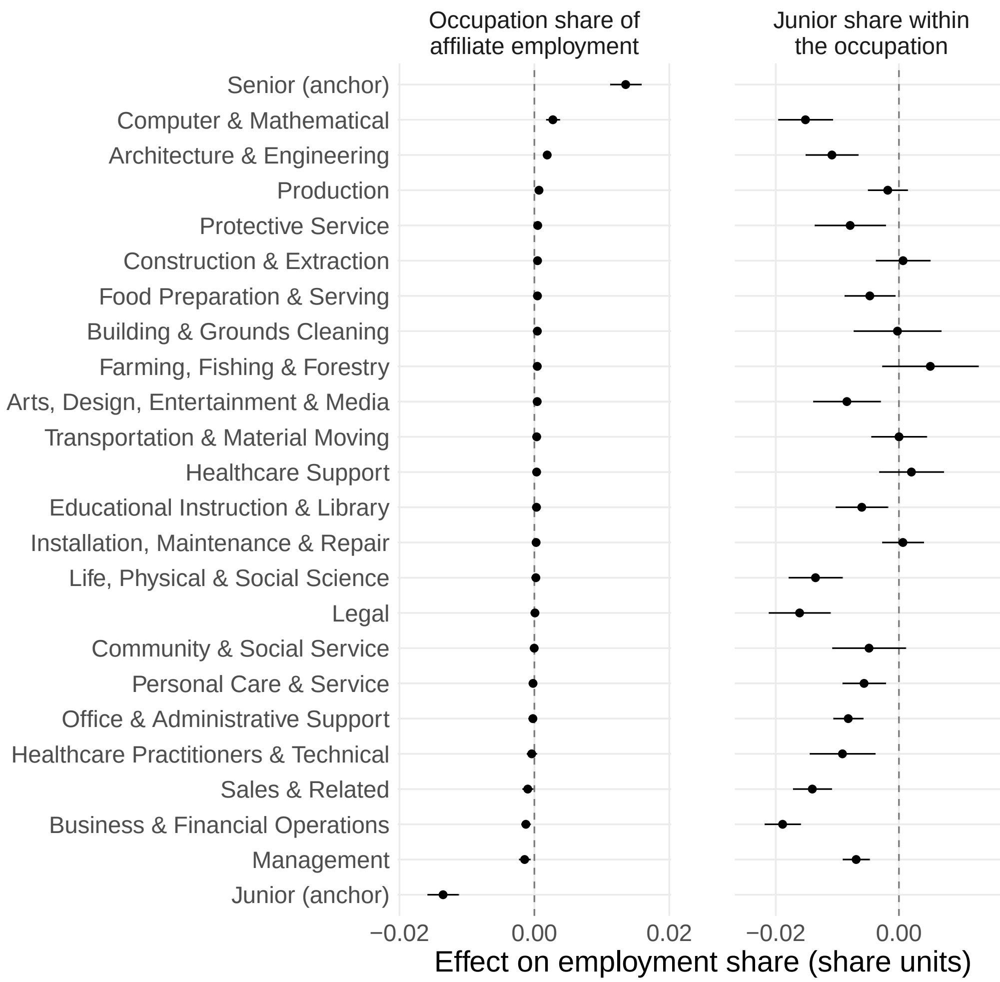

Notes: The figure repeats Figure A50 on the foreign-affiliate sample of the two-stage least squares analysis, namely the affiliates whose parent is headquartered in another country and for which the instrument is defined, the sample of Figure A37. For all other definitions, see Figure 5.

Figure A52: Effect of AI adoption on affiliate-level occupation shares, same-category control postings

