AI Is Already Eroding Wages: Quasi-Experimental Evidence From Occupational Exposure*

José Azar† Mireia Giné‡ Javier Sanz-Espín§

December 3, 2025

Abstract

We provide causal evidence on how generative AI is reshaping the internal wage structure of the firm. Using a large-scale dataset of realized worker-firm matches for 138 million U.S. workers, we document that the arrival of ChatGPT significantly reduced the wages of firms whose ex-ante occupation structure was more exposed to automation relative to less exposed firms. In particular, we find that the average wage of exposed firms fell 4.5% relative to unexposed firms. We also find that starting wages for new positions for junior and mid-level workers in exposed firms fell by 6.3% and 5.9% respectively relative to unexposed firms, while senior wages remained stable or even increased. We show that firms in exposed sectors reduced the share of junior employees, and increased the share of mid-level employees. Finally, we find that firms exposed to AI experienced a reduction in the level of education of new positions for junior and mid-level workers, relative to less exposed firms.

Keywords: AI, labor markets, wages

*We thank Jordi Gual, Enrique Ide, Joan Monrás, Eduard Talamás, Xavier Vives, and seminar participants at IESE Business School for helpful comments. Sanz-Espín acknowledges financial support from the Asociación de Amigos de la Universidad de Navarra and from the Spanish Ministry of Science and Innovation through the FPU fellowship.

†University of Navarra, IESE Business School, CEPR and ECGI. Email: jazar@iese.edu.

‡IESE Business School, CEPR, ECGI, and WRDS. Email: mgine@iese.edu.

§Toulouse School of Management and University of Navarra. Email: javier.sanz-espin@tsm-education.fr.

1 Introduction

The rapid diffusion of large language models (LLMs) has revived a central question in labor economics: how does technological change reshape the labor market? While earlier waves of automation primarily displaced routine middle-skill occupations (Katz and Murphy, 1992; Autor et al., 2003; Acemoglu and Autor, 2011), LLMs possess a distinct capability: they automate standardized cognitive and reasoning tasks–drafting, summarizing, coding, and basic analysis. In professional settings, these tasks are disproportionately concentrated in junior roles. This distinction raises the prospect that generative AI acts as a form of seniority-biased technological change, potentially eroding entry-level opportunities, compressing wages at the bottom of the distribution, and weakening the apprenticeship model through which workers progress into senior positions (Hosseini and Lichtinger, 2025; Brynjolfsson et al., 2025).

The impact of generative AI on firms' wage structure is still an open question. To contribute to this debate, we analyze a high-frequency dataset of roughly 450 million U.S. positions (corresponding to 138 million workers)–a scale that permits a granular view of the firm's internal labor market. Constructed from Revelio Technologies data, our sample captures anonymized employment histories with salary information derived from 200 million salary observations from job vacancies, publicly available labor certification applications, as well as self-reported salary platforms. Revelio's detailed information on each position's job title and seniority allows us to group workers into three seniority categories: juniors, mid-level workers, and seniors. Our identification strategy exploits the sudden arrival of ChatGPT in November 30, 2022 as a quasi-experimental shock. By interacting this sharp timing with pre-determined firm-level occupational exposure to LLMs, we estimate the causal effect of generative AI on the wage structure.

We measure firm-level exposure by combining the pre-shock occupational mix of each firm with the LLM-exposure scores derived by Eloundou et al. (2024). Treating this structure as predetermined, we classify firms into exposure quintiles and implement a dynamic difference-in-differences design to estimate the causal impact of the ChatGPT shock. We validate our Revelio-based wage measures by replicating the main analysis using administrative data from the BLS Quarterly Census of Employment and Wages (QCEW); across all sources, we document consistent patterns of exposure-induced wage declines.

We document three main findings. First, generative AI exposure has caused a sharp decline in starting wages of new positions. Relative to unexposed firms, new position wages fall on average 4.5% and up to 7.7% in the most exposed quintile, following the release of ChatGPT. This aggregate decline is driven by the bottom and middle of the experience ladder: junior wages fall by 6.3% and mid-level wages by 5.9%, while senior wages remain stable or even increase. Event-study estimates confirm flat pre-trends and a persistent divergence immediately post-2022.

Second, LLM exposure fundamentally alters the firm's seniority structure. In exposed firms, the share of junior new positions falls by 4% in 2024, while the share of mid-level new positions rises by approximately the same amount. Interestingly, this employment expansion at the mid-level occurs alongside falling wages– a pattern that implies that the labor supply effect dominates the demand shift. We interpret this as evidence of rapid reallocation: as AI automates entry-level tasks, the effective supply of mid-level workers expands, driven by accelerated promotions and displaced juniors seeking mid-level roles. This supply expansion for mid-level positions is more rapid than firm demand, exerting downward pressure on wages even as mid-level headcount grows.

Third, we show that this reorganization of wages and seniority is accompanied by a systematic decredentialisation of the bottom and middle of the hierarchy. In exposed firms, the average education of junior and especially mid-level new positions declines over time, while education requirements for senior roles remain broadly unchanged. This suggests that firms respond not only by shifting hiring toward mid-level titles, but also by redefining what those titles mean.

Our findings provide large-scale causal evidence consistent with recent theoretical work on AI in the knowledge economy, which models firms as knowledge hierarchies with scarce expert time and tacit know-how (see, e.g., Garicano (2000); Garicano and Rossi-Hansberg (2006, 2015); Fuchs et al. (2015)). Ide and Talamàs (2025) propose a model of "Basic Autonomous AI" where AI agents substitute for entry-level knowledge tasks. While this technology alleviates knowledge bottlenecks, their theory highlights a distinct general equilibrium effect: by competing with humans in routine cognitive work, AI displaces junior workers "upward" into problem-solving roles for which they are less prepared, or forces them to compete with low-cost computational substitutes. Our findings provide large-scale causal evidence consistent with this theoretical prediction. In particular, we document that junior wages of AI-exposed firms fall, while senior wages increase. We also show a decline in new junior positions and an increase in new mid-level and senior positions.

The remainder of the paper is organized as follows. Section 2 describes the data and the construction of the AI exposure measure. Section 3 presents the difference-in-differences design. Section 4 presents the main results on wages and hiring composition. Section 5 concludes.

2 Data and Measurement

Our primary data source is the Revelio Technologies dataset, which provides anonymized longitudinal records of employment spells for 138 million U.S. workers starting in 1950. The data contains 111 million unique users with wages since 2014, and 78 million unique users with wages in the balanced panel of firms that we use for the main diff-in-diff analysis. Unlike studies relying on vacancy postings, which capture labor demand intent, this dataset allows us to observe realized worker-firm matches and their associated compensation. The data are reconstructed from online professional profiles and augmented with wage information derived from publicly available labor certification applications (Labor Condition Applications), self-reported salary platforms, and job postings. For each employment spell, we observe the key components of the labor contract: the employer, job title, location, start and end dates, and starting salary.

To characterize the firm's internal hierarchy, we utilize Revelio's standardized seven-level experience classification derived from job titles. We aggregate these into three broad tiers: Juniors (levels 1–2), Mid-level (levels 3–4), and Seniors (levels 5–7). We define the seniority of a match based on the level at the time of the creation of the new position. This distinction allows us to isolate the impact of AI on starting wages for new recruits separately from the internal promotion dynamics of incumbent workers.

We define our main outcome as the log starting real wage for newly created positions. For each new position, we convert the starting compensation to a full-time annualized amount and deflate it using the CPI-U. We then compute firm–year averages separately for junior, mid-level, and senior new positions. While Revelio draws wage data directly from 200 million observed wages from job postings, publicly available labor certification applications, and self-reported wage platforms where available, and it imputes the remainder using a predictive model.

To measure exposure to generative AI, we employ the occupation-level automation index from Eloundou et al. (2024), which quantifies the technical feasibility of using GPT-4 to substitute for or significantly accelerate specific work activities. In our Revelio sample, this index captures the distinct “cognitive” nature of the shock: the most heavily exposed occupations include mathematicians, blockchain engineers, and writers—roles previously considered insulated from automation. Conversely, manual service roles (e.g., cooks, drivers) and outdoor occupations appear at the bottom of the distribution. This variation allows us to test whether the arrival of ChatGPT disproportionately affected the price of cognitive labor.

We construct firm-level exposure by computing an employment-weighted average of these occupational scores based on the firm's workforce composition immediately prior to the shock. Specifically, for each firm f, we define exposure as:

$$\text{Auto}_f = \sum_o s_{f,o}^{2022} \cdot \beta_o \tag{2.1}$$

where $s_{f,o}^{2022}$ denotes the share of firm f's employment in occupation o as of October 31, 2022 (we define occupations based on Revelio's reported 8-digit O*NET code for each position in the data). By fixing the weights at this pre-ChatGPT baseline, we treat exposure as a time-invariant, predetermined firm characteristic that does not mechanically respond to any subsequent endogenous reallocation of labor.

Our difference-in-differences design relies on a discretized version of this AI exposure measure: we classify firms into exposure quintiles and define the bottom quintile (least exposed) as the control group. This approach allows us to estimate the causal impact of AI adoption non-parametrically, avoiding functional form assumptions about how exposure scales into wage effects.

We aggregate the worker-level records into a balanced annual panel of firms from 2014 to 2025. For each firm-year, we calculate total employment and the composition of new positions and employment stocks across the three seniority tiers (Junior, Mid-level, Senior). Our primary wage outcome is the mean log starting real wage for new positions, constructed separately for each seniority group. To ensure our results are not driven by small-sample noise, we weigh all regressions by pre-shock (2014) firm employment.

2.1 Wage Trends by AI Exposure Quintile

Before conducting the causal econometric analysis, we present descriptive evidence illustrating the pre-existing trends in wages across occupations sorted by their estimated exposure to Generative AI. We divide all worker-firm matches into five quintiles (Q1 to Q5) based on the occupational AI exposure measure, where Q5 represents the highest exposure. Figure 1 plots the evolution of the average real wage for new positions in these five quintiles, indexed to 1 in October 2022. Average log real wages evolved almost identically across exposure quintiles before the diffusion of LLMs. After the introduction of ChatGPT, there was a fast reduction in wages for the high-exposure quintiles, whereas low-exposure firms' wages ramain stable and start growing at a higher rate. This pattern is consistent with our DiD results: exposure to LLM automation is associated with a relative flattening of wage growth in highly exposed firms rather than a general macro shock hitting all firms equally.

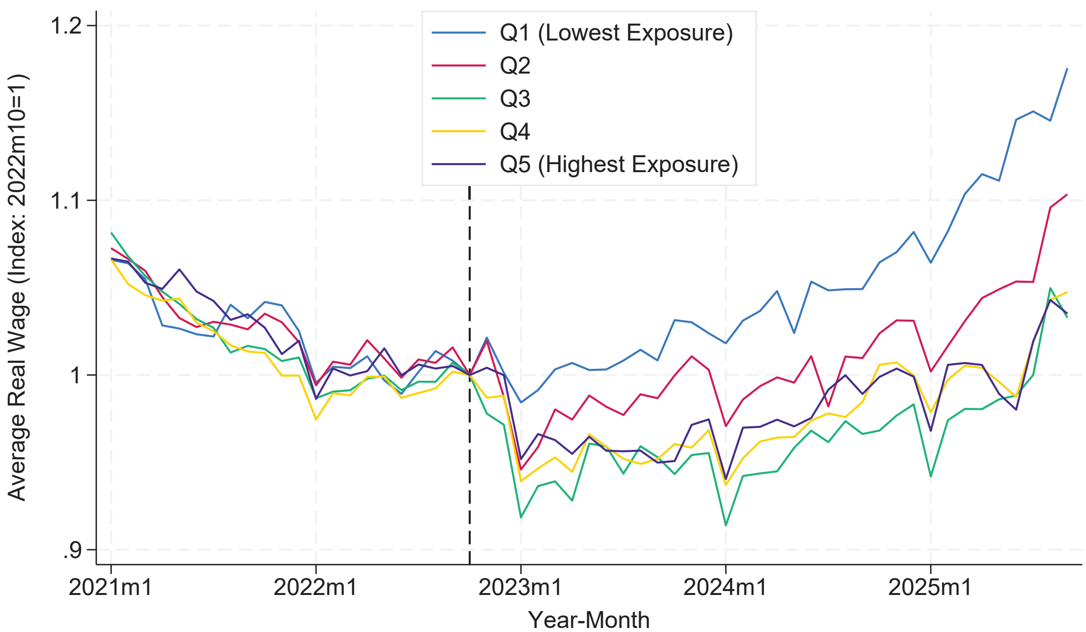

Figure 1. Average Real Wage of New Positions by Occupational AI Exposure Quintile (2022=1).
Notes: The figure plots the average real monthly starting wage for new positions, indexed to 1 in October 2022. Occupations are grouped into quintiles (Q1 to Q5) based on the estimated Generative AI exposure score, with Q5 representing the highest exposure.

3 Empirical Strategy

3.1 Difference-in-Differences Design

To isolate the causal effect of generative AI on labor market outcomes, we employ a dynamic difference-in-differences design around the public release of ChatGPT. We interpret November 30, 2022, as the onset of the shock ($t = 2022$) and define treatment based on the firm-level exposure index constructed in

Section 2. Utilizing those exposure quintiles, we define a binary treatment indicator $D_f = 1$ for firms in the top four quintiles (moderate to high exposure) and $D_f = 0$ for firms in the bottom quintile (least exposed), which serves as the comparison group. Because the exposure measure relies on the occupational composition fixed in October 2022, it is predetermined with respect to subsequent wage dynamics.

3.2 Estimation Framework

We estimate dynamic treatment effects using the estimator proposed by Callaway and Sant'Anna (2021). This framework is preferred over static Two-Way Fixed Effects (TWFE) as it is robust to treatment effect heterogeneity and strictly prohibits the use of earlier-treated units as controls for later units. Since the shock occurs simultaneously for the treated cohort in late 2022, the estimator simplifies to comparing the evolution of outcomes in treated firms relative to the control group (bottom quintile).

Formally, we estimate the employment-weighted Average Treatment Effect on the Treated (ATT) for each event-year k relative to the base year 2021 ($k = -1$). To ensure the comparability of the treatment and control groups, we estimate propensity scores using 2014 log employment and 3-digit NAICS industry fixed effects. We then use inverse-probability weighting to align the distribution of these baseline characteristics. Standard errors are clustered at the 2-digit NAICS industry level to account for correlation in shocks within sectors.

3.3 Identification

The validity of our design rests on the parallel trends assumption: that absent the release of ChatGPT, wages and hiring in exposed firms would have evolved similarly to those in unexposed firms. Two features of our design support this identifying assumption and limit the risk that our AI exposure measure is merely capturing other pre-existing differences in firms wage trends.

First, our exposure measure is predetermined, ensuring treatment is independent of ex-post firing decisions. Second, our inverse-probability weighting approach includes 3-digit NAICS industry fixed effects in the propensity-score model. This ensures that treated and control firms are balanced within industries, so that comparisons are not driven by differences between broad sectors such as Tech and Retail. By aligning the distribution of industry membership across groups, the weighting scheme helps us attribute differences in outcomes to exposure to generative AI rather than to sector-wide shocks. We provide direct evidence for this assumption in Section 4, documenting flat pre-trends across all event-study specifications.

3.4 Heterogeneity and Mechanisms

We extend this framework to characterize both heterogeneity in the wage response and the mechanisms of adjustment. First, we estimate the dynamic effects separately by exposure quintile to document a dose–response relationship between AI exposure and wage compression. Second, we study heterogeneity by firm size and by sector–distinguishing manufacturing from services and, within services, high- and low-exposure industries–to separate the specific GenAI shock from broader sectoral or size-specific fluctuations. Third, we estimate the model separately by worker seniority (Junior, Mid-level, Senior) and relate these patterns to within-firm reallocation by analyzing changes in the shares of junior, mid-level, and senior new positions and internal worker flows. Finally, we examine heterogeneity in the education of new positions along the hierarchy, which allows us to study how AI adoption decredentialises and effectively redefines seniority levels within the firm.

4 Results

4.1 Average Impact on New Position Wages

To begin our empirical analysis, we first examine the average treatment effect on the treated (ATT) of firm-level Generative AI (GenAI) exposure on the starting log wages of all new positions. Our Difference-in-Differences (DiD) framework compares the wages of workers in new positions into high-exposure firms (treatment group) relative to those new positions in low-exposure firms (control group), following the diffusion of GenAI technology.

Figure 2 presents the dynamic event-study estimates of the ATET (Average Treatment Effect on the Treated) for the period spanning 2015 to 2025. The estimated coefficients in the pre-treatment period are clustered tightly around zero, providing strong visual evidence supporting the parallel trends assumption. This confirms that the baseline trends in starting wages were similar between the high- and low-exposure firms before the technological shock.

Following the onset of the GenAI shock, the trend sharply breaks. The estimated coefficients for the post-treatment period are consistently negative and statistically significant. The overall effect demonstrates a substantial and immediate drop in the average starting log wage in new positions in firms highly exposed to GenAI. Specifically, the effect in the period following the shock implies that the average starting salary for new positions in exposed firms fell by approximately 5% relative to the control group.

This initial finding shows that the arrival of GenAI was immediately capitalized into the labor market via the price of labor for new firm-worker matches. The negative effect suggests that, on average, the automation potential of GenAI currently dominates any potential productivity-induced increases in demand or wages for the tasks performed by the marginal new position. This initial look sets the stage for the more detailed heterogeneity analysis in the following subsections, where we investigate which types of firms and which seniority levels drive this aggregate wage compression.