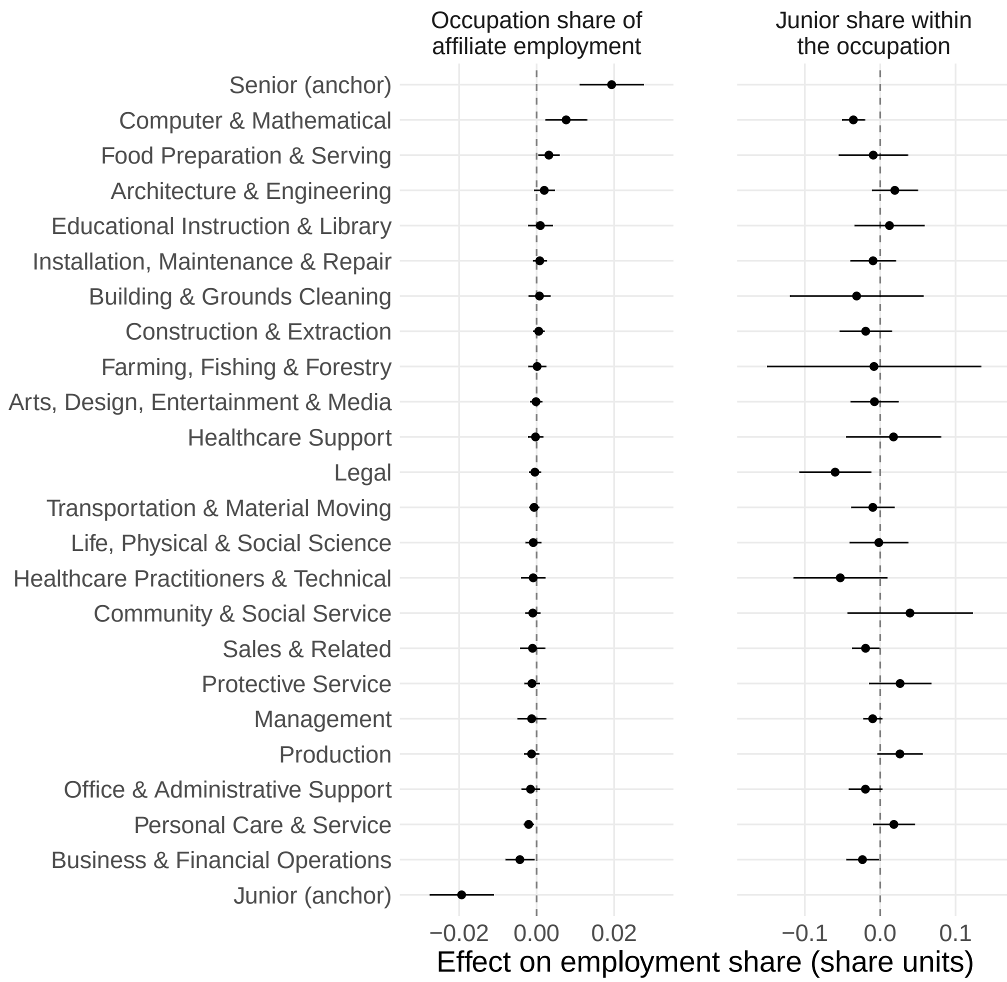

Notes: The figure repeats Figure 5 on occupation samples re-matched after restricting the control pool to affiliates whose companies posted at least one job advertisement in the affiliate's country from November 2022 onward in the occupation category of the treated affiliate's company's first classifiable AI posting anywhere in our data, the sample-construction check of Figure A39. For all other definitions, see Figure 5.

Figure A53: Effect of AI adoption on affiliate-level occupation shares, placebo launch date