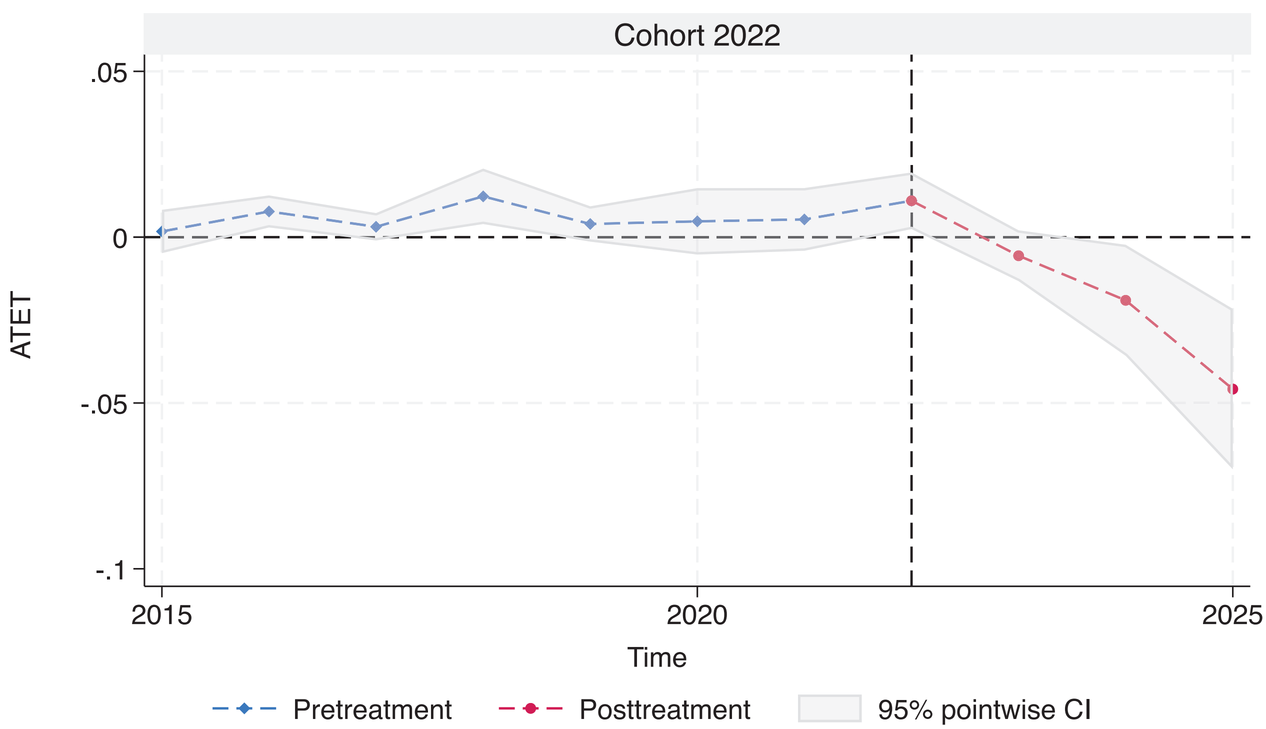

Figure 2. Dynamic ATT of Firm-Level AI Exposure on the Average Log Starting Wage of New Positions. *Notes:* The figure plots the estimated coefficients from the event-study regression of the log starting wage of new positions on the firm's AI exposure, controlling for worker, firm, and time fixed effects. The coefficients represent the difference in average log starting wages between high-exposure and low-exposure firms in a given quarter, relative to the pre-treatment period. The vertical line marks the introduction of ChatGPT in November 2022.

4.2 Heterogeneity by Exposure Intensity (Dose-Response)

To verify that the wage decline scales with the intensity of AI exposure, we estimate the dynamic treatment effects separately for each exposure quintile (Q2 through Q5), relative to the least exposed quintile (Q1).

Figure 3 displays the results. We observe a clear monotonic relationship in the magnitude of the post-treatment effect. Firms in the second quintile (Q2, Panel A) show a negligible or small negative effect. The magnitude of the wage decline increases progressively for Q3 (Panel B) and Q4 (Panel C) with a relative fall of approximately 5%, with the largest and most significant drop for the most exposed firms in Q5 (Panel D), that shows a relative fall in new position wages of approximately 8%. This dose-response relationship strengthens the causal interpretation of our results, indicating that higher concentrations of automatable cognitive tasks lead to sharper wage compression relative to unexposed firms.

4.3 Heterogeneity by Firm Size

A key question regarding the overall wage decline for exposed sectors is whether the effect is concentrated solely in large, technologically advanced firms (which might include AI developers or early adopters with massive R&D budgets) or if it represents a more general shock that has diffused throughout the economy. To address this, we split our sample into two groups based on firm size. Firms with less than 50 employees in 2014 are classified as "Small", whilst firms with more than 50 employees in 2014 are classified as "Large".

Figure 4 presents the dynamic event-study results for the starting log wages of new positions across these two subsamples. In both panels, the pre-treatment coefficients maintain the parallel trends assumption across the full sample and have a statistically significant negative treatment effect in both small (Panel A) and large (Panel B) firms.

In Panel A, the small firms, which are less likely to be producers of the technology but are active consumers, experience a negative wage effect that is comparable in magnitude to the aggregate finding reported in Section 4.1. This indicates that access to and utilization of GenAI tools (e.g., via commercial subscriptions) quickly diffused the automation potential across firms of all sizes.

The finding that the wage decline is not confined to the largest corporate players is highly significant.

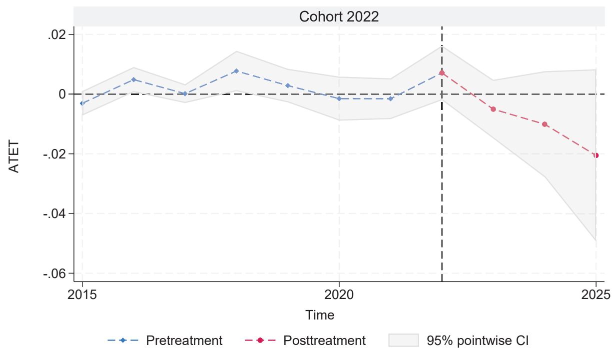

(a) Quintile 2 (Low-Mid Exposure)

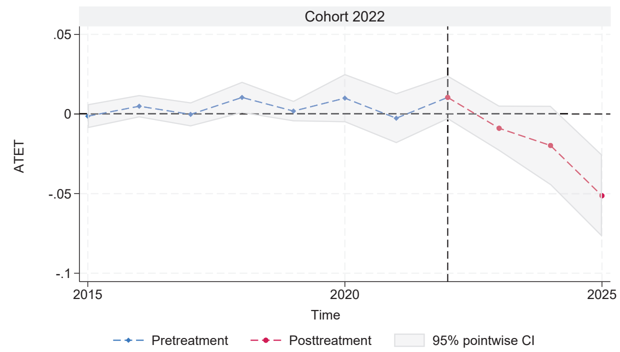

(b) Quintile 3 (Mid Exposure)

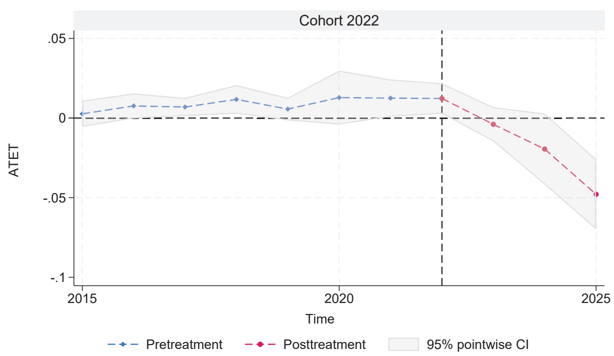

(c) Quintile 4 (Mid-High Exposure)

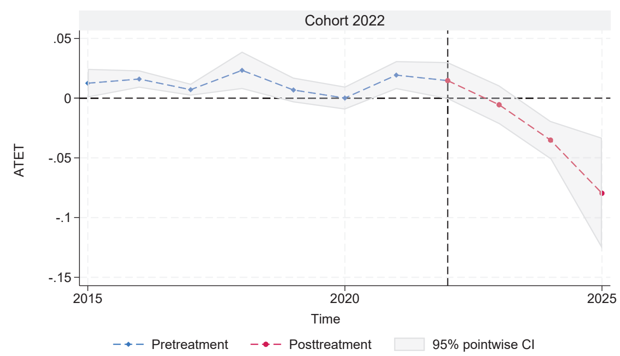

(d) Quintile 5 (Highest Exposure)

Figure 3. Dose-Response: Wage Effects by Exposure Quintile.
Notes: The figure plots the dynamic ATT of firm-level AI exposure on starting wages separately for exposure quintiles 2, 3, 4, and 5. The reference group is Quintile 1 (Lowest Exposure).

It supports the interpretation of GenAI as a General Purpose Technology (GPT) shock that rapidly alters the production function for cognitive tasks across the entire size distribution of the knowledge economy. The universality of this effect confirms that the aggregate wage compression is a robust phenomenon, not a feature exclusive to the tech giants who pioneered the change.

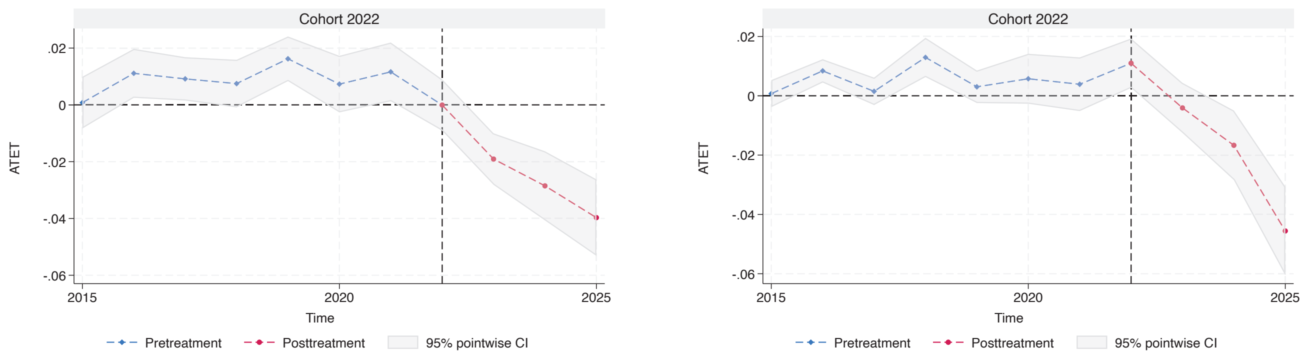

(a) Small Firms (< 50 Employees)

(b) Large Firms (> 50 Employees)

Figure 4. Dynamic ATT of Firm-Level AI Exposure on the Average Log Starting Wage of New Positions, by Firm Size.

Notes: The figure plots the estimated event-study coefficients for the log starting wage of new positions. Panel A reports results for firms in the bottom half of the employment distribution (Small); Panel B reports results for firms in the top half (Large).

4.4 Heterogeneity by Sector: Manufacturing vs. Services and Granular Analysis

4.4.1 Manufacturing vs. Services

To estimate the heterogeneity by sectors we first divide firms into two broad categories: the Manufacturing sector (NAICS 31-33) and the Services sector (all non-manufacturing industries, including Information, Finance, and Professional and Business Services).

Figure 5 presents the dynamic event–study estimates for these two groups. In both the Manufacturing sector (Panel A) and the Services sector (Panel B), starting wages in exposed firms fall after 2022 by a very similar amount. The manufacturing estimates are somewhat noisier and less precisely estimated, while the services coefficients are tighter and closely mirror the aggregate pattern. Given that services in our data concentrate a larger share of cognitive, knowledge-based tasks, indicating that wage compression following the arrival of ChatGPT does not arise only in a narrow set of service industries.

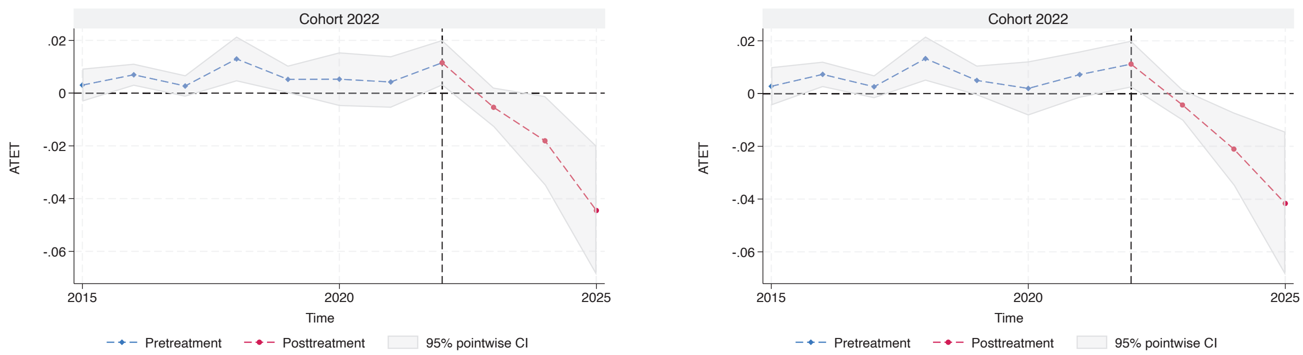

(a) Manufacturing **(b)** Services

Figure 5. Dynamic ATT of Firm-Level AI Exposure on the Average Log Starting Wage of New Positions, by Sector.
Notes: The figure plots the estimated event-study coefficients for the log starting wage of new positions. Panel A reports results for firms in the Manufacturing sector (NAICS 31-33); Panel B reports results for firms in Services industries (all non-manufacturing NAICS codes).

4.4.2 High- and Low-Exposure Granular Sectors

To refine this analysis, we examine a set of granular sectors, comparing those known to be highly exposed to GenAI with those with low exposure.

High-Exposure Services Sectors Figure 6 shows that the wage effects are concentrated in the most cognitive-intensive service industries. Information (Panel A), which includes software, media, and data processing, displays the largest and most persistent relative decline in starting wages, with a fall of approximately 13%. Finance and Professional Services (Panel B) also experience a substantial negative effect of approximately 4%, consistent with the idea that LLMs directly substitute for standardized analysis, documentation, and reporting tasks in these sectors.

Low-Exposure Sectors Conversely, Figure 7 shows that sectors with lower inherent AI exposure exhibit no systematic wage response. Education (Panel A) and Health (Panel B) both display coefficients clustered around zero in the post-treatment period, with wide confidence intervals that generally include zero. This pattern confirms that in a more granular sectoral breakdown, some low-exposure industries such as education and health do not experience a wage response to AI exposure.

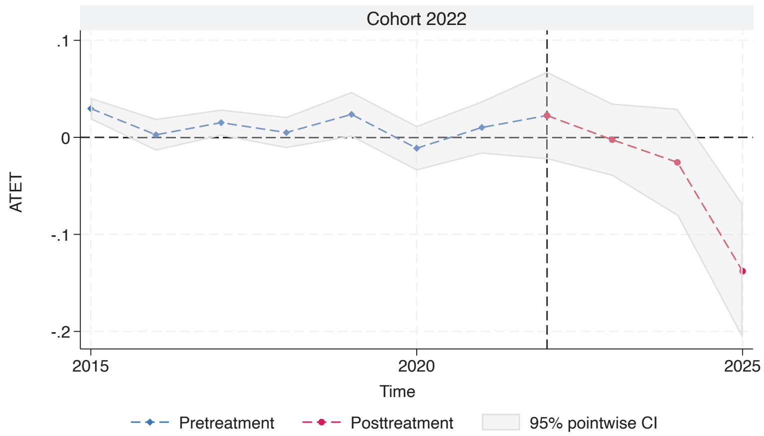

(a) Information

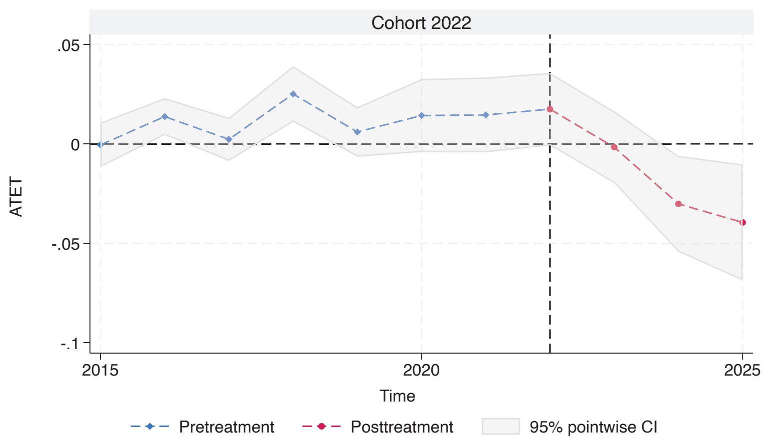

(b) Finance & Professional Services

Figure 6. Dynamic ATT on Average Log Starting Wage for High-Exposure Service Sectors.

Notes: The figure plots the estimated event-study coefficients for the log starting wage of new positions in high-exposure sectors. Panel A reports results for Information; Panel B reports results for Finance & Professional Services.

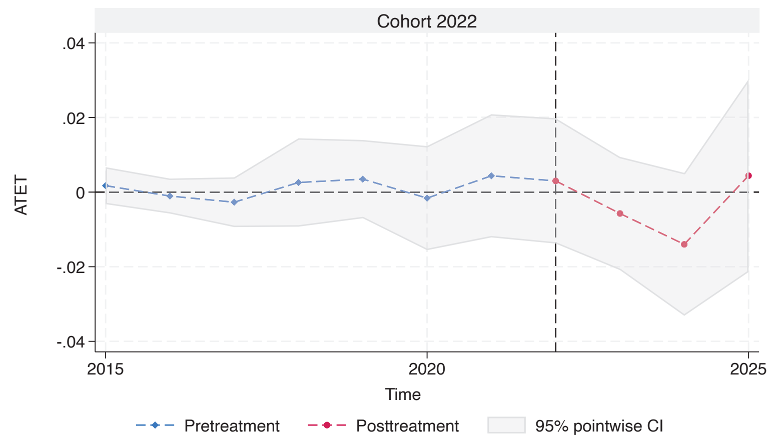

(a) Education

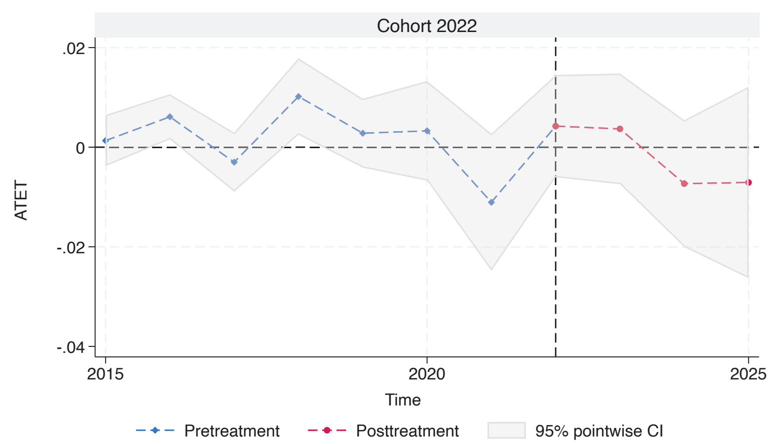

(b) Health

Figure 7. Dynamic ATT on Average Log Starting Wage for Low-Exposure Sectors.
Notes: The figure plots the estimated event-study coefficients for the log starting wage of new positions in low-exposure sectors. Panel A reports results for Education; Panel B reports results for Health. Coefficients remain close to zero and are generally statistically indistinguishable from zero in the post-treatment period.

4.5 Heterogeneity by Worker Seniority

We divide new positions into three seniority levels: Junior (entry-level), Mid-Level (experienced workers), and Senior (management/executive roles).

Figure 8 clearly illustrates this effect. For Juniors (Panel A), the decline in starting wages is rapid, substantial, and statistically significant, with a drop estimated at 6.3%. This confirms that entry-level roles absorb a large part of the aggregate wage shock. The effect on Mid-Level new positions (Panel B) is also negative and significant, with a decrease of 5.9%. In contrast, the starting wages for new positions of Seniors (Panel C) remain stable and statistically insignificant, suggesting that senior, less routinized, and supervisory roles are largely shielded from this initial wage compression. This heterogeneity is consistent with GenAI automating tasks heavily concentrated at the start of the career ladder.