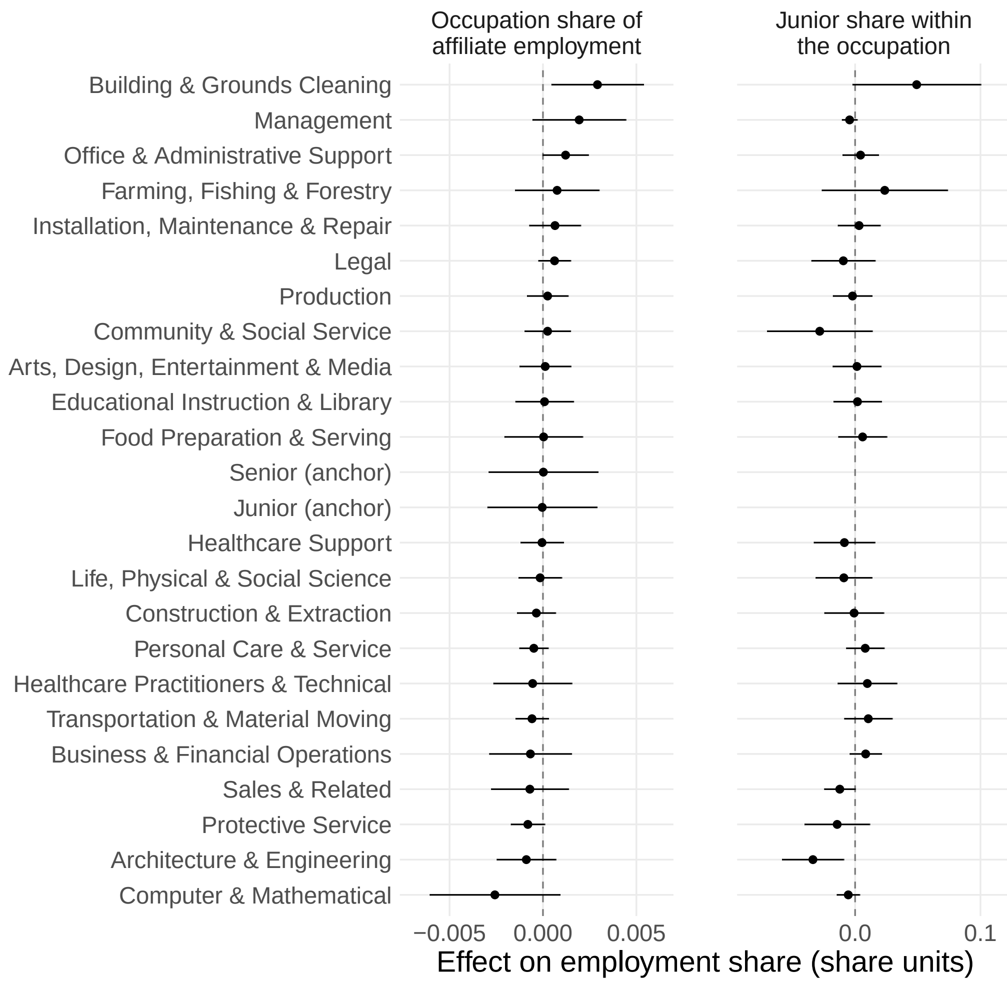

Notes: The figure repeats Figure 5 under the placebo launch date of January 2022 used in Figure A42: each effect is the change from the December 2021 placebo baseline to October 2022, the month before ChatGPT's actual launch, so that the real post-launch period never enters. Sample, matching, instrument, and clustering are held at their main-specification values, and occupations whose first stage is too weak are omitted. For all other definitions, see Figure 5.

Figure A54: Effect of AI adoption on affiliate-level occupation shares, no posting requirement for controls

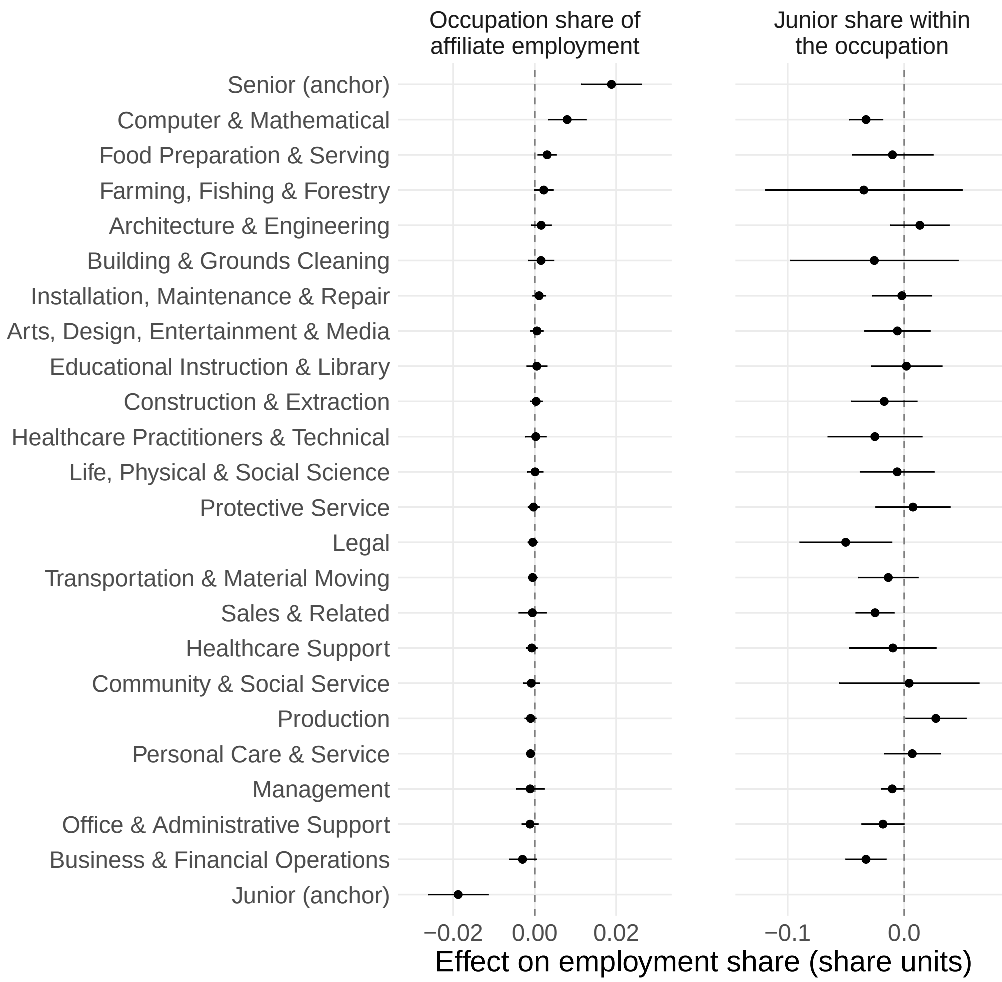

Notes: The figure repeats Figure 5 on occupation samples re-matched without requiring each control affiliate's company to post a job advertisement anywhere in our data from November 2022 onward, the sample-construction check of Figure A43. Occupations whose first stage is too weak in this sample are omitted. For all other definitions, see Figure 5.