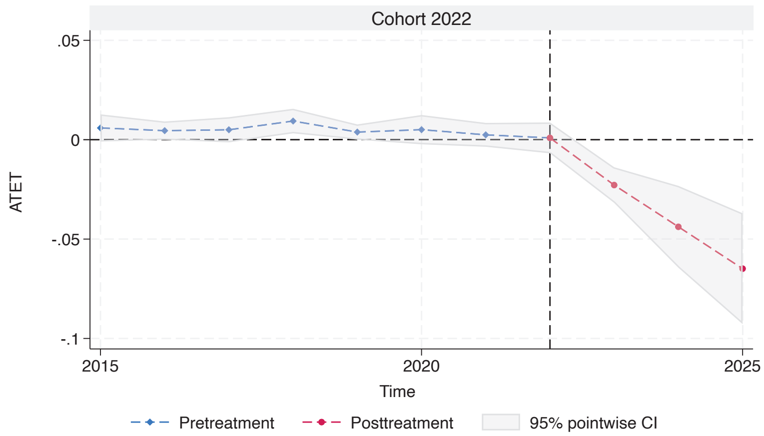

(a) Junior

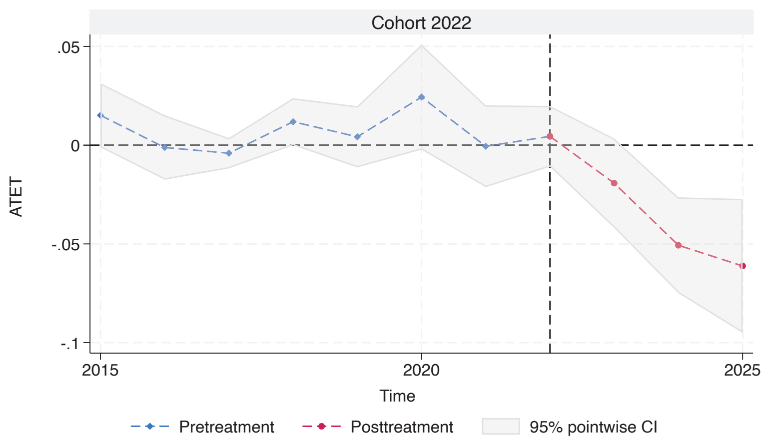

(b) Mid-Level

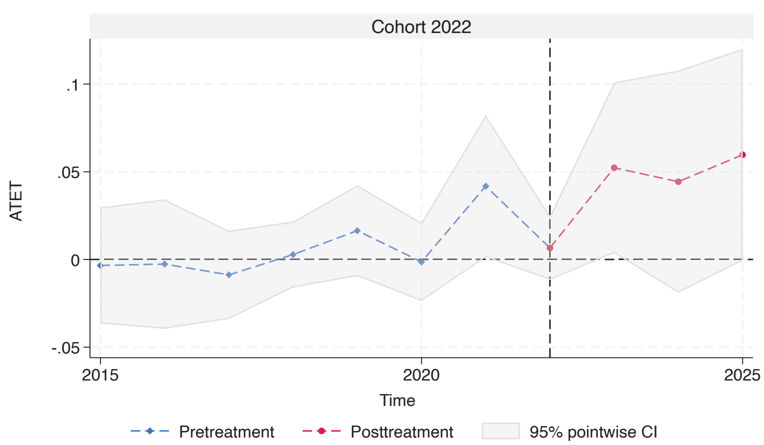

(c) Senior

Figure 8. Dynamic ATT of Firm-Level AI Exposure on the Average Log Starting Wage of New Positions, by Seniority Level.

Notes: The figure plots the estimated event-study coefficients for the log starting wage of new positions, segregated by worker seniority.

4.6 Shares and Reallocation Mechanism

4.7 Heterogeneity by Seniority and Within-Firm Adjustment

The wage effects by seniority show that AI exposure does not generate a uniform compression across the hierarchy, with juniors and mid-level workers being more affected than seniors. To understand whether this differential impact is accompanied by changes in the structure of jobs, we analyze the shares of junior, mid-level, and senior positions and worker flows to assess how firms reorganize their hierarchies in response to AI adoption.

4.7.1 Composition of New Positions

Figure 9 shows the dynamic ATT on the share of new positions by seniority. The Junior New Position Share (Panel A) shows a drastic negative change, confirming that exposed firms are reducing the share of new junior positions relative to less exposed firms. Conversely, the Mid-Level New Position Share (Panel B) increases by an amount very similar in magnitude to the decline in the junior share, indicating that a bulk of new labor demand is being redirected toward the mid-level. The Senior New Position Share (Panel C) remains positive but much smaller, consistent with a largely unchanged role for senior hires. Taken together, these patterns hint at a substitution away from junior labor toward mid-level labor in the hiring and promotion process.

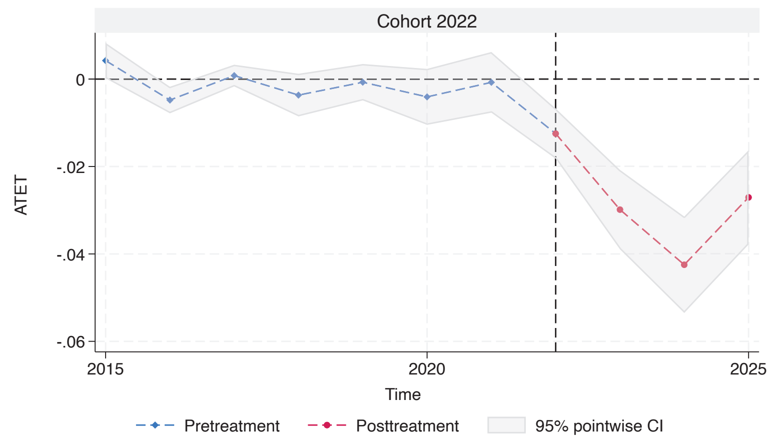

(a) Junior New Position Share

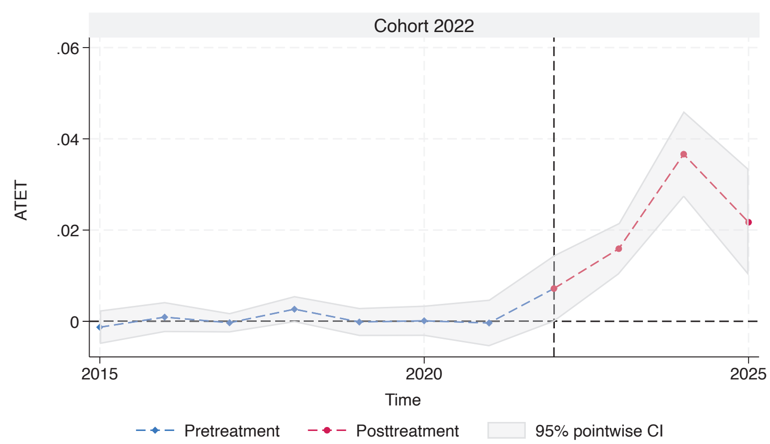

(b) Mid-Level New Position Share

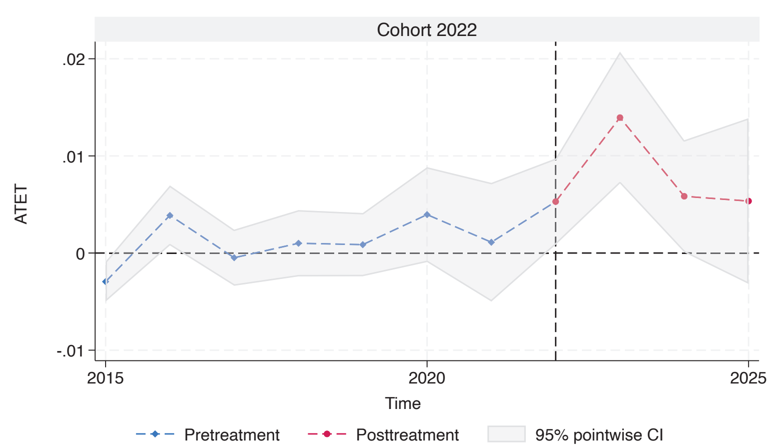

(c) Senior New Position Share

Figure 9. Dynamic ATT of Firm-Level AI Exposure on the Share of New Positions, by Seniority.
Notes: The figure plots the estimated event-study coefficients for the share of new positions within each seniority level.

4.7.2 Stock Composition

Figure 10 examines the long-run changes in the firm's labor stock composition. The Junior Stock Share (Panel A) shows a sustained negative decline, meaning the stock of junior workers relative to the total workforce is shrinking. The Mid-Level Stock Share (Panel B) increases, reflecting the sustained creation of new positions into this band. The Senior Stock Share (Panel C) remains stable, indicating that the core structure of senior leadership is largely preserved, while the composition of the early and intermediate career levels shifts significantly.

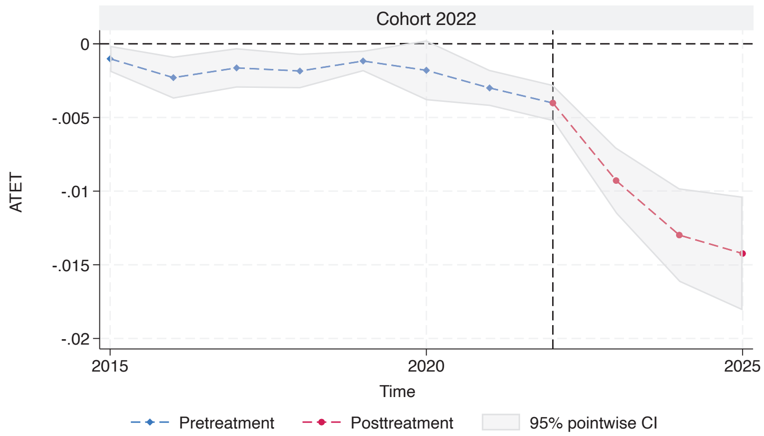

(a) Junior Stock Share

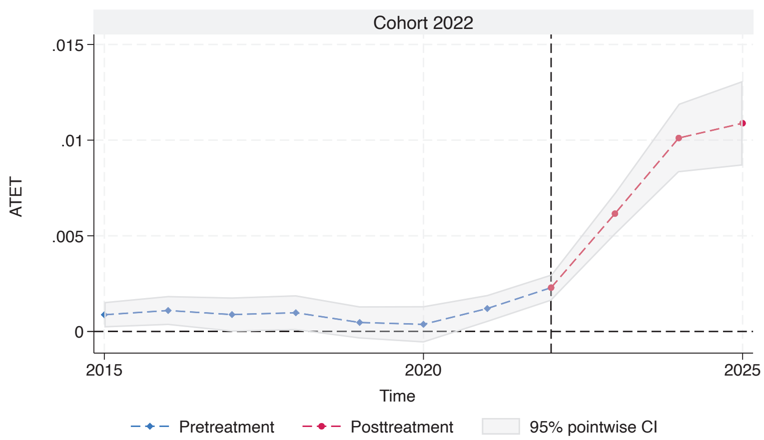

(b) Mid-Level Stock Share

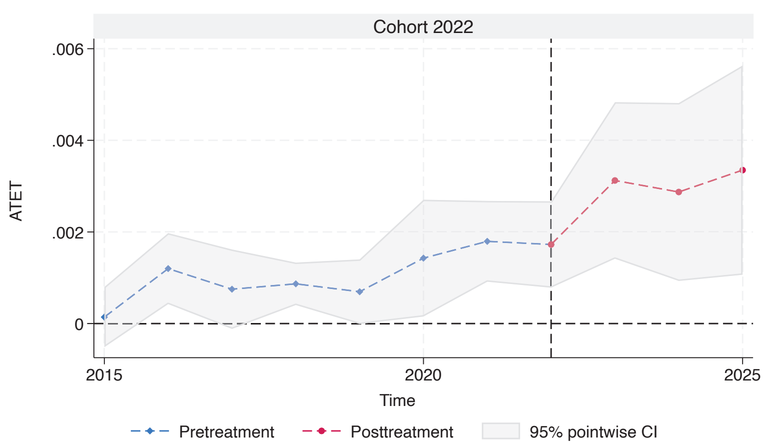

(c) Senior Stock Share

Figure 10. Dynamic ATT of Firm-Level AI Exposure on the Stock Share of Employment, by Seniority.
Notes: The figure plots the estimated event-study coefficients for the share of the firm's total employment stock within each seniority level.

4.7.3 Career Transitions

We next investigate the impact on promotions. Figure 11 reports the dynamic treatment effects on the transitioning between seniority levels.

Panel A tracks the share of transition from Junior to Mid-Level (J $\rightarrow$ M). While the point estimates exhibit an increase in 2024, the effect returns to zero by 2025. The increase in 2023 and 2024, together with flat pre-trends, is consistent with the reallocation effects shown previously.

Panel B shows that the share of transition from Mid-Level to Senior (M $\rightarrow$ S) is estimated as a null effect. The coefficients remain close to zero throughout the post-treatment window, confirming that the upper seniority level are unaffected from the shock.

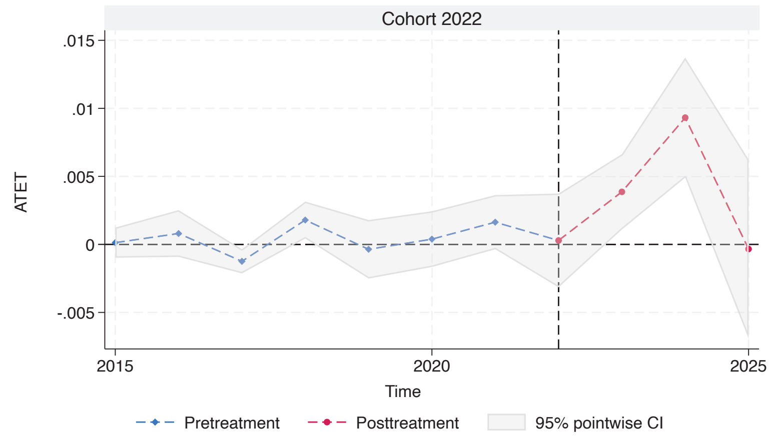

(a) Junior to Mid-Level share of Transition

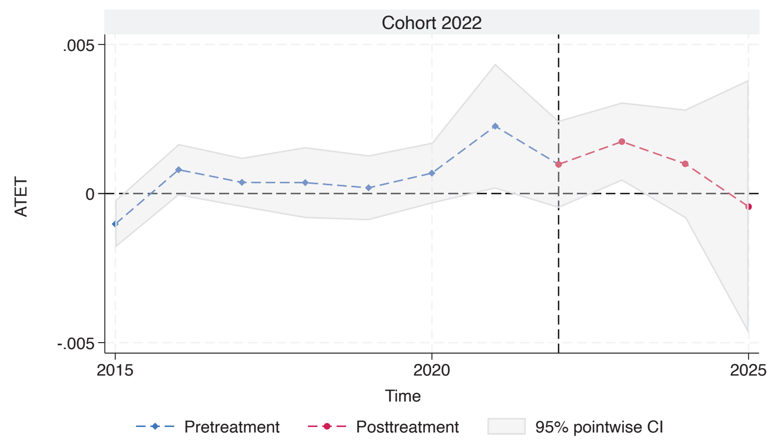

(b) Mid-Level to Senior Share of Transition

Figure 11. Dynamic ATT of Firm-Level AI Exposure on Transition.
Notes: The figure plots the estimated event-study coefficients for the share of career transitions between seniority levels.

4.8 Effect on Seniority Level Definitions

We first look at the average college graduate share of new positions, pooling all seniority levels. Figure 12 reports the event-study estimates. Pre-trends are close to zero, whereas after the arrival of ChatGPT the average college graduate share of new positions in AI-exposed firms declines steadily, indicating a gradual overall de-credentialisation of new positions of workers.

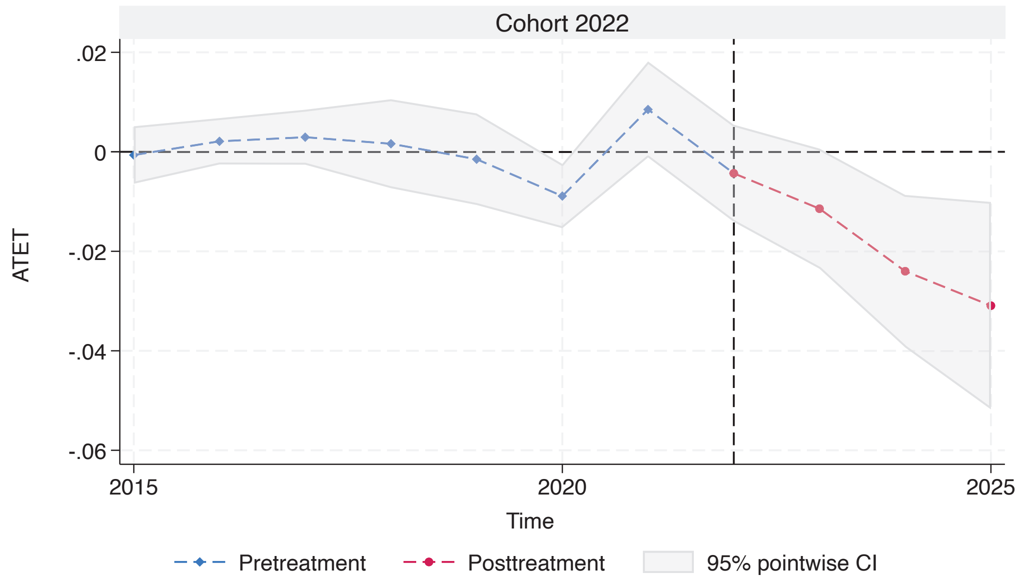

Figure 12. Effect of AI exposure on the college graduate share of new positions (overall).

Finally, we study how AI exposure affects the education of new positions along the hierarchy. Figure 13 shows estimates for junior, mid-level, and senior positions. In AI-exposed firms, the average education of junior and especially mid-level new positions declines over time relative to less-exposed firms, while the education of senior new positions remains broadly unchanged. This indicates a decredentialisation at the lower and middle rungs, but not at the top of the hierarchy.

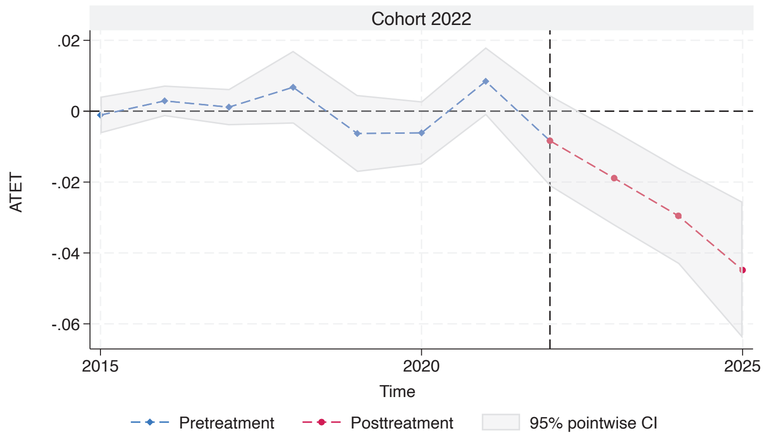

(a) Juniors

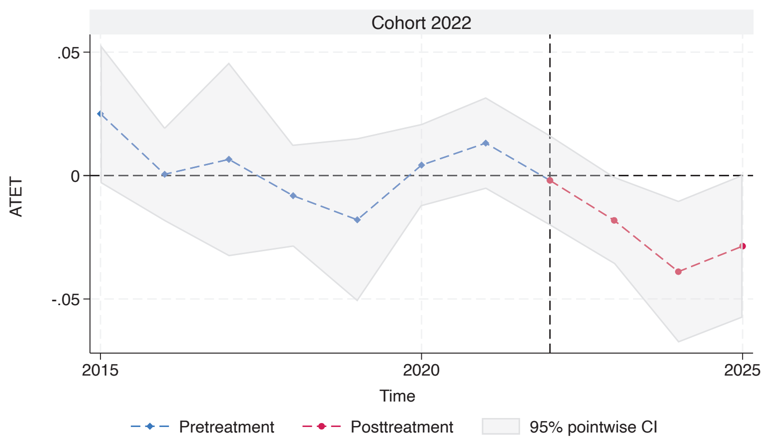

(b) Mid-level

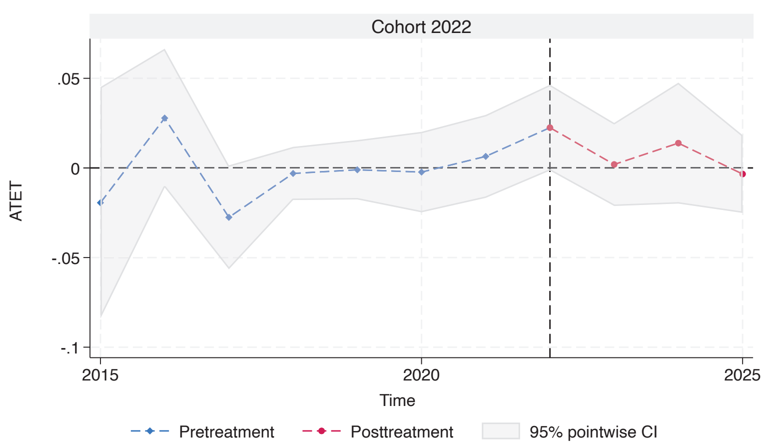

(c) Seniors

Figure 13. Effects of AI exposure on the college graduate share of new positions by seniority level.

These patterns are consistent with a knowledge-hierarchy view of the firm.[1] In a Garicano and Rossi-Hansberg (2006) framework, production is organized as a ladder of increasingly difficult problems: workers at the bottom handle relatively routine tasks and escalate unsolved cases to more knowledgeable managers and, ultimately, to a small group of top experts. Formal education is one way to acquire the knowledge needed to solve problems without escalation. When a basic autonomous AI arrives that can perform a large share of the routine cognitive tasks previously handled by humans, it effectively injects "cheap knowledge" at the lower rungs of the hierarchy. The marginal value of additional formal education at the junior and mid-level layers falls: an averagely educated worker armed with AI can now do much of what a more educated worker used to do. Our results indicate that firms respond not only by changing who they hire or promote, but also by redefining what it means to be junior or mid-level – expanding mid-level positions populated by less educated workers that would have been assigned to junior work pre-AI.

In the language of Ide and Talamàs (2025), our setting corresponds to the case of *basic autonomous AI* that substitutes for the tasks of less knowledgeable agents but not for the hardest, most tacit parts of production.[2] Their model predicts that such AI makes intermediate and lower types more substitutable while preserving the scarcity of top expertise. Our education results map directly into this asymmetry: firms appear to de-credentialise and relabel the bottom and middle of the hierarchy, filling expanded "mid-level" roles with less educated workers, while keeping education requirements for genuinely senior positions broadly unchanged. In other words, AI adoption compresses the returns to formal education where AI can stand in for codified knowledge and allows firms to stretch titles downward, but leaves intact – or even reinforces – the role of education at the top of the hierarchy, where human judgment and tacit knowledge remain critical.

5 Conclusion

The rapid diffusion of generative AI has reopened the debate on technology and labor, shifting the focus from the displacement of routine manual work to the automation of cognitive tasks. In this paper, we provide large-scale causal evidence on how this shock is reshaping the internal wage structure of

[1] See Garicano and Rossi-Hansberg (2006) and Caliendo et al. (2015), who show empirically that firms restructure layers effectively redefining seniority levels as technology and scale evolve.

[2] See Ide and Talamàs (2025)

the firm. We document that the arrival of ChatGPT acted as a form of seniority-biased technological change: it severely eroded the price and quantity of entry-level labor while leaving senior roles largely untouched relative to workers in less exposed firms.

However, we show that this effect is not uniform. Wage compression is strongest in sectors where workers perform cognitive, knowledge-based tasks that are particularly exposed to Large Language Models. Furthermore, the adjustment occurs exclusively on the new positions margin. While exposed firms have sharply reduced starting wages for junior roles relative to less exposed firms, the progression through the early stages of the career ladder has accelerated. Incumbents and new hires are being promoted to mid-level roles at faster rates, effectively expanding the intermediate tier to absorb the increased productivity of AI-augmented staff.

Our findings suggest that firms are responding to the shock by restructuring their knowledge hierarchies. By automating routine information-processing tasks, AI effectively injects "cheap knowledge" at the bottom of the ladder. Firms appear to respond by redefining seniority: they expand mid-level roles to include less-credentialed workers performing AI-augmented tasks, while maintaining a premium for the tacit knowledge and judgment that characterize senior positions. This generates a supply shock that depresses wages for the middle tier, not because mid-level skills are obsolete, but because they are becoming less scarce.

The most profound implication of these results concerns the future of human capital formation. The traditional professional service firm operates as an apprenticeship system: junior workers accept lower initial productivity to acquire the knowledge required for senior roles. By automating the entry-level rung of this ladder, generative AI risks disrupting the pipeline through which workers acquire skills. If firms stop hiring apprentices today to increase short-run efficiency, they risk creating a broken rung in the career ladder, raising urgent questions about how the next generation of experts will be trained in an era of automated cognition.

References

Acemoglu, Daron and David Autor (2011) "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, 4B, 1043–1171.

Autor, David H., Frank Levy, and Richard J. Murnane (2003) "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118 (4), 1279–1333.

Brynjolfsson, Erik, Bharat Chandar, and Daniel Chen (2025) "Canaries in the Coal Mine? Six Facts about the Employment Effects of Artificial Intelligence," *Stanford Digital Economy Working Paper*.

Caliendo, Lorenzo, Ferdinando Monte, and Esteban Rossi-Hansberg (2015) "The Anatomy of French Production Hierarchies," *The Quarterly Journal of Economics*, 130 (4), 1655–1708.

Callaway, Brantly and Pedro H. C. Sant'Anna (2021) "Difference-in-Differences with Multiple Time Periods and Heterogeneous Treatment Effects," *Journal of Econometrics*, 225, 200–230, 10.1016/j.jeconom.2020.12.001.

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

Fuchs, William, Luis Garicano, and Luis Rayo (2015) "Optimal Contracting and the Organization of Knowledge," *Review of Economic Studies*, 82 (2), 632–658, 10.1093/restud/rdu033.

Garicano, Luis (2000) "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 108 (5), 874–904, 10.1086/317671.

Garicano, Luis and Esteban Rossi-Hansberg (2006) "Organization and Inequality in a Knowledge Economy," *The Quarterly Journal of Economics*, 121 (4), 1383–1435, 10.1093/qje/121.4.1383.

——— (2015) "Knowledge-Based Hierarchies: Using Organizations to Understand the Economy," *Annual Review of Economics*, 7, 1–30, 10.1146/annurev-economics-080614-115424.

Hosseini, Sara and Georg Lichtinger (2025) "Generative AI as Seniority-Biased Technological Change: Evidence from US Resume and Job Posting Data," *SSRN Working Paper*.

Ide, Enrique and Eduard Talamàs (2025) "Artificial Intelligence in the Knowledge Economy," *Journal of Political Economy*, forthcoming.

Katz, Lawrence F. and Kevin M. Murphy (1992) "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 107 (1), 35–78.

A Summary Statistics

Table 1. Summary Statistics: Balanced Firm-Year Panel (2014-2025)

Variable	Observations	Mean	Std. Dev.	Min	Max
Average Log Firm Wage					
Overall	3,038,256	9.98	0.37	8.51	12.36
Junior	2,909,268	9.78	0.30	8.49	12.13
Mid	2,150,600	10.22	0.30	8.24	12.19
Senior	1,850,952	10.81	0.40	8.60	13.21
Employment Flow (Share of New Positions)					
Junior	3,038,256	0.683	0.263	0.00	1.00
Mid	3,038,256	0.184	0.198	0.00	1.00
Senior	3,038,256	0.132	0.181	0.00	1.00
Employment Stock (Share of Employees)					
Junior	3,038,256	0.599	0.202	0.00	1.00
Mid	3,038,256	0.201	0.126	0.00	1.00
Senior	3,038,256	0.200	0.152	0.00	1.00
Employment Level					
Overall	3,038,256	212.85	3,273.94	1.00	1,528,218
Junior	3,038,256	133.94	2,146.36	0.00	959,534
Mid	3,038,256	45.57	809.01	0.00	398,093
Senior	3,038,256	33.34	414.03	0.00	170,591
Eloundou et al. (2024) Automation Score					
Automation Score (β)	3,038,160	0.466	0.099	0.00	0.95

Notes: Wage variables are in logarithmic form. Share variables range from 0 to 1 representing proportions. Employment variables show stock measures of worker counts by seniority level. Sample restricted to balanced panel.

Table 2. Summary Statistics: Unbalanced Firm-Year Panel (2014-2025)

Variable	Observations	Mean	Std. Dev.	Min	Max
Average Log Firm Wage					
Overall	19,836,784	10.04	0.58	7.54	13.98
Junior	15,026,633	9.74	0.38	7.54	12.25
Mid	7,256,961	10.23	0.34	8.24	12.34
Senior	6,624,581	10.90	0.48	8.51	13.98
Employment Flow (Share of New Positions)					
Junior	19,836,784	0.632	0.412	0.00	1.00
Mid	19,836,784	0.187	0.321	0.00	1.00
Senior	19,836,784	0.181	0.333	0.00	1.00
Employment Stock (Share of Employees)					
Junior	53,251,215	0.475	0.414	0.00	1.00
Mid	53,251,215	0.226	0.340	0.00	1.00
Senior	53,251,215	0.299	0.388	0.00	1.00
Employment Level					
Overall	53,251,215	16.40	783.67	1.00	1,528,218
Junior	53,251,215	10.00	513.69	0.00	959,534
Mid	53,251,215	3.52	193.56	0.00	398,093
Senior	53,251,215	2.89	99.24	0.00	170,591
Eloundou et al. (2024) Automation Score					
Automation Score (β)	46,227,227	0.464	0.151	0.00	1.00

Notes: Wage variables are in logarithmic form. Share variables range from 0 to 1 representing proportions. Employment variables show stock measures of worker counts by seniority level.

B Data Validation

B.1 Construction of Wage Measures in Revelio

Revelio does not directly observe a wage for every employment spell. Instead, the wage variable used in our analysis combines *directly observed* salaries with *predicted* salaries obtained from a statistical imputation model.

According to Revelio's technical documentation, wages are constructed in two steps. First, whenever possible, Revelio assigns a wage based on direct information from external sources, including: (i) U.S. H-1B / Labor Condition Applications (LCA), which report the offered salary together with the employer, job title, and location; (ii) job postings that list salary ranges; and (iii) crowdsourced salary platforms such as Levels.fyi, where workers self-report their compensation. These observed salaries are converted to a full-time annualized measure under the assumption of a 40-hour work week and 52 working weeks per year.

Second, for employment spells without directly observed wages, Revelio estimates a predictive model that imputes a salary based on a rich set of covariates. The model uses information on job title and its standardized cluster in Revelio's job taxonomy, seniority level, employer identity, geographic location, years of tenure at the firm, and calendar year of observation. When a given firm lacks sufficient direct wage observations, the model borrows strength from "similar" firms in the same industry and location. The resulting predicted wage is stored alongside an estimated confidence interval (lower and upper bounds) and a time stamp indicating when the prediction was last updated. In addition, Revelio provides sampling weights designed to correct for differential coverage across occupations and locations so that aggregate wage statistics better reflect the underlying population.

This procedure implies that a non-trivial share of the wages in our worker–firm panel are model-based predictions rather than raw administrative reports. It also introduces several limitations noted by Revelio and in user discussions. First, all wages are expressed on a full-time-equivalent basis, which may overstate earnings for any truly part-time positions. Second, the underlying sources and coverage vary across geography and occupation, so some regions and niche roles are measured with more noise than others. Third, by construction, the imputation model smooths over idiosyncratic firm-level shocks, which can attenuate measured dispersion in wages.

B.2 Comparison with QCEW Administrative Data

Because of this imputation step, it is important to verify that the Revelio wage series captures genuine movements in labor costs rather than artifacts of the prediction model. To assess the reliability of the imputed wage data, we validate our aggregate trends against administrative benchmarks from the Bureau of Labor Statistics.

Figure A.1 compares the average real log starting wage of new hires in our Revelio sample against the real average weekly wage reported in the Quarterly Census of Employment and Wages (QCEW) for the private sector. While the two series measure related but distinct concepts—Revelio captures the marginal price of new hires, whereas QCEW captures the average earnings of the stock of employed workers—their temporal dynamics are strikingly similar. The series track closely during the pre-pandemic period and capture the same nominal acceleration and subsequent real adjustment post-2020. The correlation coefficient between the two annual series over the 2014–2024 period exceeds 0.9.

This tight comovement provides strong evidence that the Revelio imputation model successfully recovers genuine variation in U.S. labor costs rather than idiosyncratic noise. In particular, despite the fact that many individual wages are predicted rather than directly observed, the aggregate wage path in Revelio closely replicates the trajectory of administrative earnings in QCEW, which supports the use of the Revelio wage series in our difference-in-differences analysis.

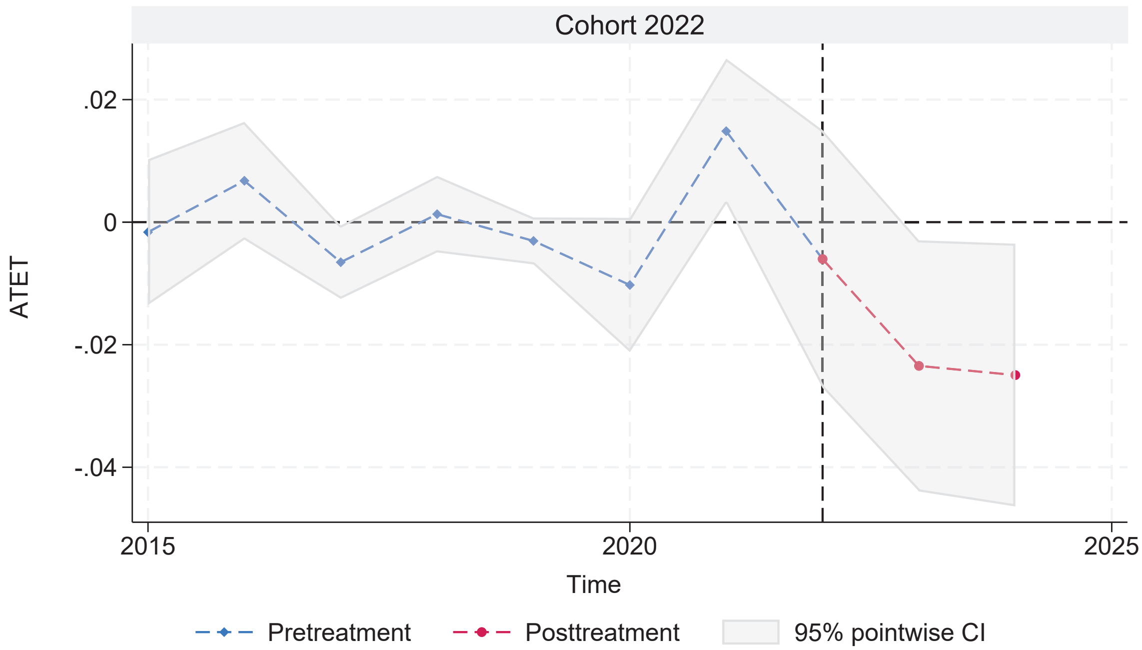

Figure A.1. Dynamic ATT of Firm-Level AI Exposure on the Aggregate Log Wage of New Hires Using QCEW Data.
Notes: The figure plots the estimated coefficients from the event-study regression of the log wage on the firm's AI exposure, controlling for worker, firm, and time fixed effects using QCEW data. The coefficients represent the difference in average log starting wages between high-exposure and low-exposure firms in a given quarter, relative to the pre-treatment period. The vertical line marks the introduction of ChatGPT in November 2022.

C Robustness: Unbalanced Panel

Our main analysis uses a balanced panel of firms to track within-firm adjustments over time. To ensure that our results are not driven by survivor bias or by the specific sample-selection criteria implied by the balanced panel, we replicate our main wage analysis using an **unbalanced panel** that includes all firms that appear in the data for at least part of the sample period.

Figure A.2 presents the results. Panel A shows the aggregate effect on starting wages of new positions, while Panels B, C, and D decompose the effect by seniority. The results are qualitatively identical and quantitatively very similar to the balanced-panel estimates reported in the main text. We observe a significant aggregate decline that is entirely driven by Junior (Panel B) and Mid-level (Panel C) hires, while Senior wages (Panel D) remain stable. This confirms that the seniority-biased wage erosion is a robust feature of the labor markets response to Generative AI, independent of firm entry and exit dynamics.

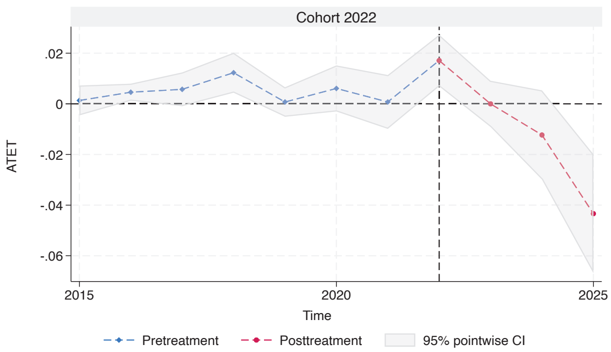

(a) Overall

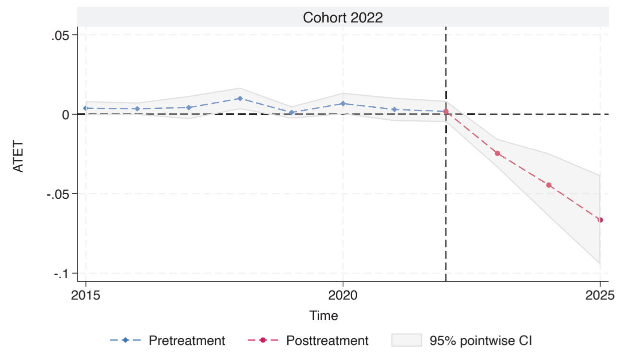

(b) Junior Hires

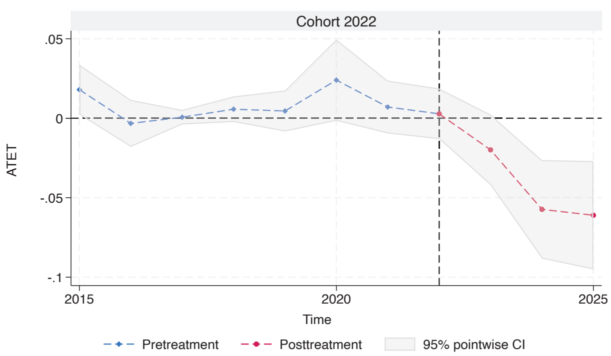

(c) Mid-level

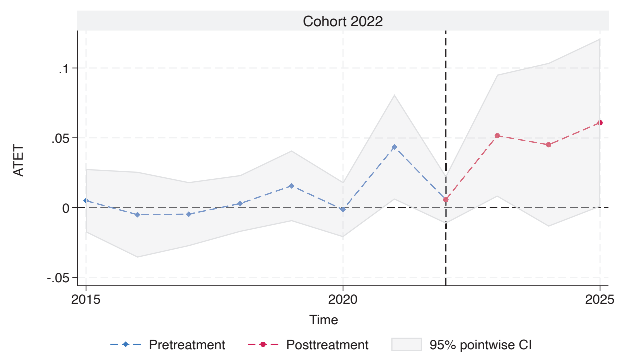

(d) Senior

Figure A.2. Robustness Check: Wage Effects in an Unbalanced Panel.
Notes: The figure replicates the main wage analysis using an unbalanced panel of firms. Panel A plots the aggregate effect; Panels B, C, and D plot the effect by seniority level.